



Generative AI Services 2025 RadarView™ Excerpt

The evolution of service providers
from AI enablers to ecosystem
orchestrators

November 2025






Table of contents

- About the Generative AI Services 2025 Market Insights report 3
- Executive summary 6
- State of the market 9
- Future outlook 16
- Infosys profile 18
- Appendix 21
- Key contacts 25


About the Generative AI Services 2025 Market Insights report




The *Generative AI Services 2025 Market Insights* is a companion report to the *Generative AI Services 2025 RadarView*. While the RadarView report concentrates on evaluating service providers, this report takes a broader perspective, offering a comprehensive overview of the current state of the generative AI (Gen AI) services market.



Amid growing interest in the Gen AI services space in recent years, this report aims to inform key stakeholders, including enterprises, service providers, and Gen AI leaders, about the major trends shaping this dynamic landscape. It explores the primary growth drivers, evolving strategic priorities, and the broader transformation of the Gen AI ecosystem.



This report was created using detailed data collected through our ongoing market tracking and surveys, public disclosures such as US Securities and Exchange Commission filings, annual reports, and quarterly earnings calls. We also used interactions with enterprise executives leading outsourcing initiatives and influencing service provider decisions and conversations with service provider practice leaders.



The *Generative AI Services 2025 Market Insights* report offers a comprehensive analysis of key market trends and developments. It provides valuable insights into Gen AI adoption across various enterprise types, key industries, prevalent use cases, and the regulatory landscape shaping the Gen AI ecosystem.

Key reports of Avasant's Generative AI Services research



Market Insights™ 2025

This report provides a comprehensive overview of the Gen AI services landscape. It delves into the current state of the market and analyzes recent trends shaping the Gen AI services industry.

RadarView™ 2025

This report examines how leading service providers are delivering Gen AI services. It covers key supply-side trends and features a deep-dive analysis of providers that Avasant has recognized for their excellence and innovation in the space.

Defining generative AI


Artificial intelligence (AI) is the simulation of human intelligence in machines to perform tasks that typically require human intelligence.

Machine learning (ML) is a subfield of AI that uses algorithms to learn from data and make predictions or decisions.

Deep learning is a subfield of ML that uses neural networks with multiple layers to learn and make decisions.

Generative AI is a type of AI that focuses on generating new content, such as images, music, or text, based on patterns and examples in the training data.

Agentic AI is a form of AI that combines a large or small language model (LLM/SLM) with tools, APIs, and memory to interpret context, make decisions, and execute multistep tasks across enterprise environments.





Executive summary

Key enterprise Gen AI service trends shaping the market

Large enterprise deals involving Gen AI/agent AI rose 61% in 2025, with more contracts now embedding measurable Gen AI productivity

- Over the past 12 months, large enterprise deals involving Gen AI and agentic AI grew by 61%, driven by demand for AI-ready infrastructure, unified AI operations, robust data foundations, and sovereign AI deployments.
- In the same period, contracts explicitly quantifying Gen AI productivity gains rose by 29%, as providers tied managed service renewals and multiyear transformation deals to measurable workflow acceleration, efficiency metrics, and ROI benchmarks.

Enterprises are scaling Gen AI adoption across industries while proprietary LLM creation is rising for greater control and contextual intelligence

- High-tech and telecom lead Gen AI adoption at 20%, followed by BFS at 13%. These industries are deploying LLMs and AI agents to accelerate application deployment, improve multimodal data processing, and enhance omnichannel customer engagement.
- At the same time, about 10%–12% of Gen AI initiatives involve building proprietary LLMs to strengthen industry-specific nuance, regulatory compliance, contextual reasoning, edge-based processing, and strategic data independence.

Agentic AI is gaining traction across business functions, led by automation-heavy domains, with slower adoption in complex and compliance-driven areas

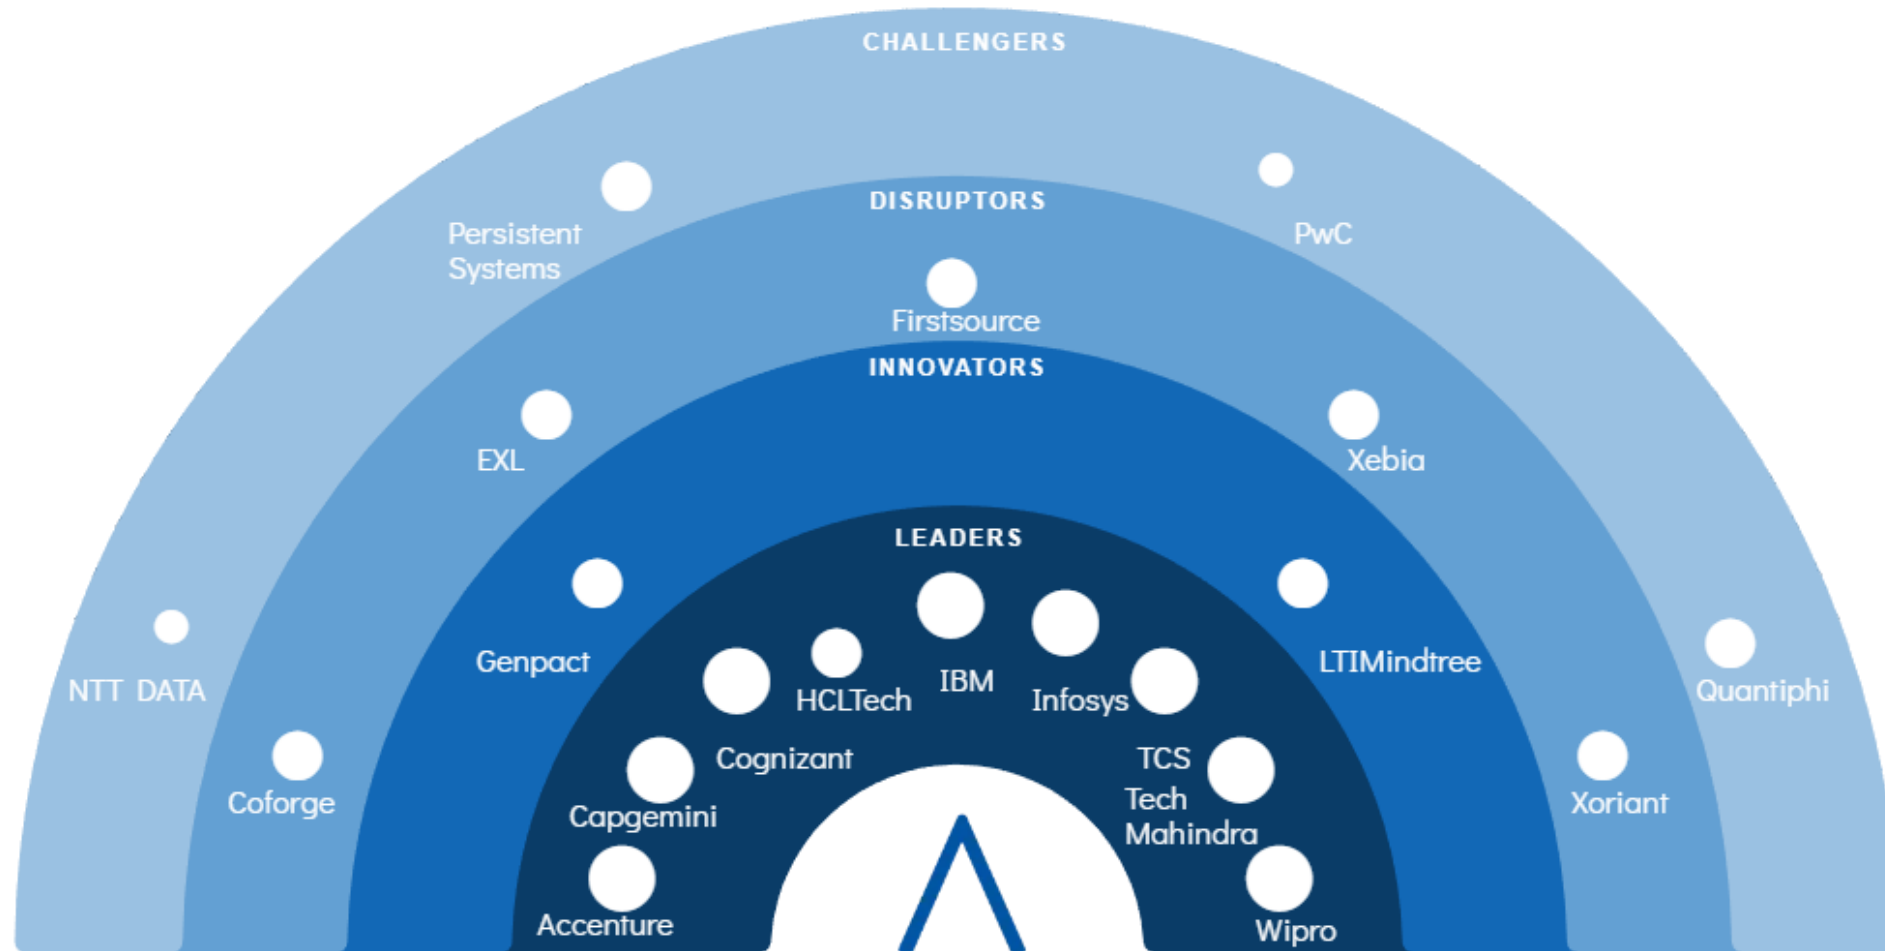
- IT, finance, operations, and customer service lead adoption, driven by structured workflows and measurable efficiency gains. Use cases include software development life cycle automation, fraud detection, performance tracking, and chatbot assistance.
- Supply chain, legal, and R&D adoption remains slower due to fragmented data, evolving regulatory frameworks, and higher model validation needs across use cases such as autonomous procurement, compliance audit, and experiment workflow automation.

Sovereign AI initiatives on rise as nations invest in compute, data, and foundational models to achieve technological self-reliance

- Nations are prioritizing six key pillars of sovereign AI—compute infrastructure, chip fabrication, foundational models, national data repositories, governance frameworks, and AI talent ecosystems—to build resilient and localized AI systems.
- These investments span domestic chip manufacturing, enterprise-focused datasets, AI testing sandboxes, sovereign cloud infrastructure, regulatory frameworks for responsible AI, and large-scale talent development programs, collectively aimed at achieving technological and data sovereignty.

Avasant's *Generative AI Services 2025 RadarView* recognizes 20 top-tier providers supporting the enterprise adoption of Gen AI

Practice maturity 






Note: Please refer to Avasant's *Generative AI Services 2025 RadarView*™ for a detailed analysis of the providers and supply-side trends.



State of the market






The number of large deals involving Gen AI and agentic AI has increased by 61% over the last year, reflecting robust enterprise demand for AI-native transformation

Factors driving Gen AI and agentic AI-backed large deals

AI-ready infrastructure 	Data foundations 	AI-first operating model 	AI governance 	Sovereign AI 
<p>Enterprises are establishing dedicated AI infrastructure footprints (GPU clusters, power, networking, & data center zones) as part of long-term transformation contracts.</p>	<p>Data lakehouse build-outs, integration, and governance are now co-funded through Gen AI programs, resulting in larger bundled deal sizes and contract durations.</p>	<p>Organizations are shifting from multiple small pilots to unified AI platforms (model hubs, agent orchestration, LLMOps, & observability) governed by a single partner.</p>	<p>Responsible AI, auditability, and traceability have become contractual requirements, considerably expanding security, risk, and compliance scopes.</p>	<p>Sovereign AI workloads necessitate localized AI deployments, which involve additional infrastructure build-outs and compliance layers within transformation initiatives.</p>

61%
Increase in number of large deals involving Gen AI and agentic AI in 2025


Illustrative examples

 <p>In November 2025, Deutsche Telekom signed a €1.0 billion multiyear contract with NVIDIA to build an AI factory in Munich, designed to deliver large-scale compute infrastructure for telecom, manufacturing, and enterprise clients across Europe.</p>	 <p>In October 2025, the NHS Business Services Authority awarded Infosys a £1.2 billion (15-year) contract to modernize its workforce and payroll systems through an AI-driven data platform. It will consolidate disparate datasets, enforce data governance, and enable analytics-led decision-making.</p>	 <p>In May 2025, a global agribusiness leader signed a \$450 million (seven-year) deal with LTIMindtree to implement an AI-powered operating model, integrating predictive automation, AI-led applications, and digital infrastructure to drive end-to-end operational modernization.</p>	 <p>In September 2025, the US Department of Defense signed a \$100 million (five-year) agreement with Scale AI to deploy an enterprise-class AI platform on its classified networks, embedding governance, audit-ready tooling, and secure data controls for mission-critical operations.</p>	 <p>Under its \$1 billion multiyear contract with AWS, e& launched the UAE Sovereign Launchpad in November 2025, a locally managed sovereign AI cloud platform designed to ensure in-country data residency, zero-trust security, and compliant GPU infrastructure for regional AI workloads.</p>
---	--	--	--	--

Note: Large deals are defined as contracts with a total contract value (TCV) exceeding USD 50 million.







In the same period, the share of managed service contracts explicitly itemizing Gen AI productivity gains rose by 29%, reflecting a growing focus on clear Gen AI ROI

Gen AI productivity benefits as part of managed service contracts



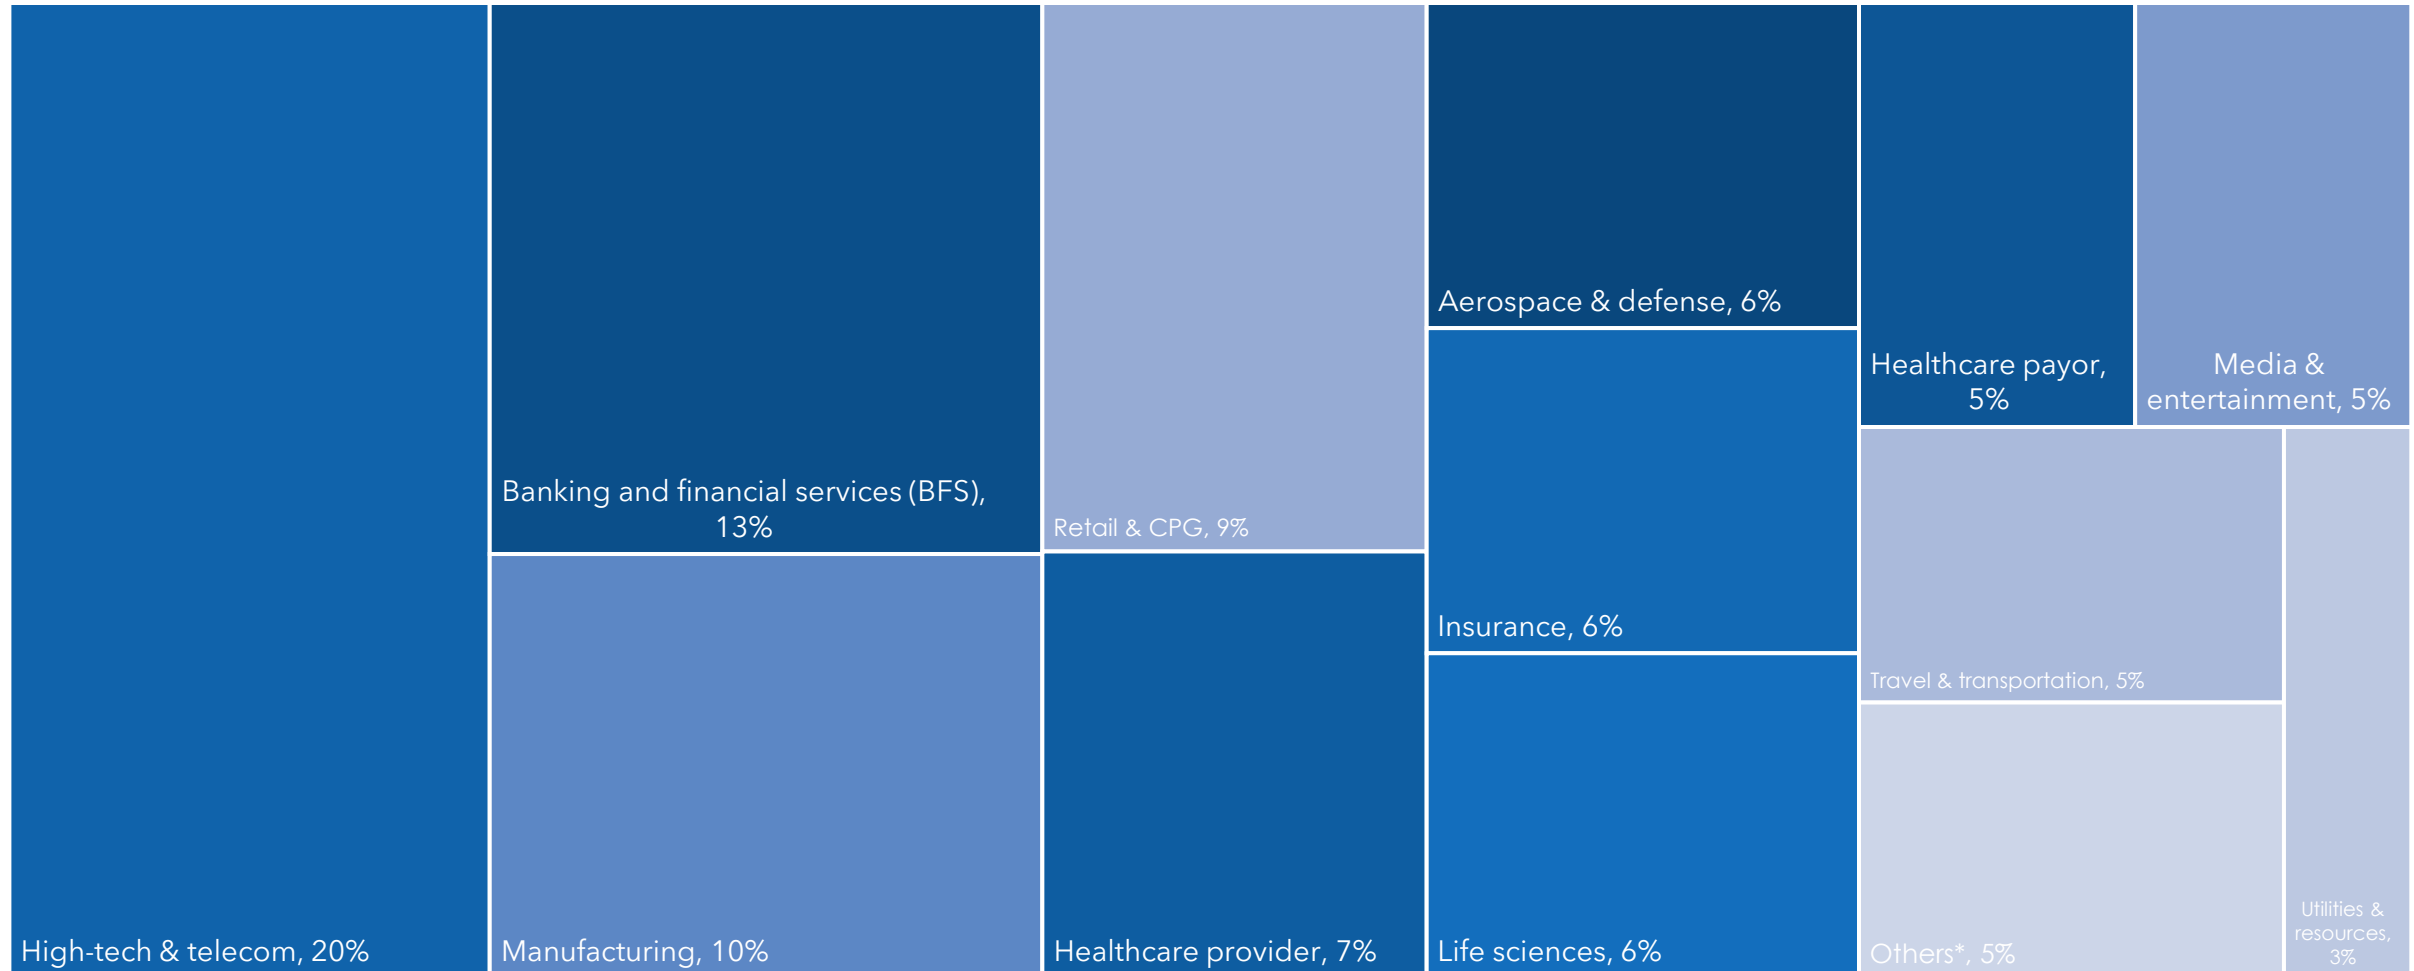
- No specific line item for Gen AI, but its productivity impact is bundled with other technologies
- Yes, showing Gen AI as a separate line item

- With Gen AI becoming a separate line item, managed service deals are being renegotiated to embed AI-driven productivity gains.
- Upfront savings from reduced human effort and operational efficiency are now expected as part of renewal pricing.

Enterprise	Use case	Incremental Gen AI impact	Implementation description
 A financial services firm	Application modernization	70% acceleration in modernization timelines	Deployed a Gen AI platform to streamline app modernization, accelerating migration timelines and reducing technical debt.
 Healthcare and agriculture firm	IT service desk automation	~25% reduction in operational costs	Implemented a Gen AI-based assistant, reducing incident volumes by 20%-25%, improving response times by over 90%, and achieving 25% cost savings.
 A US-based healthcare firm	Patient claims processing	35% enhancement in document processing time	Deployed a Gen AI-driven claims processing solution, enhancing classification accuracy by 30% and straight-through processing by 25%.
 A Japanese automotive manufacturer	Aftermarket operations optimization	Productivity improvement by 10%-15%	Implemented a Gen AI-powered aftermarket solution, reducing parts inventory turnover by 5%-7% and increasing order fill rates by 10%-15%.
 A Swedish multinational manufacturer of luxury vehicles	Warranty claims optimization	40% enhancement in assessor productivity	Deployed a Gen AI-driven claims warranty analysis platform, improving assessor productivity by 40%, identifying 16% false claims, and reducing warranty spend.
 A global homestay provider	Automated data annotation	60%-70% productivity improvement	Implemented a Gen AI-powered chat labelling and categorization solution, achieving 15-20x faster turnaround and 60%-70% productivity improvement.

High-tech and telecom industries are the biggest adopters of Gen AI, focusing on faster app development, followed by BFS, which prioritizes improved CX

Enterprises are leveraging LLMs and AI agents for use cases such as application deployment and testing, multimodal data processing, risk and compliance monitoring, automated credit analysis, and omnichannel customer relationship management.







Note: The percentages in this figure represent the distribution of 1,267 service provider clients in the Gen AI space across industries. Others refer to (professional services, nonprofit, government, etc.)

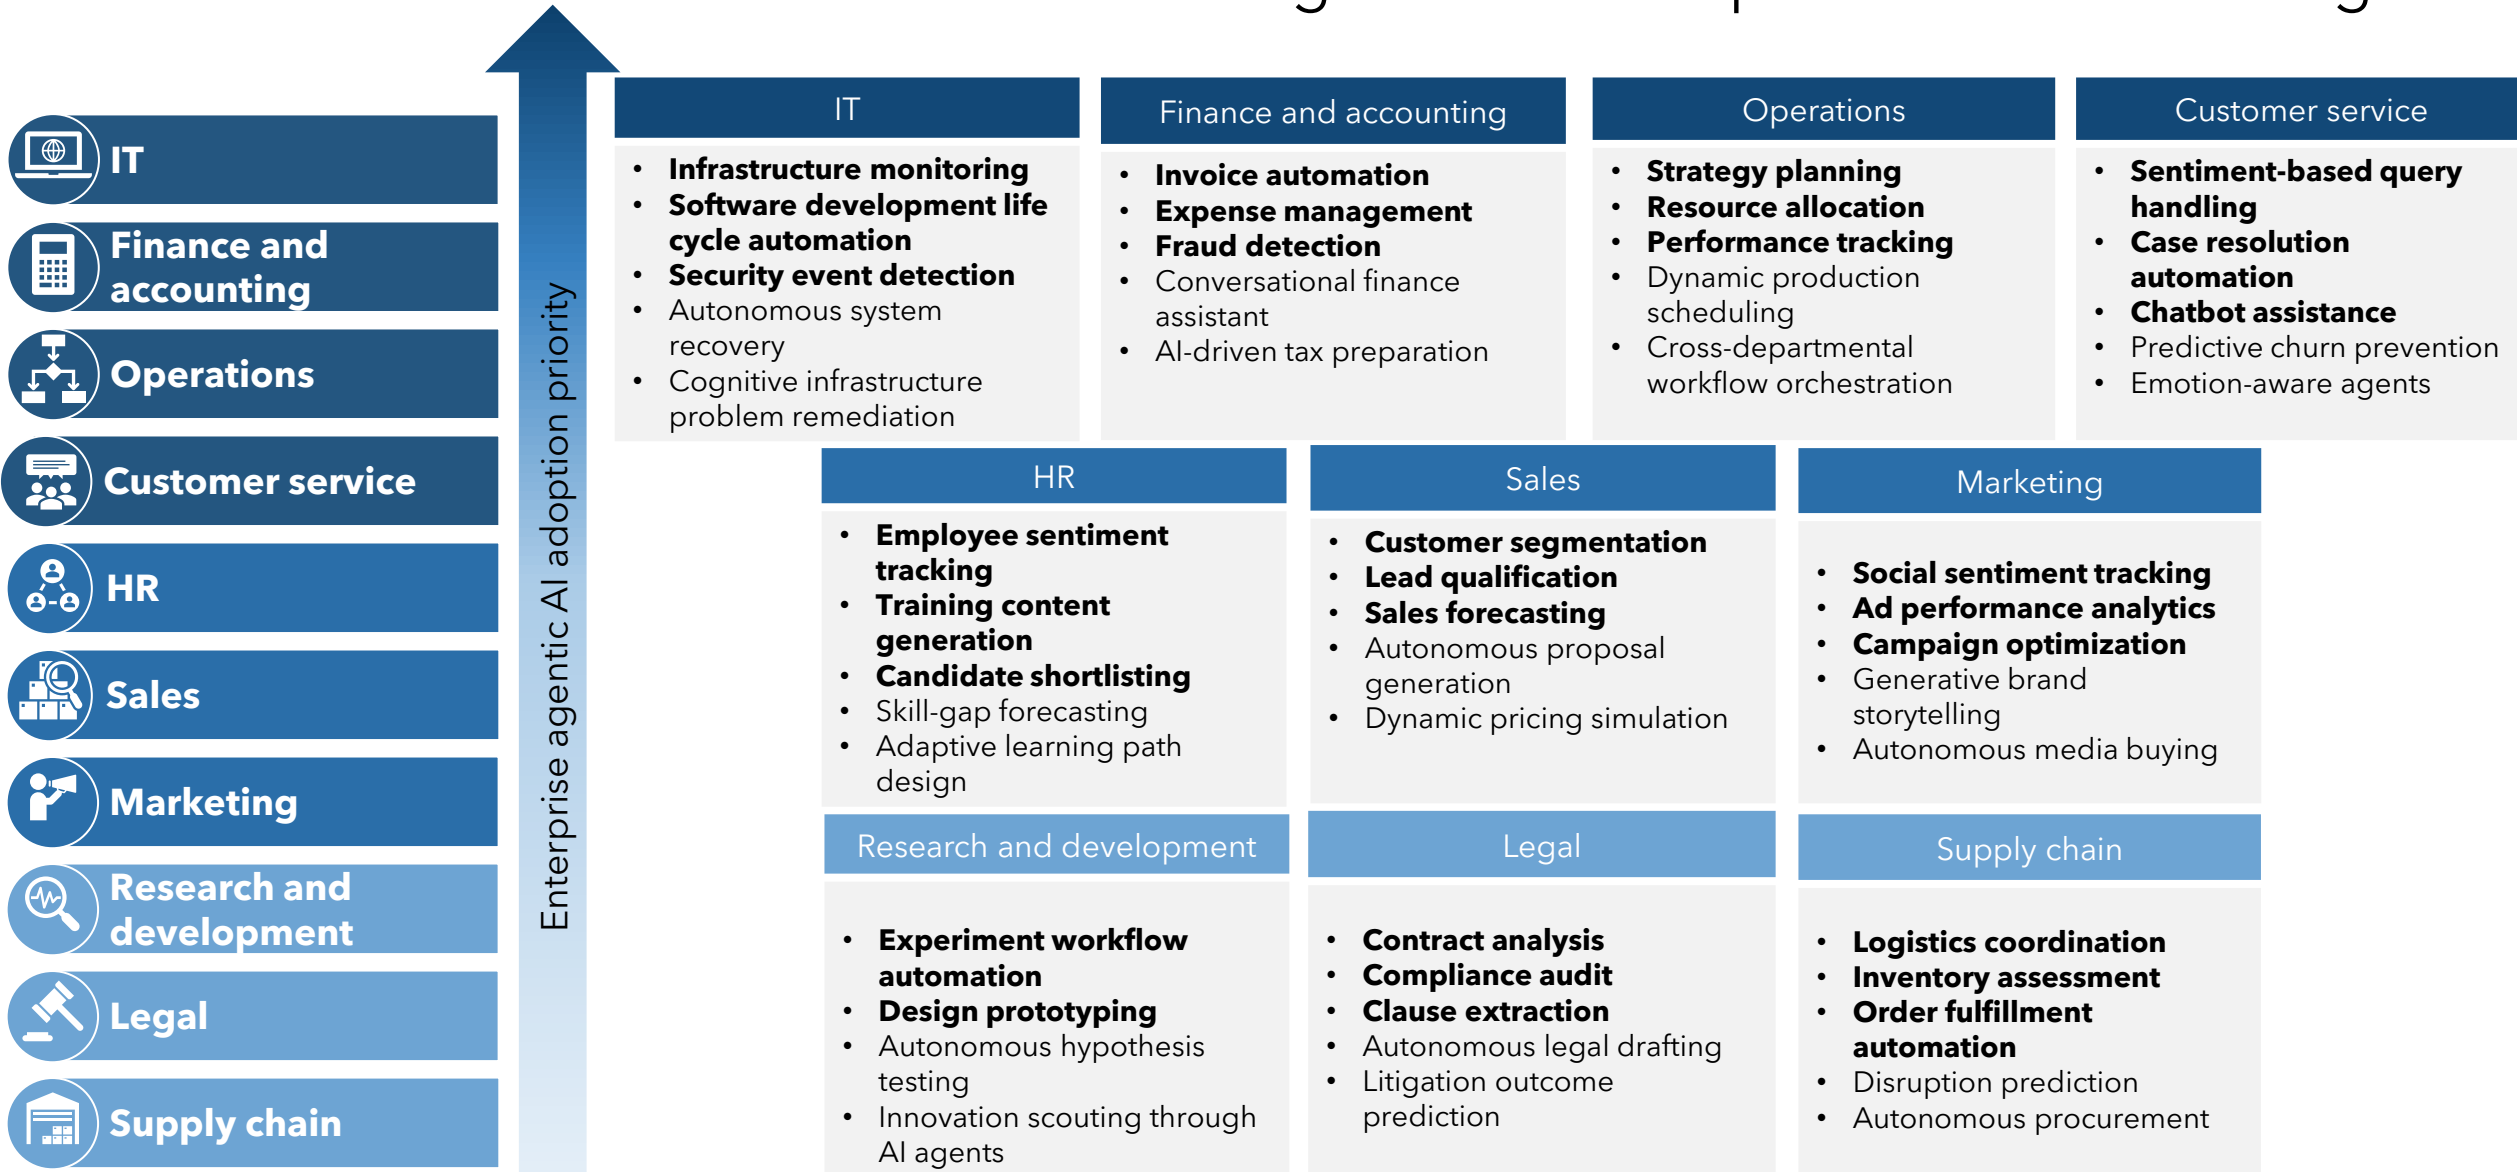
About 10%-12% of enterprise Gen AI projects now involve building custom LLMs and SLMs, reflecting a shift toward domain understanding and contextual nuances

10%-12%
of enterprise
Gen AI projects
now involve the
development of
proprietary LLMs
and SLMs

- Enterprises are increasingly building proprietary LLMs and SLMs to overcome the constraints of general-purpose models, focusing on accuracy, data confidentiality, and compliance with industry regulations.
- These models are trained on internal, domain-specific datasets to generate more precise insights for complex functions such as drug discovery, supply chain optimization, and personalized customer engagement.
- Advancements in compute infrastructure, model compression, and open-source frameworks are making it technically and economically feasible for enterprises to develop and deploy their own LLMs at scale.

Enterprise	Reason for LLM development	Implementation description
 MERCK	Industry-specific nuance understanding	Merck built a life sciences-specific LLM for clinical trial summarization and R&D hypothesis generation across drug discovery, regulatory, and translational science teams and analytics globally.
FOXCONN	Regulatory and sovereignty compliance	Foxconn developed FoxBrain, a proprietary Chinese LLM trained on production-line telemetry, defect images, and supply-chain data, to enable predictive maintenance and process optimization across factories.
Walmart 	Contextual reasoning from internal datasets	Walmart developed Wallaby, a retail-focused proprietary LLM, for contextual product search, personalized recommendations, and conversational experiences across shopping, merchants, and support teams.
 LG	Edge-based data processing and knowledge management	LG Group built EXAONE, a large-scale proprietary LLM running on custom chips for advanced manufacturing, supply chain optimization, and product intelligence across its electronics and industrial business units
 Z O H O	IP control and strategic independence	Zoho deployed Zia LLM, a proprietary LLM hosted in Zoho data centers, to enable private fine-tuning, on-premises inference, branded agents across Zoho applications, and customers worldwide.


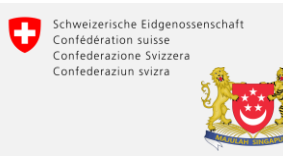


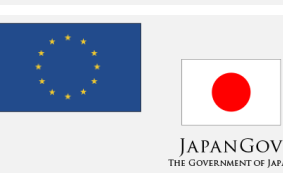
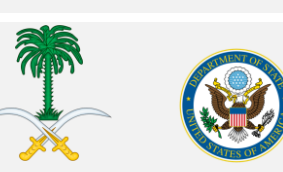
Enterprises are advancing agentic AI adoption in IT, finance, operations, and customer service to drive contextual intelligence and adaptive decision-making



Nations are advancing sovereign AI through targeted investments in compute, data, and foundational models to strengthen self-reliance and strategic AI leadership

Pillars of sovereign AI

Illustrative examples of sovereign AI initiatives worldwide

<p>Chip and compute hardware</p>	<ul style="list-style-type: none"> In August 2025, China's government mandated that publicly-owned data centers use at least 50 % domestically-produced AI chips, and companies such as Huawei and Alibaba began ramping up production of AI accelerators and domestic chips. In July 2025, the US released America's AI Action Plan, which includes directives and funding priorities to secure domestic AI hardware supply chains and expand local compute capacity (federal support for data-center/AI-accelerator projects). 	
<p>Foundational models</p>	<ul style="list-style-type: none"> In September 2025, the Swiss government, together with EPFL*, ETH Zurich, and Swiss National Supercomputing Centre, unveiled Apertus, a national LLM trained on public data from over 1,000 languages, with open-weight release. In May 2025, the Singapore government, via Infocomm Media Development Authority and A*STAR I2R, unveiled the Multimodal Empathetic Reasoning and Learning in One Network (MERaLiON), an LLM trained on eight languages, including Malay, Tamil, Thai, Bahasa Indonesia, Vietnamese, English, Mandarin, and English. 	
<p>Cloud and compute infrastructure</p>	<ul style="list-style-type: none"> In July 2025, the UK government published its UK Compute Roadmap, committing up to £2 billion to build a sovereign compute ecosystem (including over £1 billion to expand the AI research resource) by 2030. In September 2025, India announced its Sovereign AI Computing initiative with funding of \$2.4B, aimed at building domestic AI infrastructure (including data centers and supercomputing hubs). 	
<p>National data libraries and foundries</p>	<ul style="list-style-type: none"> In January 2025, the UK government announced the creation of a National Data Library to unlock public-sector datasets for AI research. In March 2025, the Indian government launched the AIKosha datasets platform under IndiaAI to provide machine-readable Indian language and public-domain data for training foundational models and reducing dependence on foreign AI ecosystems. 	
<p>Responsible AI and ethics</p>	<ul style="list-style-type: none"> In August 2024, the EU AI Act came into effect, establishing obligations around transparency, risk classification, and governance of AI systems. In May 2025, the government of Japan enacted the "Act on the Promotion of Research and Development and the Utilization of AI-Related Technologies," establishing a national framework for AI development and usage with requirements around transparency, safety, and alignment with public welfare 	
<p>Talent and IP retention</p>	<ul style="list-style-type: none"> In May 2025, Saudi Arabia launched the HUMAIN initiative under the Saudi Data & AI Authority, consolidating top talent from companies such as Aramco, TONOMUS, and global tech firms to scale Gen AI workforce readiness aligned with the kingdom's Vision 2030. In August 2025, the National Science Foundation announced new funding opportunities to build the AI workforce (K-12 through research career pipeline) as part of America's AI Action Plan to retain AI talent and leadership. 	

Note: EPFL denotes École Polytechnique Fédérale de Lausanne; A*STAR I2R denotes Agency for Science, Technology and Research - Institute for Infocomm Research

AVASANT

Future outlook

The next frontier of Gen AI will focus on physical AI, multi-AI agentic systems, sustainable computing architectures, and real-time interactive environments



Physical AI and robotics

- Current AI robots have limited coordination and adaptability. With the convergence of physical AI and Gen AI, robots are evolving into embodied systems that merge perception with generative reasoning.
- Recent launches such as Accenture's Physical AI Orchestrator and NVIDIA's Generative Physical AI platform (Omniverse and Project GR00T) highlight this shift toward intelligent, adaptive robotics.



Agent-to-agent communication protocols

- Current AI agents lack seamless interoperability across enterprise systems, which limits their ability to coordinate and share context. With the emergence of agentic commerce and interconnected ecosystems, these communication protocols will enable structured interaction between agents and tools.
- Frameworks such as Google's A2A Protocol and Anthropic's MCP mark the first steps toward multi-agent collaboration.



Real-time interactive 3D environments

- Real-time interactive environments evolve traditional simulations into dynamic, continuously updating systems where humans, AI agents, and devices interact in real time.
- By integrating Gen AI-driven digital twins, firms can enable adaptive scenario modeling for industrial design, remote operations, and workforce training, as shown in the May 2025 metaverse-linked smart building twin study.



Sustainable Gen AI computing systems

- Sustainable Gen AI computing systems aim to curb the soaring energy footprint of model training and inference through carbon-aware scheduling and renewable-powered infrastructure.
- Recent advances such as Google's carbon-aware workload placement, Microsoft's geothermal and fusion-powered AI data centers, and EcoServe's emission-optimized inference design highlight the shift toward greener Gen AI compute.



Autonomous AI governance with meta-AI agents

- Autonomous AI governance is emerging as the next step in self-regulating AI systems, where supervisory agents oversee compliance and coordination across multi-agent, decentralized, and cross-organizational networks.
- For instance, in April 2025, the LOKA Protocol proposed meta-AI agents that autonomously enforce ethical and policy alignment within such distributed agent frameworks.




Quantum computing-powered Gen AI

- While current Gen AI systems struggle with training data fidelity and classical computational limits, quantum-powered frameworks can overcome these constraints.
- In February 2025, Quantinuum introduced its Generative Quantum AI framework, using quantum-generated data from its H2 computer to target enterprise applications in drug discovery, financial modelling, and real-time logistics optimization.

AVASANT

Infosys profile

Infosys: RadarView profile



- Practice maturity** ★★★★★
- Partner ecosystem** ★★★★★
- Investments & innovation** ★★★★★

Specializes in industry-specific Gen AI deployments with proprietary SLMs and governance offerings. Expands Gen AI governance and agentic AI capabilities through strategic partnerships.

<h3>Practice overview</h3> <ul style="list-style-type: none"> Practice size: N/A Active clients: N/A External certifications: 42,000+ Delivery highlights: 292 delivery centers worldwide <div style="display: flex; justify-content: space-around;"> <div style="border: 1px solid gray; padding: 5px; text-align: center;">850+ Gen AI engagements</div> <div style="border: 1px solid gray; padding: 5px; text-align: center;">200+ Prebuilt AI agents</div> </div>	<h3>Client case studies</h3> <ul style="list-style-type: none"> Implemented a generative (Gen) AI solution for a US-based healthcare company to streamline integration of Gen AI across its business functions. It reduced application development efforts by 15%–20%, ticket documentation efforts by 90%, and ticket analysis and triage efforts by 20%–25%. Provided a Gen AI-driven multi-AI agentic system for a European life sciences company to enhance the review of drug-related regulated content. It improved review accuracy by 85%, reduced manual review costs by more than 50%, and accelerated regulatory submissions. Deployed a Gen AI chatbot for a European healthcare benefits provider to streamline the knowledge management process. It accelerated document search by 70%, increased operational efficiency, and reduced manual efforts by 40%. Implemented Infosys AP on Cloud driven by agentic AI for an Asian restaurant chain to streamline accounts payable processing. It translated invoice languages with 90% accuracy, reduced manual efforts by over 50%, and improved invoice turnaround time by 40%. 															
<h3>Key IP and assets</h3> <ul style="list-style-type: none"> Infosys Topaz: A unified suite of Gen AI platforms, tools, and solutions that include: <ul style="list-style-type: none"> Infosys Responsible AI suite: A suite of responsible AI guardrails to detect anomalies corresponding to large language models (LLMs) PolyAI: A vendor-agnostic Gen AI deployment architecture Agentic Foundry: A suite of prebuilt domain-specific AI agents Infosys Accounts Payable on Cloud: An AI/Gen AI-powered accounts processing solution 	<h3>Key partnerships</h3> <h4>Technology partners</h4> <h4>Universities/academia</h4>	<h3>Sample clients</h3> <ul style="list-style-type: none"> A US-based healthcare company A European life sciences company A European healthcare benefits provider An Asian restaurant chain An American firm dealing in power generation products A British professional services firm An American electronics company A pharmaceutical company A US-based airline company An oil and gas processing company 	<h3>Industry coverage</h3> <table border="1"> <tr><td>Banking</td></tr> <tr><td>Financial services</td></tr> <tr><td>Government</td></tr> <tr><td>Healthcare & life sciences</td></tr> <tr><td>High-tech</td></tr> <tr><td>Insurance</td></tr> <tr><td>Manufacturing</td></tr> <tr><td>Nonprofits</td></tr> <tr><td>Professional services</td></tr> <tr><td>Retail & CPG</td></tr> <tr><td>Telecom, media & entertainment</td></tr> <tr><td>Travel & transportation</td></tr> <tr><td>Utilities & resources</td></tr> </table>	Banking	Financial services	Government	Healthcare & life sciences	High-tech	Insurance	Manufacturing	Nonprofits	Professional services	Retail & CPG	Telecom, media & entertainment	Travel & transportation	Utilities & resources
Banking																
Financial services																
Government																
Healthcare & life sciences																
High-tech																
Insurance																
Manufacturing																
Nonprofits																
Professional services																
Retail & CPG																
Telecom, media & entertainment																
Travel & transportation																
Utilities & resources																

Darker color indicates higher industry coverage: ●●●●●

Infosys: RadarView profile

Analyst insights

Practice maturity



- Infosys supports enterprises in Gen AI implementations through its Infosys Topaz suite of AI/Gen AI services focused on Gen AI assessment and guardrail integration, domain-specific LLM and small language model (SLM) customization, and industry-specific agentic AI deployments.
- It supports enterprise Gen AI implementations with its suite of prebuilt platforms, tools, and solution accelerators. These include industry-specific AI blueprints, real-time LLM performance monitoring and optimization tools, and domain-specific solution accelerators for use cases such as code development and customer engagement. Additionally, it has also built a suite of SLMs across domains such as ITOps, cybersecurity, and finance, and industries such as banking and financial services.
- Focusing on responsible AI, it offers multiple responsible AI offerings backed by its centralized responsible AI office. These include the AI3S suite of Gen AI guardrails and the open-source RAI toolkits. These tools address input risks, such as personally identifiable information leaks, prompt injections, copyright issues, and toxic content requests, and monitor LLM outputs for inaccuracies and policy violations.

Partner ecosystem



- In May 2025, it partnered with Google Cloud to deploy over 200 AI agents across industries such as healthcare, finance, retail, telecom, manufacturing, and agriculture, leveraging Infosys Topaz and Google Vertex AI.
- In March 2025, Infosys teamed up with Linux Foundation to strengthen responsible AI capabilities for enterprise network infrastructure. As part of this collaboration, Infosys will contribute its open-source RAI toolkit and AI app development framework to two open-source networking projects,
- In April 2025, Infosys collaborated with Confluent to codevelop Gen AI and agentic AI solutions, powered by real-time data streaming capabilities. This collaboration supports enterprises in modernizing data infrastructure and scale AI applications across sectors such as banking and financial services.

Investments and innovation



- In August 2025, Infosys formed a joint venture with Telstra, acquiring a 75% stake in its cloud services unit, to strengthen its Infosys Topaz AI and Gen AI capabilities. This builds on its previous acquisition of The Missing Link in April 2025 to strengthen its Gen AI-powered cybersecurity services in Australia.
- In May 2025, Infosys launched the Agentic Foundry to support agentic AI implementations. It will assist users in building AI agents across domains such as IT and backend operations. In May 2025, its business process management division launched a suite of prebuilt AI agents for finance and accounting processes.
- It has established multiple Gen AI CoEs across its global delivery centers to codevelop Gen AI solutions with clients. It plans to support these initiatives through strategic hirings in technologies such as Gen AI, multimodal AI, computer vision, and agentic AI.

AVASANT

Appendix:
About RadarView

Avasant's Generative AI Services 2025 RadarView assesses service providers across three critical dimensions

Practice maturity

- This dimension considers the current state of the provider's Gen AI practice in terms of its strategic importance for the provider, the maturity of its offerings and capabilities, and client engagement.
- The width and depth of the client base, use of proprietary/outsourced tools and platforms, and quality of talent and execution capability are all important factors that contribute to this dimension.

Partner ecosystem

- This dimension assesses the nature of the provider's partnerships, their objectives (codevelopment and co-innovation), and the provider's engagement with solutions providers, startup communities, and industry associations.
- The kind of joint development programs around offerings, go-to-market approaches, and the overall depth of partnerships are all important aspects.

Investments and innovation

- This dimension measures the strategic direction of the provider's investments and the resultant innovations in the offerings and commercial model and how it aligns with the future direction of the industry.
- The overall strategic investments, both organic and inorganic, toward capability and offering growth, technology development, and human capital development, along with the innovations that the service provider develops with its partners, are critical aspects of this dimension.

Research methodology and coverage

Avasant based its analysis on several sources:

Public disclosures

Publicly available information from sources such as Securities and Exchange Commission filings, annual reports, quarterly earnings calls, and executive interviews and statements

Market interactions

Discussions with enterprise executives leading digital initiatives and influencing provider selection and engagement

Provider inputs

Inputs collected through an online survey and structured briefings in August 2025–October 2025

Of the 48 service providers assessed, the following providers are the final 20 featured in the *Generative AI Services 2025 RadarView* are:



Reading the RadarView

Avasant recognizes Generative AI service providers in four classifications:



Leaders show consistent excellence across all key dimensions of the RadarView assessment (practice maturity, partner ecosystem, and investments and innovation) and have had a superior impact on the marketplace. These providers have shown true creativity and innovation and have established trends and best practices for the industry. They have proven their commitment to the industry and are recognized as thought leaders in their space, setting the standard for the rest of the industry to follow. Leaders display a superior quality of execution and a reliable depth and breadth across verticals.



Innovators show a penchant for reinventing concepts and avenues, changing the very nature of how things are done from the ground up. Unlike leaders, innovators have chosen to dominate a few select areas or industries and distinguish themselves based on superior innovation. These radicals are always hungry to create pioneering advancements in the industry and are actively sought after as trailblazers, redefining the rules of the game.



Disruptors enjoy inverting established norms and developing novel approaches that invigorate the industry. These providers choose to have a razor-sharp focus on a few specific areas and address those at a high level of granularity and commitment, resulting in tectonic shifts. While disruptors might not have consistent depth and breadth across many verticals like leaders or the innovation capabilities of innovators, they exhibit superior capabilities in their areas of specialization.



Challengers strive to break the mold and develop groundbreaking techniques, technologies, and methodologies on their way to establishing a unique position. While they may not have the scale of the providers in other categories, challengers are eager and nimble and use their high speed of execution to great effect as they scale heights in the industry. Challengers have a track record of delivering quality projects for their most demanding Global 2000 clients. In select areas and industries, challengers might have capabilities that match or exceed those of the providers in other categories.

Key contacts

Primary contact:



Abhisekh Satapathy

Principal Analyst
abhisekh.satapathy@avasant.com

Contributors:



Michael Stoeckert

Distinguished Fellow at Avasant
michael.stoeckert@avasant.com



Anupam Govil

Managing Partner and Digital
Practice Lead
anupam.govil@avasant.com



Swapnil Bhatnagar

Partner
swapnil.bhatnagar@avasant.com



Chandrika Dutt

Research Director
chandrika.dutt@avasant.com

Disclaimer

Unauthorized reproduction or distribution in whole or in part in any form, including photocopying, faxing, image scanning, e-mailing, downloading, or making available any portion of the text, files, data, graphics, or other materials in the publication is strictly prohibited. Prior to photocopying items for internal or personal use, please contact Avasant to ensure compliance with Avasant or third-party intellectual property rights and usage guidelines. All trade names, trademarks, or registered trademarks are trade names, trademarks, or registered trademarks of Avasant, its licensors, or the applicable third-party owner. No express or implied right to any Avasant or third-party trademarks, copyrights, or other proprietary information is granted hereunder. Avasant disclaims, to the fullest extent under applicable law, all warranties and conditions, expressed or implied, with respect to any content provided hereunder, including, without limitation, warranties of merchantability and fitness for a particular purpose. Avasant does not assume or guarantee and hereby disclaims any and all liability for the quality, accuracy, completeness, or usefulness of any information contained herein, which shall be inclusive of any and all direct or consequential damages or lost profits. Any reference to a commercial product, process, or service does not imply or constitute an endorsement of the same by Avasant. This publication is for information purposes only. By distributing this publication, Avasant is not engaged in rendering legal, accounting, or other professional services. If legal, accounting, or other advisory services or other expert assistance is required, the services of a competent professional person should be sought.

AVASANT

EMPOWERING BEYOND

GET CONNECTED



www.Avasant.com