ISG®

September 2023

State of Applied Generative AI Market

Section 1: Introduction

The phenomenon of generative artificial intelligence (AI) has caught the attention of not only technology and business leaders but also the average individual. In fact, according to ISG Research, 85% of enterprises believe that investment in generative AI technology in the next 24 months is important or critical (from the ISG study: 2023 Future Workplace). This technology did not appear out of a vacuum; it leverages decades of work and is distinctly different from predecessor methods.

Historically, one of the first milestones of functional AI was its ability to recognize patterns via assimilating data into static algorithms. However, the use cases for this learn-only capability are narrow, largely because it requires intensive human management and training. Further advancements in machine learning and neural networks facilitated AI to learn and repeat tasks, enabling it to actively engage, interpret and replicate data inputs. The human role was less burdensome, enabling new business improvements and efficiencies.

Now, AI has a new, powerful capability: the ability to generate content.

While generative AI is likely to have a deep and lasting impact, there are many challenges to overcome – from navigating security, copyright data challenges, ethics and legal concerns to fine-tuning enterprise-grade use cases.

Our aim with this report is to look past the hype and conduct a comprehensive analysis of the current landscape for enterprise adoption of generative AI. It will be followed by further reports and snapshots of the landscape as it evolves. This study has two goals:

1. To report our findings of top trends, use cases, solutions, industries and challenges related to adopting generative AI, and
2. To share data-driven insights and actionable recommendations to assist businesses in achieving impactful transformations across diverse sectors.

Evolution of AI

Origins

AI Can **Learn**

HUMAN ROLE: Train and manage
USED BY: Analytics and data science
TO: Support decisions & optimization

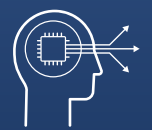

Yesterday

AI Can Learn and **Repeat**

HUMAN ROLE: Supervise and intervene
USED BY: Business process improvement
TO: Automate repetitive tasks

Today

AI Can Learn, Repeat, **Create and Re-Create**

HUMAN ROLE: Oversee and course correct
USED BY: Developers and content creators
TO: Turbocharge digital development

Section 2: Provider Engagement Methodology

Methodology

The ISG State of Applied Generative AI report outlines and analyzes the most recent and thought-provoking developments in the generative AI solutions and services market. Using the ISG Research methodology which consists of a multi-phased research and analysis process, the report explores active investments and ongoing initiatives of relevant service providers across the global market. Our multi-phased and methodical research approach uses a combination of information-gathering processes, including questionnaire-based surveys, intensive conversations and interactive discussions with providers offering generative AI solutions and services. The report offers insights intended to help enterprises make informed decisions and engage in active dialogues to drive transformational impact across their AI projects.

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

The research and analysis presented in this report includes insights from ISG Provider Lens™ provider evaluation research, 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 represents information that ISG believes to be current as of September 2023 for providers that actively participated as well as for providers that did not. ISG recognizes that several active use cases and projects that were in the pipeline have taken place since that time, but those changes are not reflected in this report.

ISG interviewed the following service and technology providers as part of our research for this study: Accenture, Capgemini, Cognizant, Genpact, HCLTech, IBM, Infogain, Infosys, Persistent Systems, Quantiphi, SLK Group, TCS, Tech Mahindra, TheMathCompany, Tiger Analytics, Tredence, UST Global and Wipro.

Process Description

The study was divided into the following steps:

1. Definition of the generative AI solutions and services market
2. Use of questionnaire-based surveys of service providers across all trend topics
3. Interactive discussions with service providers/vendors on capabilities and use cases
4. Use of ISG's internal databases and advisor knowledge and experience (where applicable)
5. Detailed analysis and evaluation of services and service documentation based on the information received from providers and other sources
6. Use of the following key evaluation criteria:
   - Strategy and vision
   - Technology competence and innovation
   - Breadth and depth of portfolio of services offered
   - Delivery, functionality and performance capabilities
   - Ethics, responsibility and privacy considerations
   - Customization, flexibility and scalability
   - Use cases, case studies and reference projects
   - Technology advancements
   - Brand awareness and presence in the market
   - Sales and partner landscape

Section 3: Key Findings & Trends

1. The hype is working

The majority of enterprises are paying deep attention to generative AI and actively want to experiment with and look for business applications. This is demonstrated not only by the breadth of vertical use cases but also by the functional use cases identified in our study (see use cases section). There is a sense of urgency in the market: the time to turn generative AI into a competitive edge is now.

2. Anti-blank-slate attitude says use what you have

We see a lot of enterprises asking their providers to lean in and consider how generative AI might be adopted into existing services such as call center and virtual assistant operations. Successful providers are looking to proactively engage with enterprises to explore these possibilities.

3. Specialist providers drive the market

Beyond the top-down imperative to embrace generative AI, most enterprises lack the focus to select the right use cases. Our study finds an immature but diverse market of inventive solutions driven mostly by providers that are actively engaging with enterprises to brainstorm and co-create. We see providers engaging by:

- Building domain-specific generative AI models developed from proprietary data sets
- Developing generative AI adoption platforms customized to enterprise needs
- Establishing new roles needed to support generative AI-driven solutions
- Creating solutions to unite analytics, AI and generative AI to solve business problems
- Continuously expanding model-training datasets via dedicated resources so foundational models can be trained with more data and get better over time
- Collaborating with enterprises as leaders achieve quick wins and re-apply learnings to fill gaps in use cases according to business needs, especially for functional process improvements

4. Robust generative AI architecture is rare

Our primary market assessment shows a gap in establishing generative AI architecture. Despite the potential for transformation, most service providers haven't yet given deep attention to architecture. While it's understandable that rethinking AI architecture is still at an early, evolving stage, it's still critical to gain a clear understanding of the best approach for your business before proceeding (See Section 8 for more details).

5. Fear of data hallucination prioritizes data isolation

Enterprises are cautious about pushing the capabilities of generative AI too far, too soon. The risk of getting it wrong looms large in the minds of leaders. Enterprises supported by customized partner solutions often prioritize isolation and control of data to train generative AI capabilities. We will likely continue to see this caution as a trend with future custom solutions.

6. Enterprises weigh risk vs reward

Despite the enthusiasm for implementing generative AI, enterprises have a few notable concerns.

- **ROI**. CFOs are getting smarter about tech investments and their returns. They know that experimentation only takes you so far, so they want to see a clear ROI. More than ever, successful providers must put some skin in the game.
- **Legacy data challenges**. All AI models need data to train. While there is an opportunity to get smarter at solving data-quality issues via synthetic data generation and by leveraging AI to extract data from unstructured sources, legacy challenges will still lead to familiar struggles in getting quality results.
- **Security, legal, copyright and ethical implications**. The dark side of AI hype haunts leaders with what-if scenarios about security, legal and ethical risks. Providers that are emerging leaders in this space must work with organizations to establish the right security and ethical guardrails for AI.

7. Transformative potential transcends function

Enterprises believe that generative AI is transformative enough to fundamentally change cross-functional business operations. Unlike traditional AI, which has use cases mostly in analytics and data science, generative AI is being built into cross-functional solutions.

Section 3: Key Findings & Trends (cont'd)

We see organizations exploring intersections between marketing, sales, demand management, connected products, industry value chains, intelligent metrics and IT/DevOps.

Summary

One clear theme throughout the study is that the art of the possible is looming large. And, now, signs of early caution and reality are starting to appear. We are approaching the very top of the hype curve. As use cases turn into real adoption, a clearer picture will emerge.

The diagram to the right calls out some common cross-functional use cases that we are seeing emerge for generative AI. For example, a use case that combines the analytics function with the IT and DevOps function could enable contextual data extraction, and if you add in the industry value chain function, the use case evolves toward performance tuning.

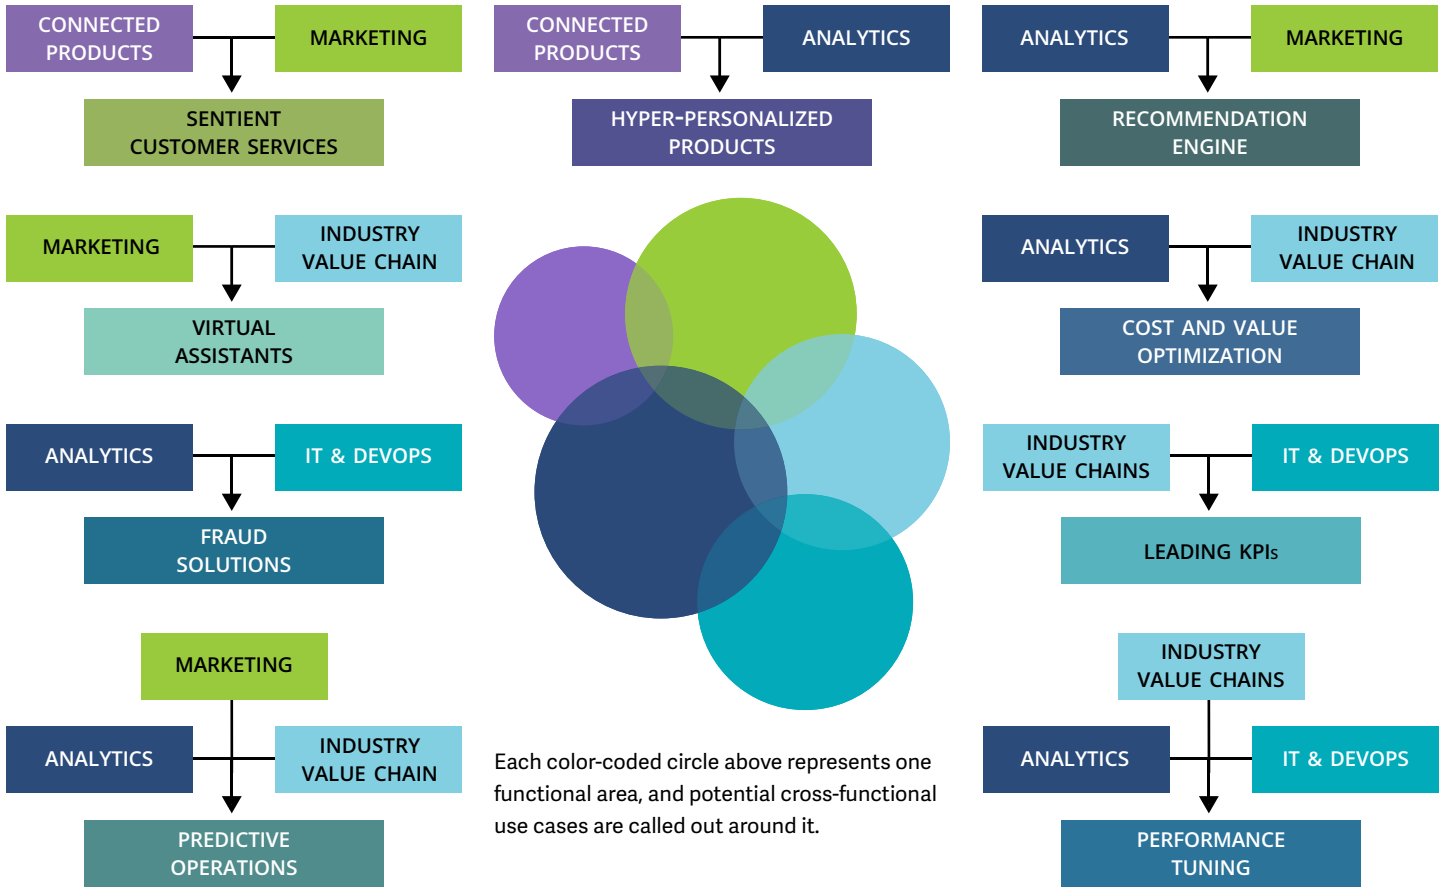

Each color-coded circle above represents one functional area, and potential cross-functional use cases are called out around it.

Section 4: Transformative Industries by Use Cases

Adoption by industry has been far from uniform. Our study shows a heavy over-indexing in financial services (including banking and insurance) at **24% of total use cases**. Manufacturing, Healthcare, Pharmaceuticals and Business Services have a more uniform distribution of 14-11% and make up a total of 37% of use cases, with the remaining 39% distributed among other industries. Let's take a deeper dive into the top four, which collectively make up 61% of observed use cases.

1. Banking, Financial Services and Insurance (BFSI): 24% of use cases

The demands on financial organizations to digitize products and offerings and make them seamlessly integrated into their offerings are greater than ever. Customers expect their assets to be safe, and they expect to be able to interact with them at will. Meanwhile, the constant specter of cyber threats is growing, pushing organizations to improve fraud detection and risk assessment. As a result, the financial sector is eagerly diving into the possibilities of generative AI.

The financial sector is blessed with a large volume of unstructured data that is readily available, which is a perfect training platform for AI models. Twenty percent of use cases start with data extraction and contextual search. Contextual data extraction is key to driving downstream solutions in fraud and risk assessment. Contextual search and analysis help financial sector companies stay ahead of the news at scale. A number of providers are training and building generative AI tools to support the financial sector. These tools gather as much data as possible and identify trends and shifts in the markets, either through their proprietary platforms or by leveraging a custom language model built on enterprise data (including open source tools like Llama).

However, the most mature use cases are in customer support and personal banking. While early results show strong accuracy from generative-AI-trained chatbots, productivity gains for organizations are still marginal, suggesting a lot of work still needs to be done in this sector.

2. Manufacturing, including Aerospace and Defense, Consumer Durables, Automotive, Conglomerates, Capital Goods, Construction Materials and Tech: 14% of use cases

While the mix of Manufacturing use cases for generative AI is quite high, the percentage of mature use cases (see footnote on p. 7) is one of the lowest at 13%. A deeper dive into our data reveals that, while there is a lot of interest in generative AI applications, most are still being developed and have not gone too far past ideation. What is interesting is that the nature of these use cases is mostly focused on cost management, efficiency, performance analysis and quality control. More mature applications are seeing good results in productivity gains by reducing waste and avoiding defects.

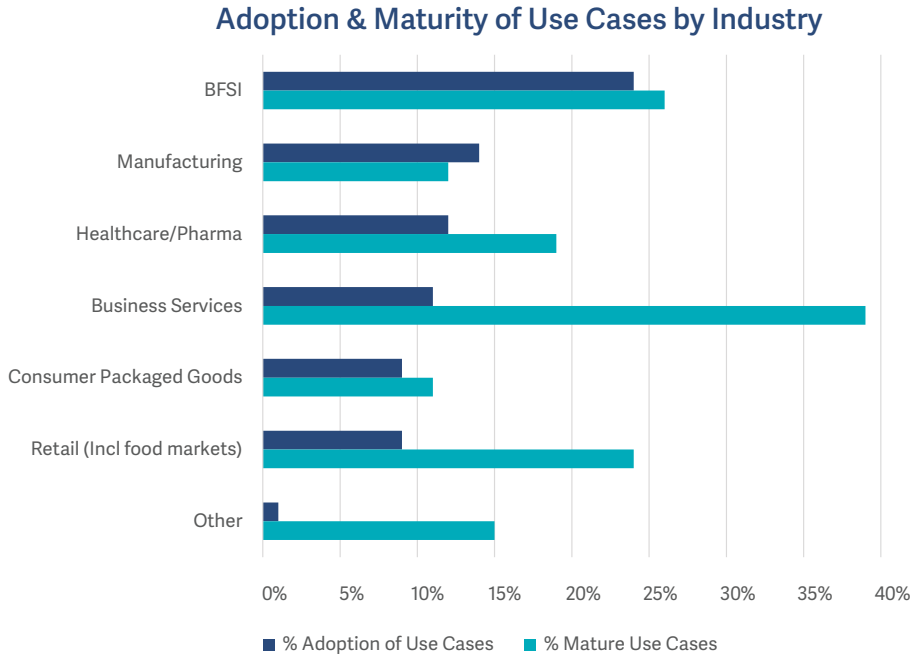

Section 4: Transformative Industries by Use Cases (cont'd)

Sourcing and supply chain management are growing problems in the current complex climate, and so they are top of mind for most manufacturing-heavy organizations. Unsurprisingly, applying generative AI to improve forecasting methods is one of the most mature applications in this sector. Other emerging applications include hyper-personalized product design, digital twins and product metadata generation from unstructured data to create descriptions, categories and classifications.

3. Healthcare and Pharmaceuticals: 12% of use cases

Much like financial services, the healthcare industry is drowning in data, mostly unstructured. At this point, it seems that most organizations' primary concern, in addition to actually delivering care and products, is how to deal with capturing, processing and organizing this treasure trove of data to actually improve services and manage costs.

Similar to financial services, we see greater rates of adoption of data extraction and contextual search use cases, but, in the case of healthcare, the use cases are centered around summarizing a lot of disparate data sources and data types into a human-ready readout. Generative AI output is often a co-pilot to a customer service agent who is getting ready to interact with a patient. A lot of use cases also propose to summarize clinical trials or medical research data.

Due to the domain-specific nature of healthcare, we see a lot of activity from specialty providers who are driving very niche solutions, such as using Generative AI to help fill patient recruitment shortfalls for medical trials or a generative model that produces drug labels. In addition, we see a heavy focus on optimizing the drug and treatment development process with AI.

4. Business Services, including Supplies and Software: 11% of use cases

ISG defines business services as those providing goods and services to businesses, i.e., B2B companies. While the percentage of total use cases identified in the study is fairly low at 11%, here we see the highest mix of use case maturity* at 39%. This is driven primarily by code generation use cases to help optimize the development and the speed-to-market of software. Accounting for 50% of total use cases for this sector, solutions include application migration recoding, tools to help effective deployment and testing solutions.

*Mature or maturity is defined as a well-defined use case with a potential client, a solution in process or developed, and quantitative key performance indicators that can be linked to a defined ROI.

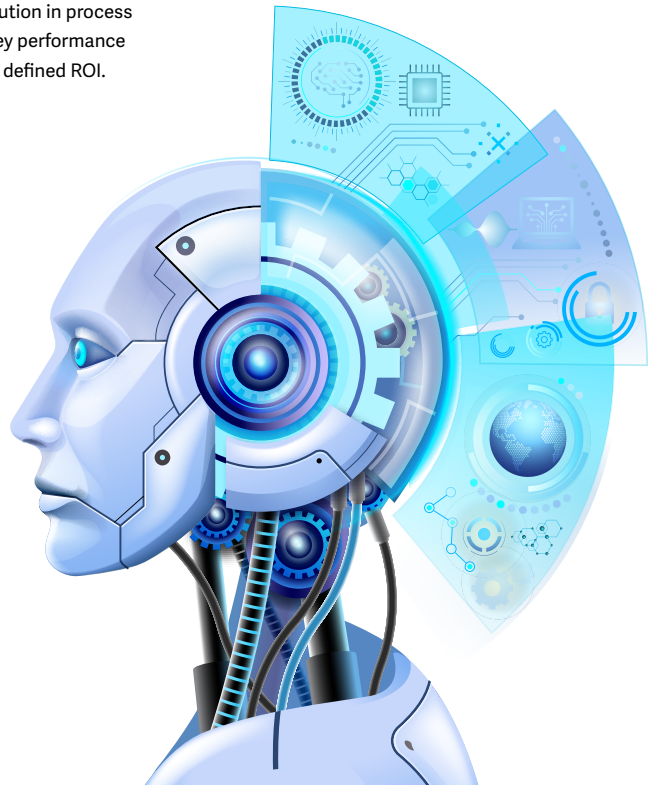

Section 5: Addressing Top AI Concerns

Enterprise concerns about adopting generative AI are defining the nature of the offerings we see from providers. Beyond the logistical challenges of adopting AI, enterprise concerns about the feasibility of ROI could benefit from a joint effort between service providers and enterprise clients.

While providers are actively working to build solutions that mitigate enterprises' concerns, we are still in the early days, and the ethical and legal implications of AI are evolving rapidly. The effectiveness of providers' ability to tackle risks and adapt to change is still somewhat unproven. The following are what our study found to be the top concerns of enterprises around adopting generative AI:

1. ROI Viability:

Is a generative AI investment profitable for an organization? While providers understand that ROI is important, most providers do not directly support enterprises in determining whether and how generative AI is a worthwhile investment for their specific organization.

How is this influencing provider offerings?
Creative providers are supporting organizations' need to validate potential generative AI projects by providing experimentation platforms and solutions that can be tested in a collaborative enterprise/provider proof-of-concept process. The collaboration and brainstorming environment should be designed to foster innovation and test financial and practical viability so organizations can invest with confidence.

2. Legacy Data Quality:

Generative AI can process unstructured data more easily than traditional analytics methods. All models need data, and the more and richer it is, the better. Organizations that have been investing in their data are going to continue to be leaders in AI as well. The question is, to what degree does generative AI allow immature data organizations to leapfrog and work with the imperfect data they already have?

How is this influencing provider offerings?
The majority of providers are revisiting or strengthening their existing offerings for data governance and data management to help support their generative AI solutions. Some are creating generative-AI-focused solutions through either synthetic data generation or by incorporating aspects of generative AI into the design and delivery of efficient warehouse management applications. Overcoming the challenge of legacy data is not a priority for most organizations as they are trying to stay creative, but mature providers are encouraging enterprises to do due diligence on their data as they prepare for generative AI.

3. Hallucinations, Data and Model Control:

One of the biggest concerns about generative AI is the potential for misinterpretation. When generative AI can't generate a correct answer to a question, it starts to invent one in a process called "artificial intelligence hallucination." Leaders also worry about data quality and the potential for models to get corrupted or poisoned – either intentionally or accidentally – resulting in a bad outcome. Therefore, while organizations see the potential of generative AI, they still don't yet fully know how to handle the risks.

How is this influencing provider offerings?
Enterprises are eager to focus on solutions that ensure accuracy and quality of the data used to train generative AI by designing domain-specific, custom-built language models that feed only on qualified internal data. Most providers surveyed (71%) have a designated data quality and hallucination solution. Mature providers are designing an architecture that is tailored to minimize enterprises' industry-specific risks. These providers are developing distinct, individualized applications for targeted use cases with a correction layer to filter relevant requests. These frameworks use stored tracking to identify the connection between a response and its corresponding training material as well as content moderation features that inspect responses for problems.

4. Ethical Concerns:

Most organizations today operate in a much more dynamic environment and generative AI opens the door to potential ethical challenges. Most organizations lack full visibility to all the risks and need support beyond technical development.

How is this influencing provider offerings?
Most of the providers we surveyed (85%) are establishing responsible or ethical AI pillars in their solution toolset. More mature providers (71%) are building ethical AI frameworks that align the provider's pillars with the values of the enterprise organizations. They also ensure that responsible AI principles are built into the design and development process of each

Section 5: Addressing Top AI Concerns (cont'd)

solution. Top providers see responsible AI as a framework designed in partnership with the enterprise, to which they bring the foundation and then customize guidelines, processes and tools within the solution to meet the enterprise's needs.

5. Security and Legal Concerns:

As the legal landscape evolves, enterprises need help achieving compliance to prevent local data breaches and litigation risks. Organizations need to design a digital strategy to define how they will successfully run applications in public vs. private cloud environments. For data not stored in an on-premises environment, enterprises must have a strategy to manage sensitive data.

How is this influencing provider offerings?
There are no transformative security offerings currently in the market designed specifically for generative AI. However, 100% of the providers we surveyed already have a strategy to help organizations establish secure AI operations. Organizations that are less mature will need to lean on providers to establish the foundational architecture for security. Most enterprises and providers working on generative AI projects are treating security and legal concerns as an integral part of the design process.

6. Organizational Copyright:

AI and generative AI tools are built using very data-hungry models. As organizations share their data with providers for the purpose of building generative-AI-based solutions, they need to know how their partners will protect their data from internal reuse.

How is this influencing provider offerings?
Our survey reveals substantial immaturity in this area with only a handful of providers recognizing the need to address copyright concerns for enterprises.

While the landscape is still being built, there is an emerging pattern among more mature providers to identify all enterprise concerns and package a solution that addresses them in a unified framework of governance and oversight. This is going to be a huge help to organizations as they move beyond obvious use cases to more transformative cross-functional solutions.

Section 6: Top Functional Use Cases

Evolution of Use Cases

When we look at any emerging technology, there is often a natural progression for how use cases emerge, and the same is holding true for generative AI. We see a four-phased maturation process – and generative AI adoption is still within the first two phases. Adoption has begun with knowledge management and functional process optimization, which manifest as:

- **Knowledge management** use cases highlight the degree to which AI and generative AI capabilities are enabled by data as well as the process of converting that data into knowledge, which is then dispersed across the organization for reasonable next actions.
- **Functional process optimization** use cases leverage AI to analyze key activities that take place within business functions and optimize them with generative capabilities. These use cases were never a fit for RPA, but with generative AI are finally coming to life as tangible, real opportunities.

As leaders get more comfortable with generative AI, they will begin to imagine more transformative possibilities leading to **product and offerings transformation** (development of true AI-first products) and complete **business transformation** (reinvented operational models that are entirely built around AI).

At the moment, generative AI use cases are, as expected, limited to knowledge management and functional process optimization. Most notably, the heaviest adoption is in knowledge management, implying that these are early days with plenty of opportunity for expansion and increased maturity.

Knowledge Management

Knowledge management tools are by far the most important investment for enterprises today (see findings from ISG study on the future workplace). The foundational use case for knowledge management is **data extraction** – the ability to extract and gather data from vast unstructured data repositories. Examples of pure data-extraction use cases focus on the automation of manual data capture and document processing such as accounts payable (AP) invoices. Given that the extraction capability is in effect a data gathering and enrichment exercise, it is also the starting point for many downstream use cases.

Function	Knowledge Management	Functional Process Optimization
HR	Labor Forecast and Planning Training & FAQ Virtual Assistant	Labor Scheduling Job Description Creation and Management
Marketing	Virtual Assistant Product Metadata Customer Marketing Recommendation Engines	Product Design & Development Media Content Management
Finance	Forecasting, Predictive, Performance Analytics	Creating Compliance Reporting and Filings
Accounting	Fraud Detection Accounts Payable Invoice Processing	Creating Accounts Payable Invoices
IT	Cybersecurity Risk Detection System Performance and Monitoring	Code Generation Autonomous System Migration System Maintenance and Support

Section 6: Top Functional Use Cases (cont'd)

One of the most common of these downstream use cases is **contextual searching and indexing** (16% of reported use cases). Some examples include creating a repository that can be used to store and search new medical research articles, or a model that tags products for efficient cataloging. Further downstream and featuring more generative capabilities are **virtual assistant** solutions (16% of reported use cases), both for customers and employees.

The most mature use cases are mostly in customer support, where organizations can easily expand on existing solutions by infusing generative AI where it makes sense to do so. Another quick win with low risk is an internal ChatGPT-type tool (probably built with newly announced ChatGPT Enterprise) that employees can use to replace less-effective frequently asked questions (FAQ) and training tools. The low mix of mature use cases for virtual assistants (20%) can be explained by the fact that, besides modest productivity gains, the benefits are hard to qualify in ROI terms as their primary goal is to improve customer service and drive employee satisfaction.

Generating intelligent business metrics using data extraction is one of the most common (16%) and mature use cases found by our study, driven by an enterprise focus on ROI and other benefits that significantly impact the organization (57% of **predictive analytics** and 21% of **performance analytics** use cases can be defined as mature; see note on p. 12). These include forecasting and demand planning as well as models that help manage resources and optimize costs.

Recommendation engine use cases, which take the data extraction and analysis enabled by AI models and generate a suggested next action, are also fairly common. These consist of 14% of surveyed use cases, of which 23% can be described as mature. The most common use cases include pricing, personalized product and services recommendations, customer segmentation and targeted sales.

Our study has identified recommendation engines that leverage generative AI to extract data and present recommendations as reports. However, we also see solutions that are structured as true co-pilots. A diagnostic co-pilot, for example, can help a doctor pick the best treatment for a patient or help underwriters identify risk and design the best insurance solution for a client.

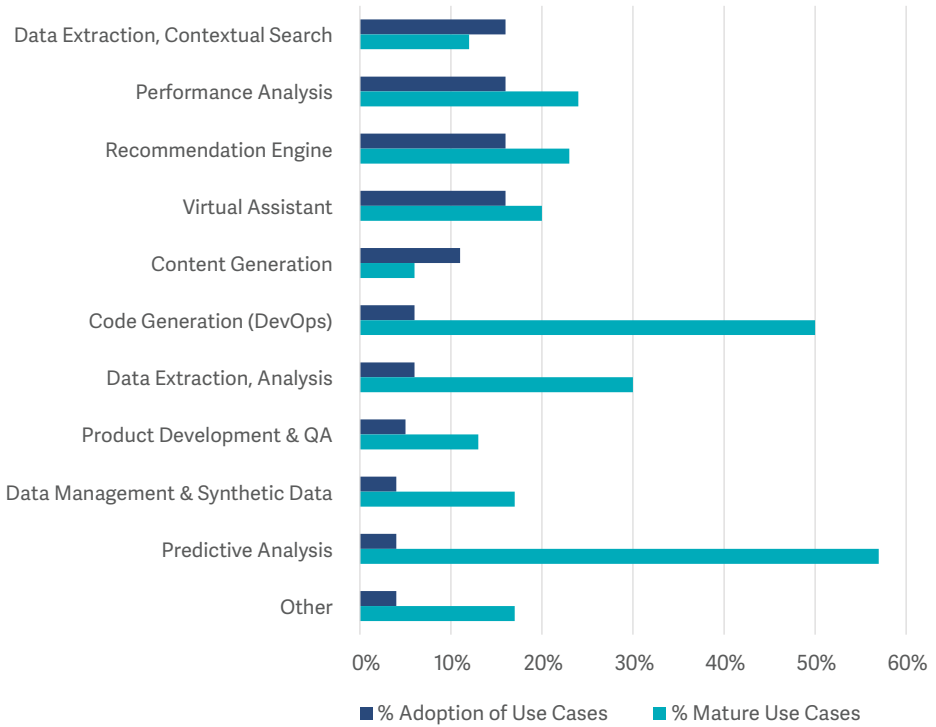

In addition, we see some early adoption of **synthetic data generation**, but the application is still limited.

Section 6: Top Functional Use Cases (cont'd)

Functional Process Optimization

With providers driving the solution market, we are seeing limited functional use cases for process optimization because these require collaboration with clients. Our study shows that, while generative AI solutions have the potential to transform key functions within organizations, their design is still being determined.

Demand Generation, Marketing and Sales. Our study found that **content generation** for marketing and product metadata is a common use case (11% of surveyed use cases) because it can significantly improve productivity by saving time and driving more consistent, higher-quality results. We also saw use cases in media content creation, distribution and management. All of this points to the fact that demand generation functions are hungry for optimization and are accustomed to technological change.

IT and DevOps Code Generation and Application Development and Maintenance. These use cases make up only 6% of the observed use cases, but 50% of them can be described as mature with a clear relationship to ROI. We also saw a lot of great use cases for generative AI to help with system migration by recoding processes from one language to another. This is an encouraging start. Organizations are spending a lot of energy, often through partnerships, to maintain their systems, including avoiding cyber security risks and managing support. Similar to customer service, providers should consider how generative AI can be infused into existing support systems and processes for enterprises' IT solutions.

HR and Labor Planning. Besides onboarding and cross training, the human resources (HR) function has the lowest visibility in terms of use cases in the study. Yet there is a lot of labor-intensive content being managed and created in HR that would benefit from generative capabilities. The creation of job descriptions and the need to align and adjust across multiple platforms is very similar to marketing. HR should also look at improved forecasting and planning tools to help with scheduling problems. Finally, with AI transforming more roles within organizations, HR should take the lead in understanding how this shifts the labor market and should lean on providers to fill the knowledge gaps.

Finance, FP&A, Treasury and Accounting. Financial functions are going to be leaning on generative AI to help drive forecasting and predictive analytics. However, adoption that will impact process optimization within this function is still limited. Organizations have to create and manage a lot of compliance-related reports and filings which require human-intensive content generation. In addition, there is a need to manage the flow of content to investors and shareholders. All of this creates a great opportunity for generative AI solutions, revealing another area of opportunity for future business process improvements.

While the study showed a lot of innovation, there are still plenty of untapped opportunities. Enterprises should continue to stay creative. Virtual assistants are a great start, but we expect that we will continue to see organizations innovate, leveraging generative-AI-enabled knowledge management for customer-facing applications and smart connected products.

*Mature or maturity defined as a well-defined use case with a potential client, a solution in process or developed, and quantitative key performance indicators that can be linked to a defined ROI.

Section 7: Solutions Spotlight

Providers are playing a large role in not only developing use cases, but also building out complete solutions that target specific industries and business functions. With most organizations still in an exploratory stage, the solutions proposed by providers are giving us a preview on how generative AI will transform the way work is done. Top solutions we're seeing include:

- **Call and customer experience center solutions delivering a full suite of tools.** These tools could include customer insights, conversational analytics, call intent prediction, agent performance analysis, virtual assistant capabilities and conversational sentiment analysis. These solutions effectively bring together the best technology to deliver a one-stop-shop that better manages, monitors and analyzes call center operations. This kind of complete package solution complements organizations that are starting their generative AI journey with customer experience tools.
- **Application development and modernization solutions, leveraging developer-friendly generative-AI-based tools, data migration and code generation.** The study findings clearly show the impact generative AI is having on application development and all aspects of coding. Creative providers are thinking beyond quick wins and are looking for ways to completely overhaul the way IT and application maintenance and development work with generative AI at the center. An example of this is helping IT quickly and easily migrate between platforms by using generative AI to recode processes.
- **Experimental generative AI lab environments for enterprise exploration, learning, education and collaborative solutioning.** This solution is less about technology and more about provider services and partnerships which explore what generative AI can do for an enterprise. What we find is that providers are as eager to explore the challenges that enterprises are facing as organizations are to solve them. Lab solutions are great for innovation with limited risk. The key to success here is to ensure that organizations are focused on ROI. To be successful, enterprises and providers must work together to align on clear goals from the get-go.
- **Proprietary generative AI solutions customized around internal processes and knowledgebases.** This solution is popular with organizations as a quick win from generative AI capabilities with relatively low risk. It is also easy to expand by feeding more sources and more data into the model. By doing so, organizations can extend generative AI's capabilities into training, onboarding or adding features like a virtual assistant or a co-pilot for customer interactions and product design. Successful providers are picking up on this trend and building customizable tools that can be set up, trained, and brought up and running very quickly, creating an approachable use case that generates excitement for other solutions across enterprises.
- **Generative AI platforms.** In addition to creating customized solutions, a number of providers are building tools specifically designed to support internal generative AI development to better support clients. These user-friendly tools could evolve into client-facing capabilties down the road. Overall, there is a lot of diversity in this space in terms of tool usability and complexity as well as practical applications beyond internal use.

With more use cases in development, it is likely that we will soon see current work-in-progress proofs of concept (POCs) grow into complete end-to-end solutions. Some likely areas are product design and development, financial planning and analysis (FP&A), accounting, public relations (PR), customer relationship management (CRM), marketing campaign management and communications, as well as many others. There is still plenty of room to completely redesign how organizations perform routine tasks, including by critically looking at solutions already in place. Stay alert to upcoming developments in this space.

Section 8: Strategic Recommendations

Industry-Specific Action Steps

Based on what we learned in the market, we are calling out some practical recommendations to avoid potholes. Here are some actionable strategies that will help you succeed on your generative AI innovation journey:

- **Effective Benchmarking:** Learn from the market by benchmarking industry best practices and the newest trends in generative AI use cases. Explore real-world use cases in your area of interest and, where possible, review success stories in detail with the parties involved to document the potential they have for your own operations. If you're not certain about investing in generative AI tech, the best way to build confidence in its capabilities is by seeing the tangible ROI other enterprises are achieving.
- **Strategic Resource Planning:** Amidst the rapid adoption of generative AI tools, it's easy to lose focus and misjudge your capabilities. We recommend you develop a structured budget, allocate resources for swift engineering support and analyze pertinent use cases to achieve the best outcomes.
- **Avoid the Hype:** Keep your focus on your mission (what business value you aim to achieve) and your chosen use case. Select the use cases that best support your goals. Stay rooted in reality, and avoid being swayed by market hype.
- **Leverage Off-the-Shelf Products:** Stay well-inform on the generative AI market so you can leverage the newest tools to provide an instant positive impact to employee productivity. Examples of these efficiency-boosting tools are ChatGPT Enterprise, Bard for Workspace accounts, Pictory, Jasper and Murf. They are easy to adopt, and they do not require immediate integration into the overall enterprise architecture. Just make sure you:
  1. Train your employees on confidentiality and data breach risk,
  2. Vet the tools for use in your environment by evaluating them for legal and regulatory compliance, and
  3. Dictate a robust security and compliance policy.
- **Continue to Innovate:** Even with current efforts underway in the market, we believe there is a still a lot of potential for AI to change the way companies perform routine tasks. Encourage your providers and leadership to look critically at their processes and determine what might be possible. Generative AI is not another automation fad. Organizations need to look beyond robotic process automation (RPA) and imagine new possibilities – some which will likely involve entirely reinventing the way things are done today.
- **Select Collaborative Partners:** As technology advances swiftly, even tech-savvy enterprises should outsource in areas of limited expertise. Consider if choosing an experienced partner offering a scalable end-to-end solution for comprehensive business transformation is right for you. Alternatively, your organization may find a better fit in a specialist partner that customizes tools to address your distinct use-case challenges. An easy approach to avoiding the "same ten partners" problem is to ask yourself these three questions:
  1. Who is doing it best?
  2. What are they doing?
  3. What are the results?
- **Build Your Infrastructure and Roadmap:** Currently, the market lacks robust architecture maps for clear generative AI integration. To navigate this challenge, prioritize creating structured architecture maps to guide your implementation. Continue reading to learn our detailed guidance on the optimal architecture and steps for adopting a transformation roadmap.

Section 8: Strategic Recommendations (cont'd)

Proposed Reference Architecture

The biggest gap we saw in the market was a lack of direction in terms of architecture. When we say architecture, we're not referring to IT architecture, but rather an architecture specific to adopting generative AI across your current business landscape. Our research found that most providers are building solution-specific architectures with a focus on the technology landscape. A few mature providers have developed generative AI concept-based structured maps. We used the information we collected to define the right structure for enterprises to use in building their strategies. Of course, this is an emerging market, and we expect this layout to evolve based on new use case success stories in the open market.

Our vision for the recommended architectural layout aims to strategically position generative AI as an industry-agnostic selection of relationships, grouping components as conceptual building blocks. At its core, a cognitive infrastructure forms the base foundation on which everything else is built. This is a conceptual infrastructure, filled with hybrid cloud technology platforms, security policies, edge computing, and any other purpose-built assets that serve and enable generative AI.

With a strong foundation in place, next we envision robust pillars for support: the AI control plane on the left and the data layer on the right:

- Think of the AI control plane as every button you'd want on an interface dashboard for your AI solution. Your AI budget, resources, ownership, large language model (LLM), AI ethics structure, value realization frameworks and governance would be among the options you would want to manipulate in your control plane.
- The data layers are comprised of data lakes and all the structured and unstructured data owned by your organization in its raw form, which is then enriched and cleansed into transformed data. These data feed LLM operations and the model resources an organization can leverage to quickly scale AI capabilities. Understanding your data plays a key role in defining your inputs for upcoming proofs of concept (PoCs).

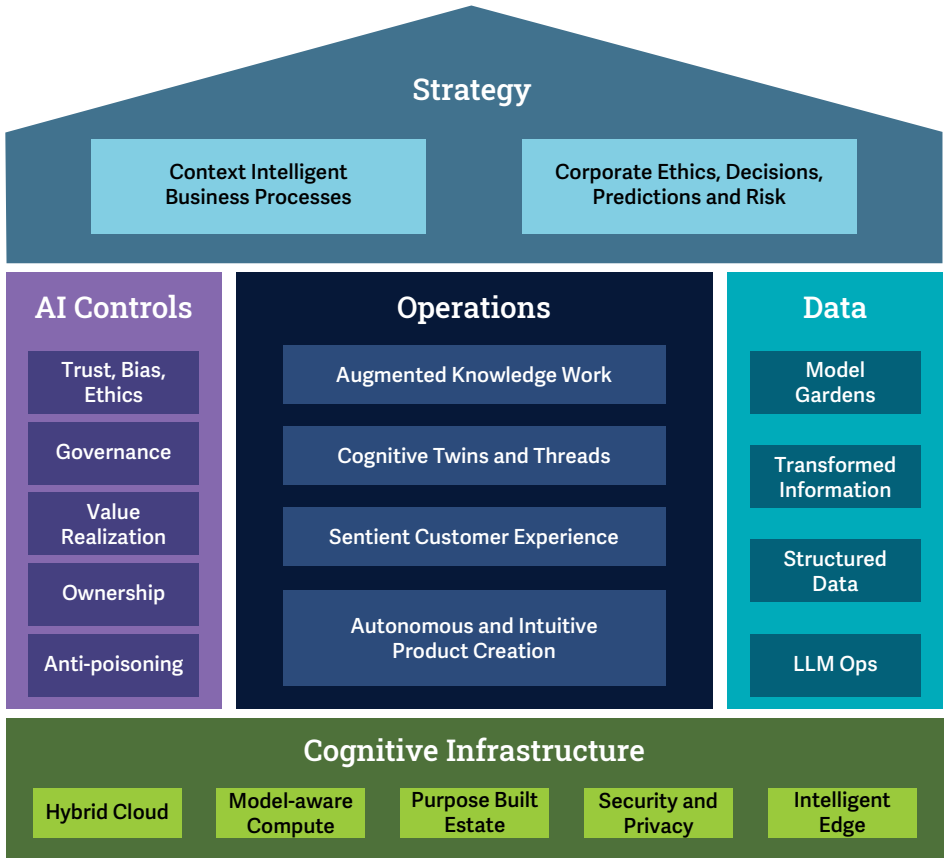

Note: This is NOT a technology architecture map.

Section 8: Strategic Recommendations (cont'd)

Within the bounds of these pillars and constructed atop the cognitive infrastructure foundation, we then address your strategy and operations framework, which you could think of as the brains and the brawn (or the mind and the mechanics). We've positioned these two frameworks in between the data and control pillars because an enterprise's data and controls must form guardrails for what an enterprise does (strategy) and how (operations).

- The topmost strategy element involves ideation by the C-suite to form a corporate vision, which serves as the cohesive spearhead that leads all other discussions. The standard ISG top-to-bottom approach is used here, as the success of any generative AI integration requires senior leadership guidance and buy-in. Note that corporate ethics and social responsibility, such as DEI initiatives and company values, lie in the strategy building block, while AI ethics, such as a bias toward Android over Apple user preference, is governed within the AI control plane as those ethics relate to AI cognitive functions only.
- The lower part of this central framework captures the operational functions that serve the functional activities, such as intelligent AI chatbots for customer service. This layer also covers the integration of generative AI with digital twins and digital threads, autonomous product creation, and contextual data extraction. This layer is where the vision and value of generative AI is executed and realized.

By grouping and accounting for each of these segments of generative AI adoption, we create a robust AI architecture that enables disparate business functions and processes to work together more effectively. The architecture map shows which parts of the process need to be in lockstep, such as governance and ethical or legal constraints, and where they plug into other parts of overall picture, such as the overall executive strategy and your cognitive infrastructure.

For any organization undertaking generative AI adoption at an enterprise scale, an architecture map is not just valuable but necessary. When the entire architecture is laid out visually, the map serves as a valuable tool for troubleshooting and diagnosing issues that may arise during the development of the generative AI solution as well as during operational phases. Having a robust architecture makes it easier to pinpoint the source of a problem and address it effectively. The architecture can also serve as a framework for effective communication among team members and stakeholders. Overall, an architecture map should serve as a common reference point, enabling individuals from diverse technical backgrounds to discuss and collaborate more efficiently.

ISG will continue to develop and research the most adopted and successful architecture layouts and continue to challenge the market to innovate and evolve in real-time. As we see providers' progress towards specific areas, such as security and iterative innovation, we will share where we see organizations successfully bridging market gaps and advise clients on new best practices. In the transformation roadmap section to follow, we will offer guidance on the order in which we recommend addressing each of these segments.

Transformation Implementation Roadmap

Much like most new technology, the generative AI phenomenon is creating a lot of excitement and a need to move quickly past exploration into the realization of first use cases. However, as we have identified in this report, the potential for AI to transform business operations from the ground up is serious enough that a more thorough strategic approach to realization is necessary. It is important to acknowledge the importance of both viewpoints and to act to accommodate the possibility of achieving both.

To achieve results with generative AI quickly, it is important to get a few key things right. First, create an inventory of use cases and evaluate them by potential ROI as well as the maturity of architecture and data that will be needed to support them. Secondly, apply risk ranking (the level to which you expose your organization to legal, security or ethical risks). Risk ranking is needed because this landscape is still evolving, so it is important to be honest about what can create push back internally and to avoid any obvious pitfalls. Thirdly, consider what level of visibility would be best for the first prototype. For example, an internal ChatGPT solution that provides an easy frequently asked questions (FAQ) or training platform for staff is a great first generative AI use case. This example

Section 8: Strategic Recommendations (cont'd)

use case gives the broader organization a taste of what is possible so that they can consider other possibilities in their own departments.

Armed with this list, you can focus on separating your priorities into both short-term, easy-win investments or intermittently funded investments for use cases that better align with your overall organizational strategy going forward. However, it is important not to let immature architecture be a roadblock to early wins. Use cases that show a lot of promise but may be too difficult to address with in-house generative AI maturity should be addressed collaboratively through a partner network. This allows the organization an opportunity to test the right ideas while still being strategic in establishing a good foundation to move into more transformative use cases. It also provides quick learnings on how AI does and does not fit into an organization.

Companies should be cautious about diving headfirst into generative AI strategies without expert guidance. Instead, they would be well-advised to lean on partners and system integrators who have already honed their skills in implementation. By doing so, companies can glean valuable insights, not just on model development, but also on the strategies these integrators are employing for other enterprises. This collaborative approach offers a unique opportunity to leapfrog over traditional designs and business models. Companies should be poised to construct their strategies atop foundational models provided by these integrators. We anticipate these foundational models to be predominantly domain-driven.

This initial investment removes pressure and allows time to thoroughly assess your current capabilities in data, infrastructure and governance (AI control plane). These separate workstreams, together with learnings from the initial prototypes allow you to formulate an AI playbook and establish a blueprint for scaling AI across the organization. Enterprises will then need to put together an empowered strategy realization office (SRO) that can navigate through these threads. SROs make the case for investments, understand the strategic intent of the AI initiatives, have the authority to make pivots, accelerations and course corrections, and act as the guarantors of value creation and delivery.

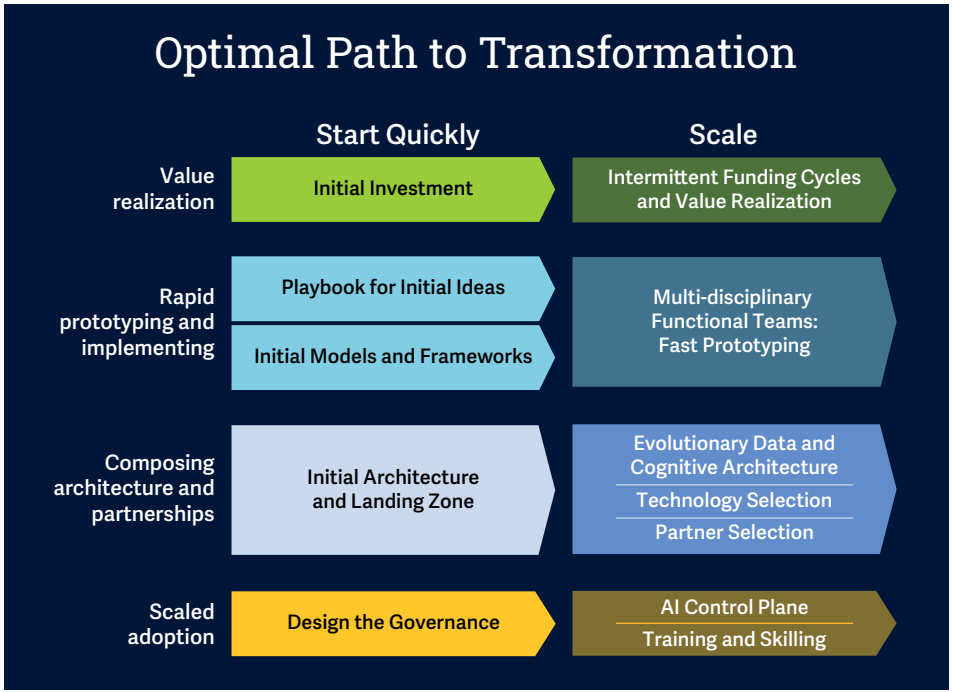

Section 9: Conclusion

In our study, we've found that generative AI is likely following a similar maturation as many other emerging technologies with four distinct phases of technology use and maturity for enterprises:

1. Managing knowledge
2. Improving functions and processes
3. Building AI-first products and offerings
4. Re-imagining businesses with AI-centric operating models

And since this technology is still relatively new, it is perhaps to be expected that ISG has placed the market exclusively in the first and second phases of this maturity process.

Use cases first emerged centered around knowledge management driven by providers. New use cases are emerging related to functional process optimization, but these use cases still have many gaps and will only mature with more and deeper collaboration between enterprises and providers to meet business needs.

We expect that leaders will begin getting more comfortable with the generative aspects of generative AI. However, leaders and early-adopters won't be able to move forward without addressing major security, legal and ethical challenges. Only after some of these challenges are addressed will the market be able to take on more innovative and deeper generative AI usage, such as the development of true AI-first products. ISG expects these AI-centric products to disrupt the market, leading to the ultimate evolution of reinvented operational models for whole business entirely built around AI, as well as emerging AI-first businesses comparable to businesses born in the cloud from ten years ago.

These evolutions in usage will drive governments, consumers, regulatory and legal entities, businesses and providers to advance their thinking on the topic of generative AI. We expect to see this play out in deepening maturity of and evolutions to our proposed enterprise architectural map and transformation roadmap as well. Looking forward at the market's trajectory, the outlined layout of generative AI integration and its impact on the enterprise holds vast potential.

Our framework guides organizations toward harnessing generative AI's transformative capabilities, while our roadmap functions as a practical guide for those headed into a transformational journey.

To conclude, this report is a robust resource aimed at empowering C-suite leaders with insights, strategies, and actionable steps to harness the substantial potential of generative AI. But this landscape is changing quickly. ISG will continue to publish findings, research and recommendations in future reports as the market evolves. The path ahead is characterized by innovation. Armed with our findings and recommendations, we hope that organizations are better prepared to foster impactful change and gain a competitive advantage in this evolving landscape. ISG is well positioned to help enterprises strategically select, source and implement generative AI with the right partner.

Contact us

ISG can help you achieve measurable business value from generative AI.

Click below to learn more.

Website:
Artificial Intelligence Advisory

Email:
Contact@isg-one.com

Authors

Olga Kupriyanova Principal Consultant

Olga is a Principal Consultant in ISG's Americas Consulting practice on the Cognitive & Analytics team. Olga supports our advisory capability in developing digital solutions with special focus on data and analytics. Olga's extensive knowledge of analytics and data engineering framework combined with hands on experience in complex transformational projects results in unique insights invaluable for effectively assessing the data analytics solutions for ISG's clients.

Ryan Hamze Principal Consultant

Ryan is a Principal Consultant in the Digital Engineering team under the Americas Consulting practice. He is a seasoned thought leader and delivers strategic roadmaps to clients across different industries and advises on advanced technology trends, enterprise architecture and digital transformation strategies. With over 15 years of IT leadership and project management experience, he recently joined ISG in 2023. He is actively involved with clients through the strategy realization office and contributes to the internal ISG research team focused on emerging technology and innovation with a focus on generative AI.

Contributors

Alex Manders
Partner

Gowtham Kumar
Assistant Director & Principal Analyst

Jan Erik Aase
Partner & Global Head of IPL

Namratha Dharshan
Chief Business Leader of India Research

Prashant Kelker
Chief Strategy Officer & Lead of Americas Consulting

Shafqat Azim
Partner & Americas Lead of Digital Transformation Practice

Shriram Natarajan
Director

Taylor Card
Marketing Manager

isg-one.com

ISG (Information Services Group) (Nasdaq: III) is a leading global technology research and advisory firm. A trusted business partner to more than 900 clients, including more than 75 of the world's top 100 enterprises, ISG is committed to helping corporations, public sector organizations, and service and technology providers achieve operational excellence and faster growth. The firm specializes in digital transformation services, including automation, cloud and data analytics; sourcing advisory; managed governance and risk services; network carrier services; strategy and operations design; change management; market intelligence and technology research and analysis. Founded in 2006, and based in Stamford, Conn., ISG employs more than 1,600 digital-ready professionals operating in more than 20 countries—a global team known for its innovative thinking, market influence, deep industry and technology expertise, and world-class research and analytical capabilities based on the industry's most comprehensive marketplace data. For more information, visit www.isg-one.com.