The executive's guide to generative AI.

Google Cloud

Table of contents

Introduction

Generative AI marks one of the most significant technological shifts in history.

Its impact on individual and business productivity can be significant, with the potential to rival the advent of the internet or the mobile device. Indeed, among organizations considering or using AI, 82% believe it will either significantly change or transform their industry.[1]

What makes generative AI different from other forms of AI that have come before is its ease of use in helping solve everyday problems in people's personal or professional lives. Anyone who knows how to ask a question of a search engine can use everyday language to interact with a generative AI chatbot or virtual agent — getting it to answer questions, create content, produce images, summarize documents, and much more.

Even better, a single generative AI platform can deliver solutions for multiple use cases, creating a network effect. As the number of users and applications increases, the model is exposed to more data and becomes increasingly accurate and useful — which in turn encourages more users.

Organizations that use generative AI to speed up, automate, scale, and improve business processes stand to reap big benefits. According to McKinsey & Company, generative AI's impact on productivity could add between $2.6 trillion and $4.4 trillion USD annually to the global economy.[2]

No technology ever takes away the fundamentals of your value proposition and the core value chain in your industry. At the end of the day in healthcare, for example, you're still trying to improve patient care.

What can change, though, is how you use this technology to enable your teams to improve core offerings, and how you solve fundamental problems that get in the way of delivering them. Indeed, with the right tools, you could even identify and deliver new points of difference.

1. Google Cloud Gen AI Benchmarking Study, July 2023
2. McKinsey & Company, The economic potential of generative AI, 2023

This guide is for business leaders who want to kickstart their organization's generative AI journey.

In the first chapter, you'll learn what generative AI is, what it is capable of, and the impact you can expect when applying it in a business context. The second chapter provides a step-by-step guide for getting started with generative AI, with recommended best practices from Google Cloud's AI experts. Next, we dive into real-world examples of pioneers across industries who have adopted generative AI to work smarter, deliver value to customers faster, and unlock new revenue channels.

With generative AI technology moving so fast, it can be overwhelming. We're here to help you chart the right path.

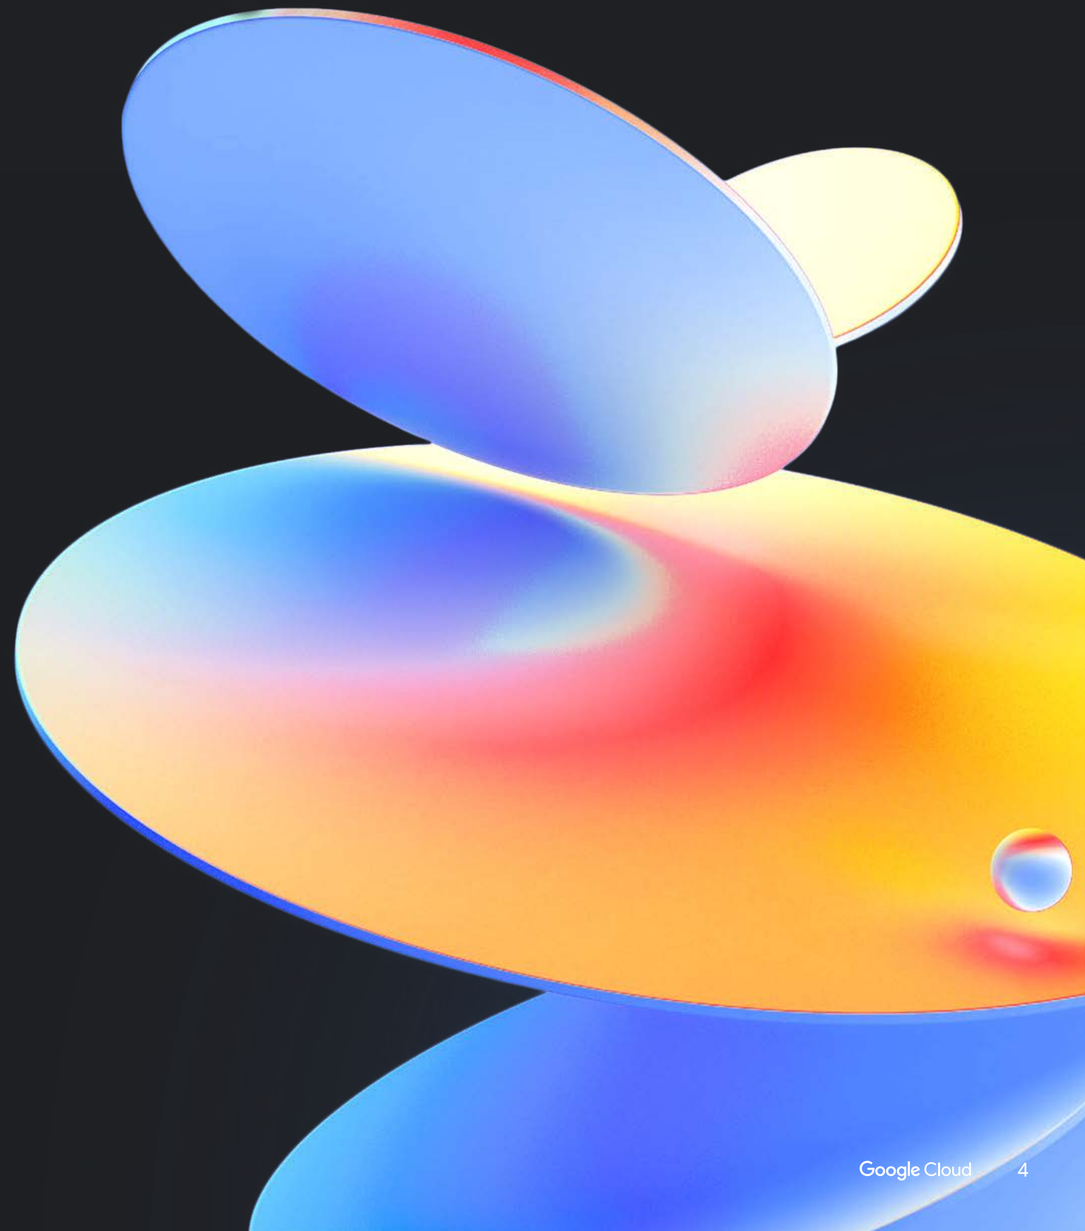

Chapter 01

A quick primer on generative AI

This generation of technology is the most approachable and flexible technology that computer science has ever built. It can solve problems that the algorithm wasn't purpose-built for. And, any business user can quickly experiment and get feedback to shape its use for their particular business problem."

Philip Moyer

Global VP, AI & Business Solutions, Google Cloud

Every day, people in your business spend time and energy digging for information to make decisions, serve customers, and move your business forward. Informed decisions require information, and collecting the right inputs can take time.

Suppose you need to know how ad spend impacts customer perceptions. Or you want to see the trends in a competitor's patent filings, R&D investments, and technology acquisitions. Information like this exists somewhere in your organization, often across a bunch of places. You need it to decide the next step — and, to obtain it, experts need to be convened, research completed, and the information compiled and synthesized. If you have a follow-up question, the whole process might have to start again.

From executives digging into strategic trends to salespeople creating product demos or new employees with benefits questions — everyone in your organization can relate to this frustration. But that's changing.

Imagine giving each person in your company not only a personal assistant, but an expert in every piece of data relevant to their job and, indeed, potentially every piece of data across your whole organization. With an assistant like that, impatient moments of indecision would dwindle. Everyone would be empowered to spend less time waiting and more time doing.

With generative AI, this is achievable. And it's just one example of the many disruptions this technology has unlocked. Always-on coding collaborators. Brainstorming assistants to draft and iterate content. Personalized self-education on any subject. Human-like interactions with customers wherever they need you, for whatever reason. Generative AI is all these things and more. In time, it will affect almost every aspect of every business.

Foundation models are the engine powering generative AI.

Generative AI applications are powered by foundation models, which are trained on vast amounts of content. For example, Large Language Models (LLMs) are one type of foundation model, trained on text or language. Other multi-modal types can be trained on images/photos, video, music, software code, medical information, or cyber security data. But access to models alone won't position your business for success.

Foundation models are best thought of as probability engines that can be nudged and shaped by human input. Because they are probabilistic, they are fundamentally different from traditional software paradigms. When today's apps need to look up product prices or validate customer information, they use deterministic functions that call a database. In contrast, foundation models use the patterns they learn during training and tuning to calculate the most probable output, such as the most likely answer to a question or an accurate caption for an image.

Since they're not confined by rows and columns in a database, foundation models are extremely powerful. They're often capable of performing many downstream tasks — such as Q&A, summarization, or open-ended content generation — with little or no additional data or tuning. Yet they can also be expensive to train and run, prone to inaccurate outputs, and difficult to work with.

For these reasons, generative applications aren't reducible to generative models. Your intelligent applications will need to mix probabilistic foundation models with traditional, deterministic (in other words, constrained) programming. Deterministic models are constrained to what can flow out of them, and are limited by endless options that have to be pre-established.

Traditional AI is purpose-built for the task at hand. It is all about optimizing and tweaking existing processes such as forecasting specific patterns that have been pre-determined by humans. It's why traditional AI can be used to automate discrete, standardized processes in specific areas of the business, such as customer service.

In comparison, generative AI models have emergent capabilities where they can perform multiple tasks, even if they were not explicitly trained for them during the instruction tuning phase. It is this ability to multitask, plus the freedom afforded by the prompting interface, that enables these models to perform in a wide variety of use cases.

Core capabilities and applications

At the core, generative AI has four capabilities:

- Creation
- Summarization
- Discovery
- Automation

And it tends to excel in four applications:

Chat

It's no coincidence that generative AI has gained rapid popularity and adoption through simple chat interfaces. Chat is a natural way to interact with powerful generative AI models. You can use it to improve customer interactions, enhance product capabilities, train employees, and more.

Search

By combining generative AI capabilities with search, you can anchor on a knowledge base — either internal or external — for more tailored and targeted interactions. Using generative AI for search can help eliminate hallucinations by sourcing information from a factual knowledge base.

Generate content

The ability to generate high-quality text, images, speech, and code has enormous potential. Whether it's speeding up processes or helping employees turn ideas into output faster, generative capabilities can be deployed into products, tools, and workflows.

Associative reasoning

This is the ability to suggest associations in information based on context, frequency, or proximity. For example, generative AI could identify the three most common reasons that a call center interaction ends negatively by parsing the large amounts of transcribed conversations.

Read our glossary on generative AI terms and concepts.

Chapter 02

The step-by-step guide to getting started

To enable your foundation models to get very smart, very fast, pick one functional domain in your business and experiment in that area.

By clustering use cases into a domain, you can start with one, and then when that's working, you can naturally expand to your second, third, and fourth use case within the same domain. The more you throw at a model, the smarter it gets.

Take customer service, for example. First, let's say you give your call center agents a generative AI tool with a conversational interface to use when taking phone calls. Customers might call in and say, "My credit card no longer works." Or "I lost my passcode." Or "When I go traveling, how do I make my cell phone work internationally?" Your agents can answer these questions naturally, querying the generative AI interface to provide an answer using similar wording to how the customer asked the question.

Then, you can look at the aggregate of those queries and ask generative AI, "What questions are being asked most frequently? What's our response time? What responses do we provide?" Now, you're moving from just answering customers' questions to summarizing their data.

Third, using this summarization, you can prompt generative AI to take the most frequently asked questions and compare them to the FAQs on your website. You can then ask it to generate answers to any questions not answered on your website, so you can publish them.

In this example, generative AI solves three use cases that enhance the customer service domain — answer the question, summarize what's being asked, and produce answers to these questions. With each use case, the model gets smarter.

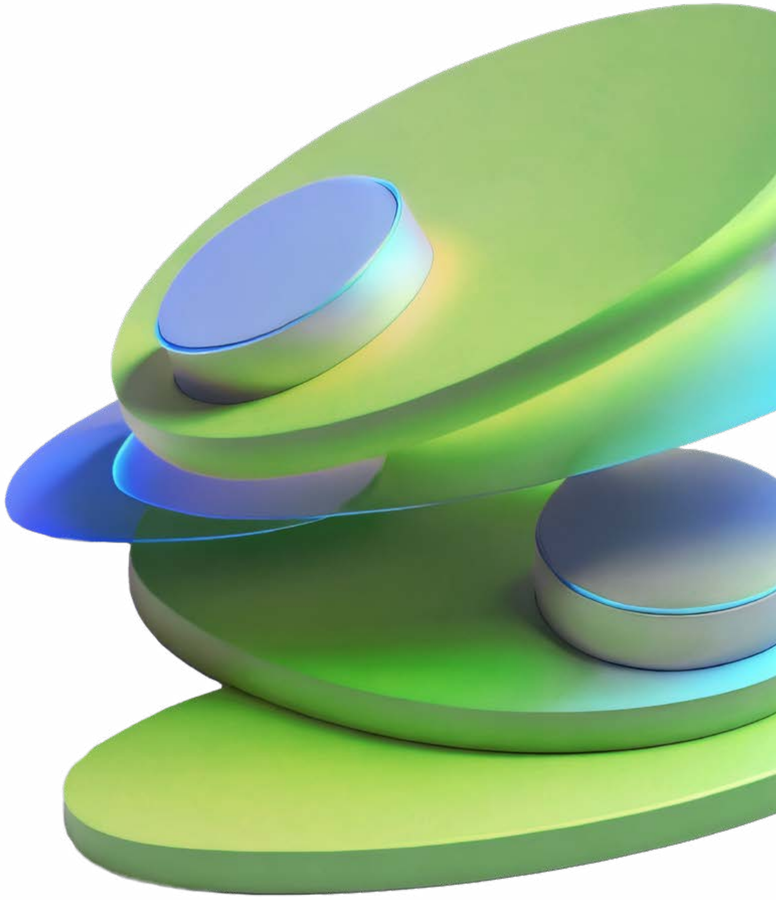

10 steps to launch your first use case in 30 days

Follow these steps for a fast, easy, and low-risk approach to starting your organization's generative AI journey.

We've included KPIs to help you demonstrate impact to your leaders, foundational processes to scale across domains, and recommendations to allow your teams to experiment in a safe, internal setting so they can gain experience with the new technology.

Step 1

Identify a specific domain

Step 2

Select a persona

Step 3

Determine data sources

Step 4

Create a tiger team

Step 5

Define your goals

Step 6

Design prompts with tiger team

Step 7

Build UX and UI

Step 8

Expand to more users

Step 9

Build an LM operations plan

Step 10

Expand to more use cases

Step 1 of 10

Identify a specific domain

Choose a domain in your company that could benefit from generative AI, such as customer service, patient intake, corporate actions, or marketing content.

Key questions:

In which areas of the business do employees spend significant time on repetitive tasks?

Is there a process or part of a role that is, in effect, already standardized (e.g., having to do a particular step or answer a question every time)?

Where do employees get stuck in the creative process (e.g., writer's block or creative block)?

Would an incorrect response or hallucination cause harm?

Which areas of the business offer the lowest-risk environment for initial use cases?

Do you have a large corpus of data that you want to activate, or make more useful?

Is there an area of the business where employees consistently have to search for existing information using internal knowledge bases and/or external search engines?

Step 2 of 10

Select a persona

Determine which job category or function within the chosen domain you want to make more productive.

Consider these three factors:

Look at job roles that are hard to retain and hard to hire.

Such roles are often repetitive and offer little career advancement. Automating these tasks can free up employees to focus on more strategic work.

Find opportunities to automate repetitive, tedious tasks that are necessary to generate revenue.

For example, the multi-trillion dollar pre-authorization industry in healthcare is incredibly frustrating for patients. It can take hours or even days to get authorization for a simple MRI or specialist visit, and the process is often paper-based. Investment memorandums are another example in which the same information must be gathered over and over again. Generative AI can help automate these tasks so employees can focus on more strategic work.

Make safety and compliance ambient.

Many industries must meet strict requirements to be compliant. In life sciences, for example, every claim about the efficacy of a medicine must be reviewed by an attorney to ensure the language is compliant. An attorney must then review the claim to ensure certain clauses are included in every instance of the fine print. This can be time consuming and repetitive, but is an essential part of the process for making medicine publicly available. Generative AI can help automate critical tasks that help organizations improve accuracy and reduce risk.

Step 3 of 10

Determine the data sources the persona needs to be productive

Your generative AI model will be trained on the data you gather. This data should be honed for the specific business or domain-level problem it is trying to solve, and accessible through enterprise data sources.

For example, if you select a marketing manager as your first persona, it's important to understand the specifics of their job role. Say they are responsible for creating digital campaign content like ebooks and reports. These assets are promoted and used to capture leads via online forms. If the individual fills out the form and opts in to receive outreach about related topics, the lead is entered into a marketing automation program and scored based on pre-determined criteria set by marketing and sales operations teams. If the lead meets the criteria for qualified sales opportunities, it may also be handed off to the sales team.

This individual job role requires several data sources, including:

- A word processing tool like Google Docs to write, edit, and collaboratively review the copy
- A design tool to lay out and format the final copy
- A web platform to publish the content online
- A marketing automation tool to track and measure marketing engagement, tasks, and workflows
- A CRM like Salesforce to ensure sales, support, and marketing are coordinated across all user touchpoints and customer interactions

By starting off with the right data to feed into and fine-tune models, your organization will be able to:

Mitigate hallucinations

AI models are trained to give users what they want, which means they occasionally spit out made-up answers that sound convincing and can be hard to spot. To avoid this problem, you can ground responses in specific data rather than relying on the LLM alone.

Enhance the explainability of AI

Generative AI models can be complex, and the ‘thinking’ that an algorithm uses to produce an output isn’t always clear. Explainable AI is like a sliding scale — there are degrees to which you can explain or reliably steer an LLM’s behavior.

Step 4 of 10

Create a three-person tiger team

Include people from both business and technology:

Individual from the business

Responsible for detailing the job requirements, workflows, challenges, and needs of the day-to-day tasks executed by the chosen persona.

Prompt engineer

Responsible for translating the business persona's needs, actions, and output into prompts for the generative AI model(s).

ML operations lead

Responsible for building and operating the application in production.

Step 5 of 10

Define your intentions, objectives, and the output you are trying to achieve

Ensure you have a human in the loop to oversee the first use cases and provide oversight.

Keep in mind, the value of a generative AI project can come from a number of sources. There's direct business value, incremental value of generative AI over legacy systems or traditional AI/ML, and the forecasted value of capabilities once scaled to other use cases.

Consider the following outcomes, which other organizations have experienced after adopting AI:[3]

3. Google Cloud Gen AI Benchmarking Study, July 2023

Step 6 of 10

Design prompts together with the tiger team

Work collaboratively with the tiger team to design prompts that will guide the generative AI model's responses.

Your three-person tiger team has expertise in the business need, the AI model(s), tuning, and application integration. Use their skills and refer to prompt samples to move quickly in this step.

Step 7 of 10

Build a user experience (UX) and user interface (UI)

Create a user-friendly experience and interface that will run the generative AI model in production for the chosen persona's use case.

Here are some tips to keep in mind:

Keep the interface and design simple.

Begin with a selection screen which allows users to choose the personality to apply to the generated text, images, or output. For example, options could include 'formal', 'casual', 'technical', or 'creative'.

Create a logical and intuitive user flow that guides users through the AI model's functionality. Ensure that the interface design aligns with the expected user journey.

Consider how the new interface fits within the larger ecosystem of existing apps, like email and chat, that might have their own built-in generative AI capabilities.

Ensure the UI/UX is responsive and accessible across different devices and screen sizes, including mobile phones, tablets, and desktops.

Step 8 of 10

Expand usage to additional individuals

Once you are getting acceptable results from tuning, invite two or three other individuals within the chosen persona to start using the model.

Continue testing, measuring, and tuning with this group until you are getting consistent, quality outputs, then expand usage to five to 10 more people within the chosen persona and continue fine-tuning the process.

With each new individual, make sure you understand the different ways each user interacts with the generative AI model. To do so, conduct user interviews, surveys, or workshops to gather insights into user preferences, pain points, and desired functionalities for the AI model's interaction.

Step 9 of 10

Build a Language Model (LM) operations plan

Develop a plan for productionizing and monitoring the AI model's output to ensure it functions effectively and safely.

Key questions to ask your management team:

- Can we quickly evaluate and experiment with generative AI?
- Do we have cost controls during evaluation and experimentation?
- How are we measuring impact? Do we have targeted goals and frequent checkpoints to ensure progress?
- Do we have a mechanism for continuous improvement? Are we able to assess, evaluate, and re-engage to go deeper within existing use cases or expand to more use cases?

An LM operations plan should include, but is not limited to:

Infrastructure setup

Prepare the necessary infrastructure for model deployment, including scalable compute resources and storage. Set up a version control system to manage model versions effectively.

Deployment and monitoring

Deploy the model in a controlled environment, such as a staging environment, to monitor its behavior before going live. Implement monitoring tools to track model performance, safety, and resource utilization during production.

Output and quality

Develop a system to capture the AI model's output and evaluate its quality, so you can measure the effectiveness of the AI-generated responses. Jump to the next section for a list of recommended KPIs that can be used to measure generative AI use cases.

Regular audits and evaluation for expansion

Establish a regular evaluation cadence to assess the quality of the AI-generated output and a plan for further expansion to other areas within the same domain.

Continuous performance improvement and model updates

Performance — which is defined by quality and latency — requires updates to the model to incorporate the latest research advancements and improvements. Conduct A/B testing to evaluate the impact of model updates on safety and effectiveness.

Security and compliance

Ensure the entire system is secure, with appropriate access controls and encryption mechanisms to protect sensitive data. Comply with relevant regulations and responsible AI guidelines.

Human-in-the-loop oversight

Set up a human-in-the-loop process to review and moderate generated content, especially in sensitive or high-risk applications. Develop a feedback loop to continually improve the model's safety and effectiveness based on human moderation.

Incident response and remediation

Develop an incident response plan to handle potential safety breaches or issues promptly and effectively.

Step 10 of 10

Expand usage to additional use cases within the same domain

In the beginning of this chapter, we explained how an organization could start with one generative AI use case and naturally expand into three use cases that all enhance the customer service domain: first, it helped customer service agents answer the question; second, it summarized the frequently asked questions agents received over the phone; and third, it produced answers that could be posted online as written FAQs.

With each use case added to the model, the model itself became more accurate in the domain.

Looking ahead at days 60-90

Once you are ready to extend your use case to external users and/or third-party data, use these methods and tactics to scale quickly and safely:

01

Host a hackathon

Harness enthusiasm across your teams by hosting a hackathon, which encourages employees to brainstorm ideas and get hands-on with AI — all within a matter of days.

02

Bring in partners

In addition to sharing their expertise, partners can consult on business value and technical implementations, provide training, and even work side-by-side with your teams to transfer knowledge as they build out your implementation.

03

Create a center of excellence

Excitement about new technologies can often lead to widespread use. A center of excellence in models, tuning, and application integration can help to standardize processes, share knowledge, and ultimately drive innovation.

KPIs for generative AI

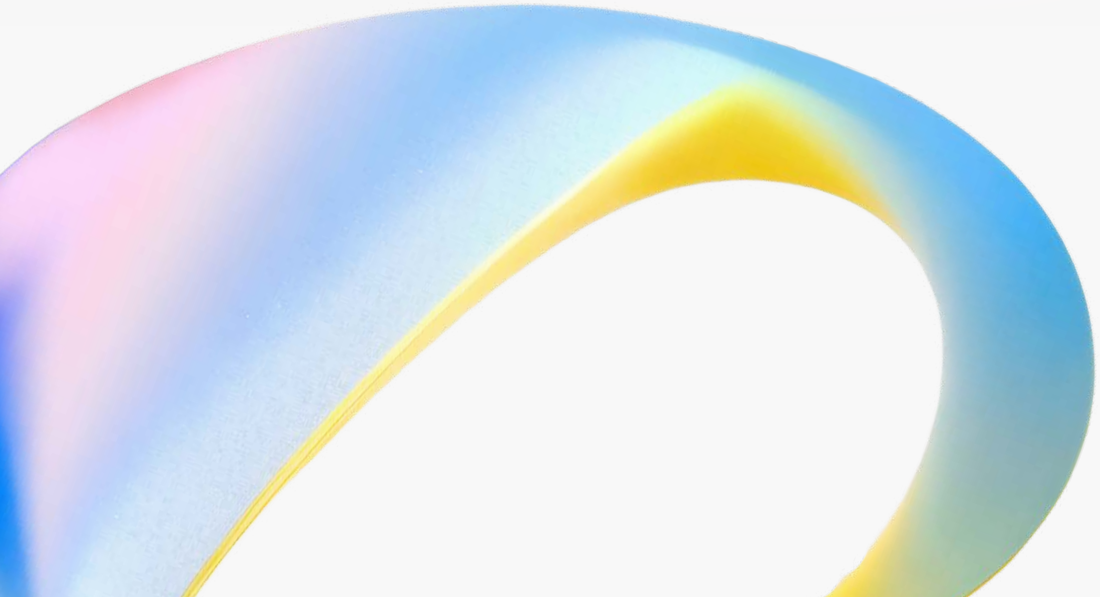

When evaluating projects, consider the feasibility, actionability, affordability, anticipated business value, and ultimate return on investment of each generative AI project.

Like any technology investment, you need to prove its worth. Embed ROI measures into every use case and project, and establish KPIs to keep a pulse on progress along the way.

Consider using these commonly used generative AI KPIs to measure and report on the value of generative AI to your organization, board members, and stakeholders. These KPIs apply to generative AI use cases across various domains and industries.

Accuracy

Measure the accuracy of the generative AI model in producing relevant and correct outputs. This can be quantified using metrics such as precision, recall, F1 score, or mean squared error, depending on the nature of the use case.

Productivity

Assess the impact of generative AI on the productivity of the target persona or department. This could include metrics like the number of tasks completed per unit of time, response time, or reduction in manual effort required.

Customer satisfaction

If the generative AI use case involves customer-facing applications, use customer satisfaction surveys or feedback to gauge how well the AI system meets customer needs and expectations.

Cost savings

Measure the cost savings achieved through the use of generative AI. This may involve comparing the costs of employing the AI system to the expenses associated with traditional manual processes or outsourcing.

Turnaround time

Evaluate the time taken for the generative AI model to generate responses or outputs compared to traditional methods. Faster turnaround times can lead to increased efficiency and improved customer experience.

Quality of output

Assess the quality of the generative AI outputs against predefined criteria. This can be done through manual review or automated quality checks, depending on the use case.

Error rate

Quantify the rate at which the generative AI model produces incorrect or undesirable outputs. Minimizing error rates is crucial for maintaining accuracy and reliability.

Business impact

Identify specific business metrics that are directly impacted by the generative AI use case, such as increased sales, reduced customer complaints, or improved employee retention.

Training time and cost

Measure the time and resources required to train and fine-tune the generative AI model. Efficient training processes can lead to faster implementation and quicker time-to-value.

Human-in-the-loop metrics

If human intervention is involved in the generative AI process, track metrics related to the efficiency and effectiveness of human oversight.

Scalability

Assess how well the generative AI model scales to accommodate increased usage or higher demands. Scalability is essential for long-term success.

Regulatory compliance

For sensitive domains like healthcare or finance, monitor how well the generative AI system adheres to relevant regulatory requirements and data privacy standards.

Chapter 03

The value of generative AI in every industry

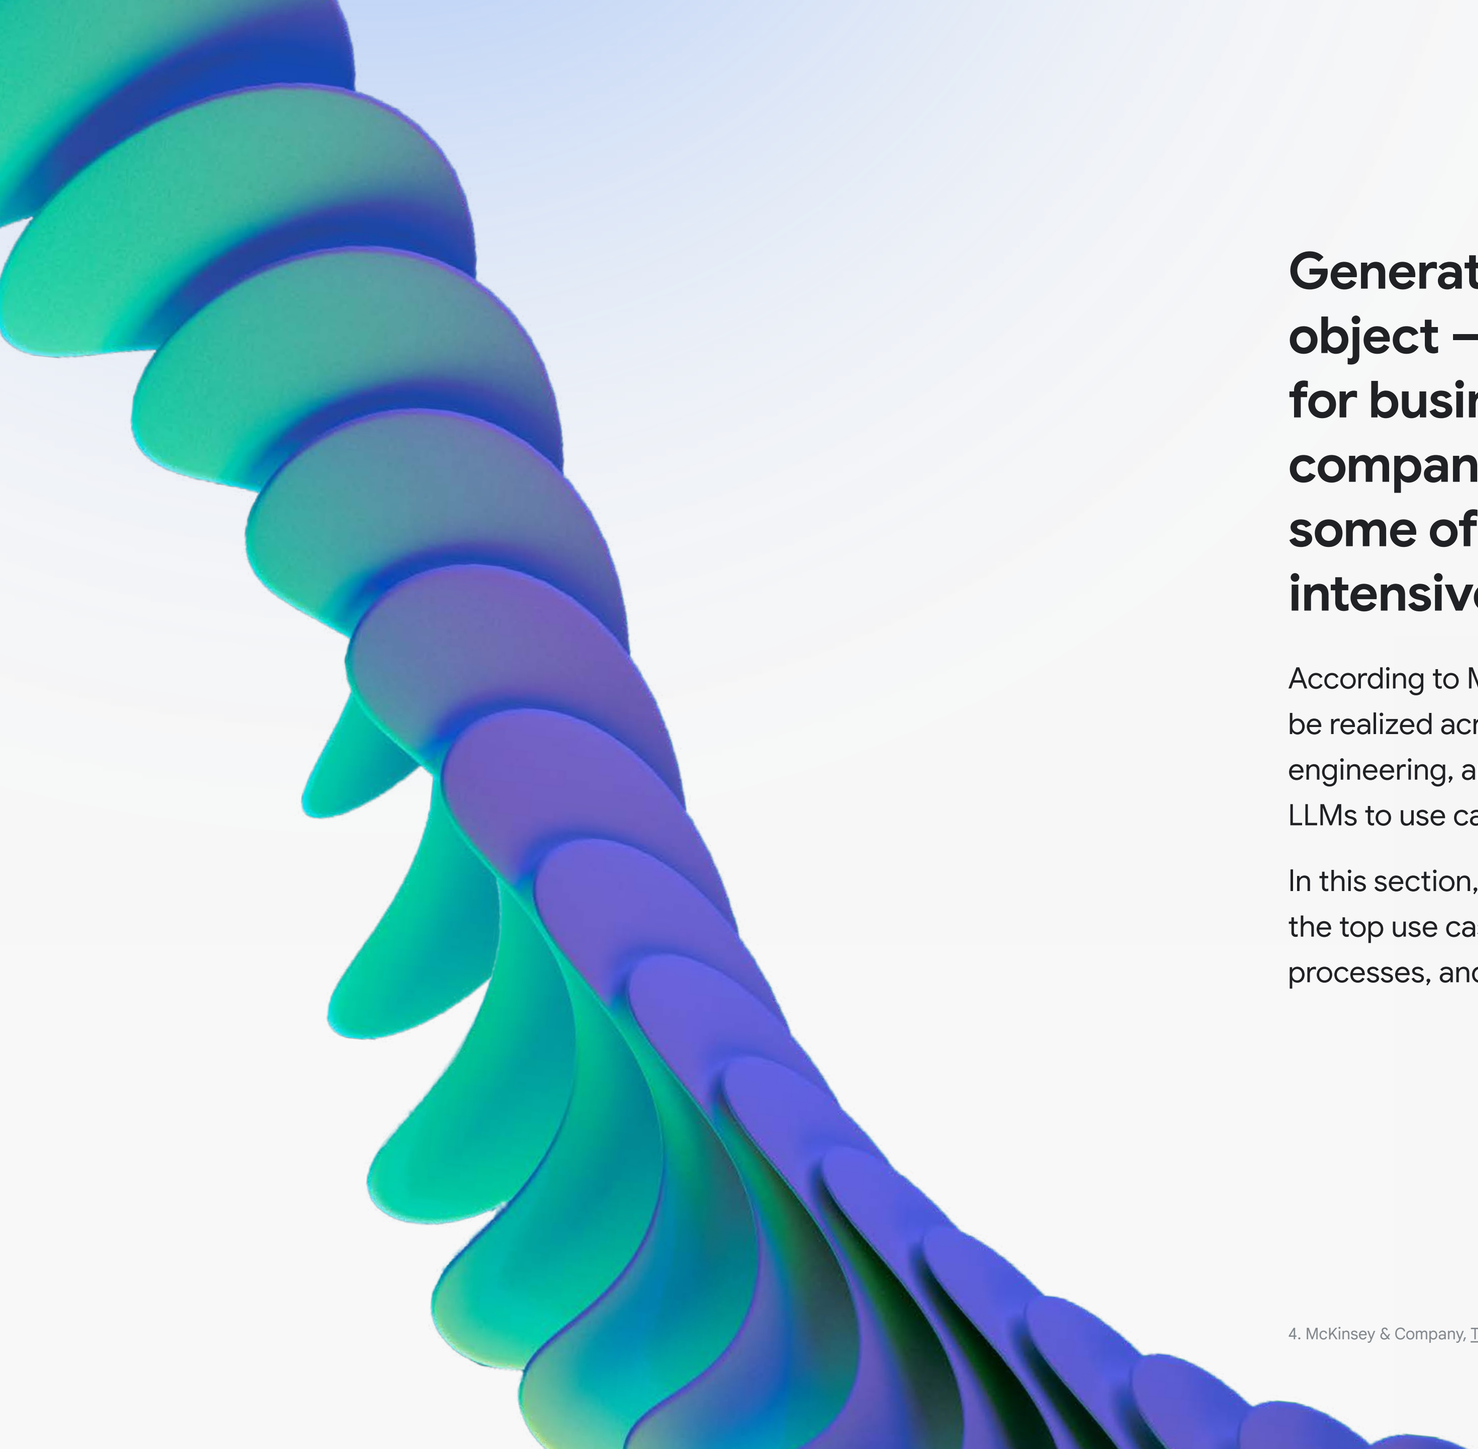

Generative AI is more than a bright shiny object — it is an entirely new value stream for business leaders. Already, leading companies are using generative AI to solve some of their most common and time-intensive problems.

According to McKinsey & Company, 75% of generative AI's value will be realized across customer operations, marketing and sales, software engineering, and R&D.[4] It's already happening, with companies applying LLMs to use cases such as conversational AI in marketing and e-commerce.

In this section, learn how industry leaders are applying generative AI to the top use cases in every industry to unlock new value chains, rewrite processes, and do business in faster time frames and at a lower cost.

4. McKinsey & Company, The economic potential of generative AI, 2023

“Generative AI opens up a new avenue, allowing people to think fundamentally differently about how business works. Whereas AI and ML were more about productivity and efficiency — doing things smarter and faster than before — now, it’s ‘I can do it completely differently than before’.

Carrie Tharp

VP Strategic Industries, Google Cloud

See what your peers are doing

To accelerate your path to value, familiarize yourself with the industry-specific use cases being adopted today. See how the leaders are realizing value in their organizations, and get inspiration for your own.

Explore them all or jump straight to your industry:

- Retail and CPG
- Media and entertainment
- Financial services
- Manufacturing
- Healthcare and life sciences
- Communications service providers

Retail and CPG

Insight

82%

of retail organizations consider customer service automation to be valuable[5]

Priority use cases

Creative assistance

Empower retail creative teams to create bespoke images and creative content for campaigns and editorial placements, and enable 1:1 personalization.

Conversational commerce

Interactively address queries, provide recommendations, and engage with customers in real time to help them make shopping decisions (for example, “Sure, here are some dresses in your size and style you may like, and here are influencer images for style inspiration”).

Customer service automation

Streamline customer service with conversation summaries and task automation.

New product development

Enhance internal consumer research with easy querying, summarization, and insight generation. Create copy concepts and claims for further testing, and visual concepts for product and packaging designs.

5. Google Cloud Gen AI Benchmarking Study, July 2023

Real-world example

Wendy's® revolutionizes the drive-thru experience

Wendy's is automating its drive-through service using an artificial intelligence chatbot powered by natural-language software developed by Google and trained to understand the myriad ways customers order off the menu.

With 75 to 80% of Wendy's customers choosing the drive-thru as their preferred ordering channel, delivering a smooth ordering experience using AI automation can be difficult due to the complexities of menu options, special requests, and ambient noise. For example, because customers can fully customize their orders and food is prepared when ordered at Wendy's, this presents billions of possible order combinations available on the Wendy's menu, leaving room for miscommunication or incorrect orders.

Google Cloud's generative AI capabilities can now bring a new automated ordering experience to the drive-thru that is intended to enhance the experience that customers, employees, and franchisees expect from Wendy's.

Wendy's is beta-testing Google Cloud's AI technology in a Columbus, OH-area, company-operated restaurant, and will use the learnings to inform future expansions to more Wendy's drive-thrus. The test will include new generative AI offerings, such as Vertex AI Search and Conversation and more, to have conversations with customers, the ability to understand made-to-order requests and generate responses to frequently asked questions.

This will all be powered by Google's foundational LLMs that have the data from Wendy's menu, established business rules and logic for conversation guardrails, and integration with restaurant hardware and the Point of Sale system. By leveraging generative AI, Wendy's seeks to take the complexity out of the ordering process so employees can focus on serving up fast, fresh-made, quality food and providing exceptional service.

Read the full story

Wendy's

"Wendy's introduced the first modern pick-up window in the industry more than 50 years ago, and we're thrilled to continue our work with Google Cloud to bring a new wave of innovation to the drive-thru experience."

Todd Penegor

President and CEO, The Wendy's Company

“Google Cloud’s generative AI technology creates a huge opportunity for us to deliver a truly differentiated, faster, and frictionless experience for our customers, and allows our employees to continue focusing on making great food and building relationships with fans that keep them coming back time and again.”

Todd Penegor

President and CEO, The Wendy’s Company

Financial services

Insight

79%

of financial services organizations consider virtual assistants to be valuable[6]

Priority use cases

Financial document search and synthesis

Help analysts find and understand information that is buried deep in contracts and other unstructured documents.

Enhanced virtual assistants

Help customers get the answers they need with less human intervention.

Capital markets research

Use as a research assistant to sift through millions of source documents to identify and summarize key information.

Regulatory and compliance assistant

Help business and technical teams monitor regulatory changes that impact their business, and ensure that controls and compliance are consistently implemented (in software and business processes).

Personalized financial recommendations

Improve cross-sell and retention with 1:1 messaging. Tailor financial product recommendations with hyper-personalized and conversational language.

6. Google Cloud Gen AI Benchmarking Study, July 2023

Real-world example

Deutsche Bank accelerates financial decision-making

Deutsche Bank is testing Google's generative AI and LLMs at scale to provide new insights to financial analysts, driving operational efficiencies and execution velocity. There is an opportunity to significantly reduce the time it takes to perform banking operations and financial analysts' tasks, empowering employees by increasing their productivity.

Watch the full video

Deutsche Bank

"Generative AI is transforming how we find, sort and analyze information at scale — helping us to support our client's global ambitions."

Bernd Leukert

Chief Technology,
Data and Innovation Officer,
Deutsche Bank

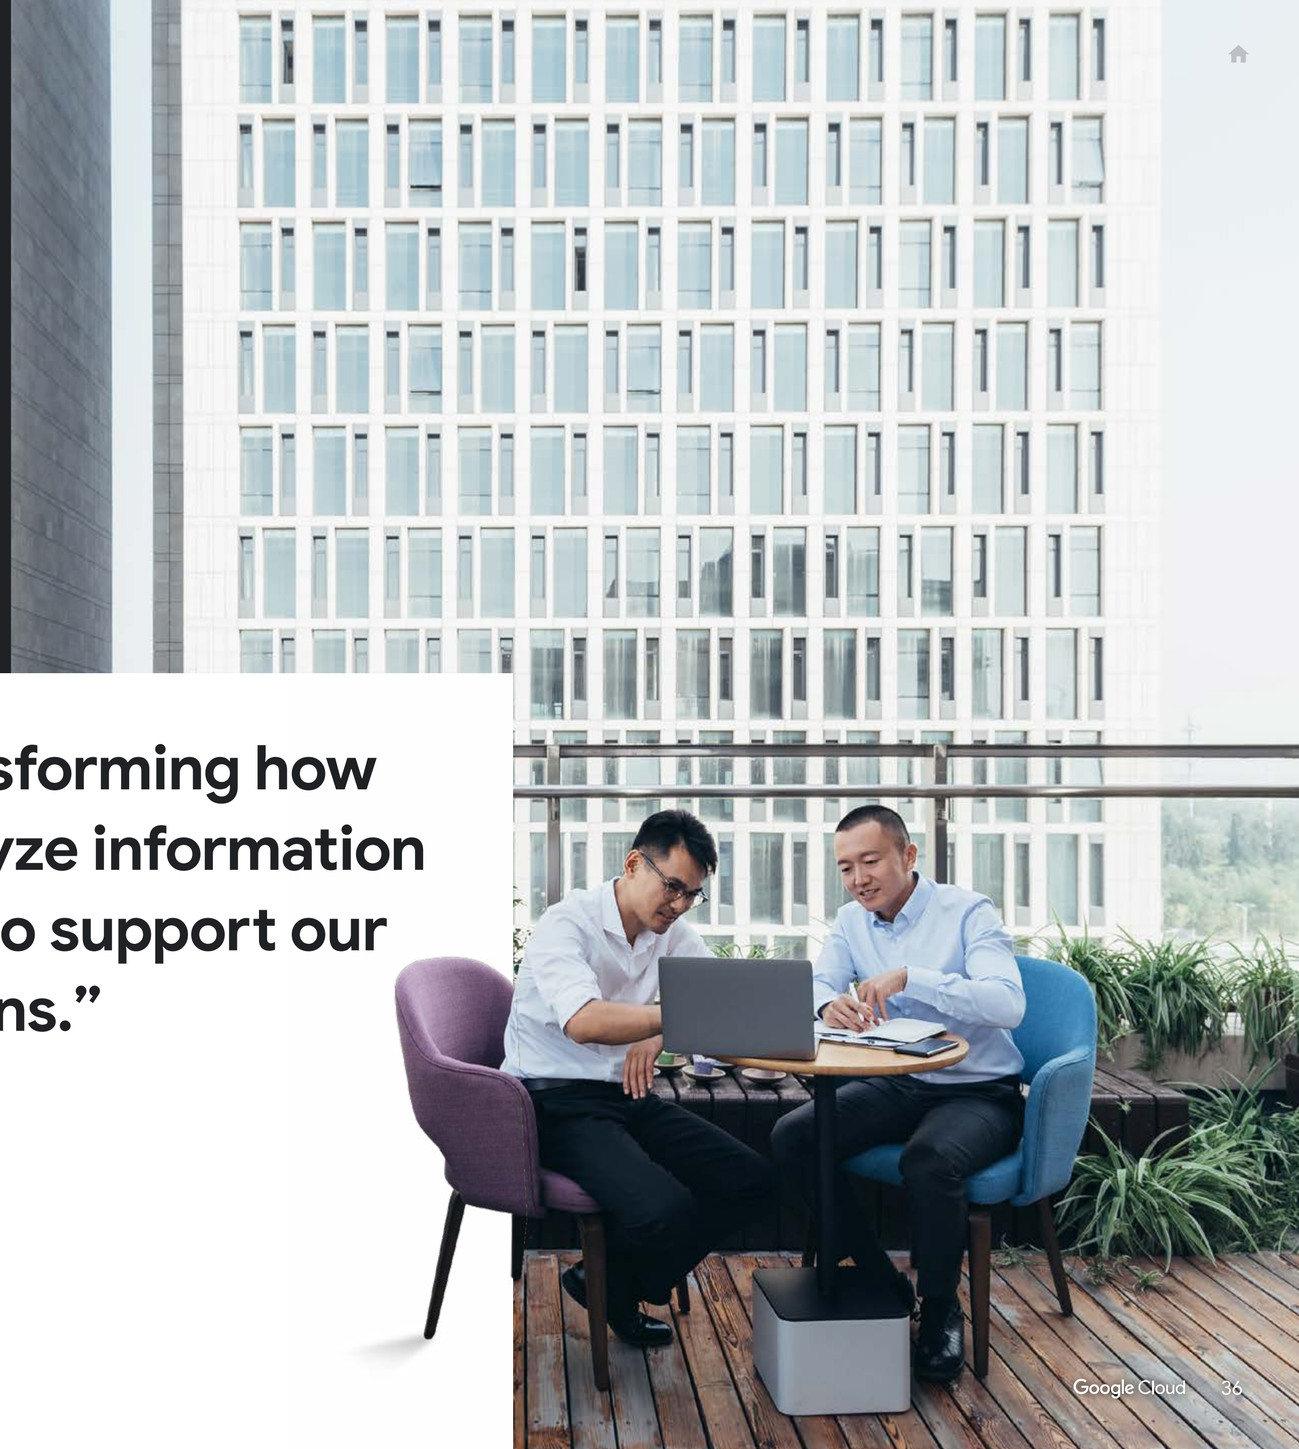

Healthcare and life sciences

Insight

75%

of healthcare organizations consider digital patient concierges to be valuable[7]

Priority use cases

Digital patient concierge

Easily locate, summarize, and generate health plan responses. Clearly explain plans and benefits to members and potential members.

Public and private contextual search

Query and extract insights from public and private datasets, and summarize research into plain language.

Expedite Prior Authorization (PA)

Reduce clinicians' admin time on drafting PA letters for procedures, medication, or medical devices, and accelerate patient care.

Clinical trial report generation

Accelerate generation of clinical studies and reports, complete with safety/efficacy claims.

7. Google Cloud Gen AI Benchmarking Study, July 2023

Real-world example

US national medical center transforms healthcare with generative AI

One of the leading national medical centers and hospitals is transforming healthcare with generative AI. The organization started by using Vertex AI Search and Conversation to improve the efficiency of clinical workflows, helping clinicians and researchers to find the information they need, and ultimately improve patient outcomes.

Healthcare professionals often rely on information from various sources such as medical records, research papers, and clinical guidelines to aid them in diagnosing and treating patients.

However, this data is scattered across multiple formats and locations, making it challenging for clinicians to find the necessary information promptly.

Vertex AI Search and Conversation unifies data across dispersed documents, databases, and intranets, making it easier to search, analyze, and identify the most relevant results.

Learn more about how generative AI can improve healthcare patient service.

Media and entertainment

Insight

87%

of media and entertainment organizations consider media content discovery to be valuable[8]

Priority use cases

Media content discovery

Help users discover new content with personalized, conversational search results based on previous behaviors.

Creative assistance

Make it easier for content creators to repurpose content in different formats, helping accelerate time to value and revenue.

Internal document and media search

Enable internal editing and operational teams to find the right content at the right time.

Branded consumer interactions

Use intellectual property on media assets to create unique and personalized audience experiences.

Content summarization and metadata

Seamlessly extract metadata from media to enable personalization, monetization, and insights, and easily summarize long-form content.

8. Google Cloud Gen AI Benchmarking Study, July 2023

TIME

Real-world example

TIME wants to build community, not just creativity, with LLMs

With trusted sources and chat resources, TIME wants to do more than deliver headline news — playing a bigger role as a beacon for accuracy. As media companies explore the possibilities of generative AI, the publisher sees an opportunity to strengthen its role as a trusted source and community builder.

For years, it's been using AI-powered recommendations to build affinity and loyalty with readers. Now, with generative AI, TIME hopes to turn a one-way conversation into a dialogue.

Read the full story

"As publishers, what we've done for a hundred years, it's been a one-way street: we put out content for the consumer and they consume it. With generative AI prompts and chat, we actually become able to understand, and have an interaction with, the consumer — creating experiences that are two-way in many ways. That's why I actually see generative AI as a powerful tool for building community."

Burhan Hamid

Senior VP Data, Product, and Engineering, TIME

Canva

Real-world example

Canva is solving for AI-powered design, for everyone

Canva is using the latest AI technology to empower their customers and make the design process as frictionless as possible. From enabling users to translate their designs into over 100 languages with just a few clicks, to turning short videos into longer and more compelling clips with Google PaLM technology, they are unlocking the magic of AI with Google Cloud.

Watch the full video

“There is a quote that I love from the science fiction writer Arthur C. Clarke, ‘Any sufficiently advanced technology is indistinguishable from magic.’ Canva has always been about removing as much friction from the design process as possible and AI technology allows us to make the design process even easier with even less friction.

We are incredibly excited to be working with Google Cloud as we test and explore ways to bring even more magic to our community. Earlier this year we launched Magic Translate and what that enables you to do with just a couple of clicks turn any of your designs into over a hundred languages. Magic video is a time-saver for marketers and teachers who want to be able to cater to a diverse audience.”

Melanie Perkins
Co-founder & CEO, Canva

Manufacturing

Insight

80%

of manufacturing organizations consider machine generated events monitoring to be valuable[9]

Priority use cases

Machine-generated events monitoring

Interpret telemetry from equipment to reduce unplanned downtime, optimize operations, and maximize utilization.

Customer service automation

Provide an easy, informative, and value-added customer service experience that automates and accelerates time-to-resolution for common interactions.

Document search and synthesis

Retain generations of documents throughout the product lifecycle, and use them to generate new content as needed.

Product/content catalog discovery

Efficiently match requirements to the specifications of products purchased.

Supply chain advisor

Optimize fulfillment by recommending best-suited suppliers based on relevant criteria.

9. Google Cloud Gen AI Benchmarking Study, July 2023

Real-world example

Leading global airline supplier, GA Telesis, integrates generative AI technology

As a major supplier of essential equipment in the airline industry, where long-term relationships and trust are the bedrock of many business transactions, GA Telesis' sales staff receive inquiries from global customers requesting quotes for all sorts of commercial aircraft and jet engine replacement parts.

The typical inquiry is not standardized, requiring sales representatives to quickly decipher the relevant aircraft or jet engine model, applicable codes, quantity required, preferred condition and provenance, and often most importantly, where the part is needed and when. Additionally, in order for airlines to meet their on-time performance metrics, inquiries are often urgent and logistics have to be factored into the equation. GA Telesis' team is expected to accomplish what can resemble an impossible feat in minutes, not hours.

GA Telesis has selected Google Cloud's Vertex AI Search and Conversation platform, which is designed to help businesses tune and deploy machine learning models, to help it quickly build innovative AI applications. Leveraging a new data extraction solution the GA Telesis technology team built internally, GA Telesis will be able to automatically synthesize purchase orders and quickly provide customers a quote, eliminating the need for sales teams to manually cross-reference emails with their inventory availability.

Read the full story

"In aerospace, GA Telesis will deploy Google Cloud's generative AI technology to revolutionize the sales and service processes for the parts the company supplies to major global passenger and cargo carriers."

Abdol Moabery
CEO, GA Telesis

Communications service providers

Priority use cases

Customer or employee service automation

Make online customer service more conversational with human-like support and search.

Network planning and operations

Easily access and understand complex data on network performance, faults, inventory, infrastructure, and anomaly detection.

Advertising and creative content assistance

Generate interactive and relevant content with highly personalized messaging.

Employee knowledge search

Make it easier and more effective for staff to get things done, with human-like bots providing IT support, self-service, and T2/T3 guidance to field techs.

Test or code script generation

Generate and test experiments using real work experiences.

Contract analysis and negotiation

Automate contract negotiations with suppliers by analyzing bills, trends, and other supply data.

Innovate faster with generative AI for business

When a new technology moves as fast as generative AI does, it can be hard to keep up.

As a strategic partner to our customers, Google Cloud helps leaders chart their path forward with the appropriate frameworks, tools, and governance structures — and ingrain a responsible, consciously cautious approach to AI across your organization.

Google is an AI-first company. Having already built some of the industry's leading AI capabilities, we continue to focus on making it easy and scalable for everyone to innovate with AI.

We support the needs of generative AI in your organization in a number of ways.

We have the most comprehensive platform now available and ready-to-go with strong support from leading organizations — helping you create amazing content, synthesize and organize information, automate processes, and build engaging customer experiences.

Your data is your data. We do not use our customers' data to train Google's models. The question we hear most is, "Do I have control of my data, brand, IP risk, and ability to meet regulatory requirements? The answer is "Yes".

Everyone can be an AI developer. All users with varying levels of expertise can create innovative enterprise search, chat, and vision apps. We enable both business and technology practitioners to be more productive using AI assistants.

We deliver infrastructure that is optimized for AI workloads by giving you access to the latest GPUs and TPUs, a rich choice of deep learning VMs, and the ability to easily build custom AI software.

The Google Cloud AI portfolio supports all stages of your generative AI journey. With a rapidly growing suite of generative AI technologies being made available — along with new educational and consulting programs, blueprints for specific industry use cases, and our growing partner ecosystem — we are ready to get you and your teams learning, building with, and deploying generative AI.

We're building a capability that allows you to get on the bus for generative AI and go where it's going. Start building out your own enterprise skill sets and capabilities, such that when you find the right use cases, and the right value levers — you have the capability to do that."

Carrie Tharp

VP Strategic Industries, Google Cloud

Hit the ground running with generative AI.

Contact us to get started today.

cloud.google.com/ai/generative-ai

Google Cloud