



6 key steps to take on your generative AI journey

Leveraging the full power of artificial intelligence
for leading software companies



Table of contents

INTRODUCTION

Forging ahead	3
What is artificial intelligence, machine learning, and generative AI?	4
Why generative AI?.....	5

THE GENERATIVE AI JOURNEY

Step 1: Championing an innovative culture	7
Step 2: Make data your differentiator.....	9
Step 3: Finding the right business problem to address	11
Step 4: Upskilling teams.....	13
Step 5: Scaling beyond pilot projects	16
Step 6: Measuring the results	18

CONCLUSION

Taking the next step with AWS	20
Solving the biggest artificial intelligence challenges	22

INTRODUCTION

Forging ahead

When deployed with the right strategy, generative artificial intelligence (AI) has the potential to change the anatomy of work. By automating work activities that absorb 60–70 percent of employees' time, organizations can improve productivity down to the individual worker. This can increase agility, streamline processes, boost revenue, and allow for better, faster decision-making for leading software companies.

The lack of a single universally accepted playbook for generative AI success is keeping some businesses on the sidelines, unsure of how to take the next (or even the first) step on their AI journey. By democratizing generative AI, Amazon Web Services (AWS) is helping organizations overcome barriers to adoption and forge ahead with confidence. This eBook outlines a proven path—from taking the first step to measuring results—with insights from Amazon best practices and its experience helping thousands of customers realize their own initiatives.

In a global survey of software and digital native companies, all respondents said there is a need to increase investment in AI. Especially to stay ahead of the competition.

In fact,

83%

agree that in the next five years, the only competitively viable publishers of business and consumer software will be those executing an effective AI product strategy.¹



WHY GENERATIVE AI?

What is artificial intelligence, machine learning, and generative AI?

AI is a way to describe any system that can replicate tasks that previously required human intelligence. Almost all AI systems are built using machine learning (ML). ML utilizes large amounts of data to create and validate decision logic, forming the basis of an AI model. The AI application then feeds input data into that model, and the model outputs human-like decisions. The rapid advancement of ML, massive proliferation of data, and easy availability of scalable compute capacity can help businesses accelerate how they use AI and, particularly, how they leverage generative AI.

Generative AI is a type of AI that can create new content and ideas, including conversations, stories, images, videos, and music. Like most AI, generative AI is powered by ML models—very large models that are pretrained on vast amounts of data and commonly referred to as foundation models (FMs).

Why generative AI?

Before diving into the steps of your generative AI journey, let's explore why software companies should go on that journey in the first place. According to Goldman Sachs, generative AI could drive a 7 percent (or almost \$7 trillion) increase in global GDP and lift productivity growth by 1.5 percentage points over a 10-year period.² Globally, we have reached an inflection point where most customer experiences and applications can be reinvented with generative AI.

Developing a strategy that includes generative AI initiatives is imperative to ensure successful business outcomes. Even after completing the steps outlined in this eBook, a business will need to regularly remind itself what it's working toward—staying focused on the precise business benefits that can be unlocked by fully leveraging generative AI technology.



Businesses are already realizing the impact of:

Boosting employees' productivity

Generative AI is being leveraged for its transformative value to help software companies reach new levels of productivity for their customers. The latest advancements in generative AI can be used to continuously **increase employee productivity** with the help of generative–AI–powered conversational search, content creation, text summarization, and code generation. Real-time monitoring company, Dataminr, used AWS Inferentia to grow their AI models. They increased processing speed, reduced costs and improved accuracy. All of which led to bigger, better innovation on new projects. And allowed them to provide a faster and more accurate service.

Improving customer experiences

Today's organizations can take advantage of software solutions that incorporate generative AI to enhance customer engagement, increase personalization, and attract new users through deeper experiences. AWS customers have improved their own **customer experiences** through the effective use of chatbots, virtual assistants, intelligent contact centers, personalization, and content moderation. One of those customers—like

restaurant management platform **Upserve**—worked closely with the Amazon ML Solutions Lab to build a predictive model ML solution. This is designed to predict how many people will visit per night and which menu items will be popular. All of this is made possible by integrating table management, point of sale and other data systems in real time.³

Transforming content creation for greater creativity

Companies like game-industry-enabler **Scenario** are using generative AI to turbocharge real-world production processes across multiple types of creative content—including art, music, text, images, animations, and video. The AI-powered capabilities for automating writing, media design, and character modeling allow for unprecedented creative exploration and speed.

Improving outcomes for real, competitive advantage

For software and technology leaders to stay relevant and competitive, generative AI should be embedded into every business plan.

Customer engagement platform Braze has made generative AI part of its daily activities to successfully send billions of customer messages on behalf of household brands. Using AWS, Braze can access real-time data so its customers can deliver best-in-class end-user experiences.⁴


Now that we've outlined the *why* of generative AI, it's time to explore the *how*

The next sections will demonstrate the steps in the AI journey using AWS best practices and those of AWS customers to exemplify the necessary changes that take place in order to successfully implement, deploy, and scale AI solutions.



³ Upserve Case Study, AWS Machine Learning Blog, 2016

⁴ Braze Delivers Personalized Customer Engagements Using AWS, Braze, 2022



The generative AI journey

The generative AI journey is not necessarily a straightforward path. Achieving success with generative AI requires more than just great technology—it also means ensuring that the organization is aligned to the right goals. Identifying and reaching those goals will necessitate broad changes in processes, management, and culture. The next sections will explore how organizations can overcome common challenges that often impede progress and take the right steps to implement AI in efficient, sustainable ways.

STEP 1

Championing an innovative culture

Unlocking the full business potential of generative AI requires cultural changes in team organization, objectives, and outlook.

For generative AI to proliferate through an organization, both business and technical teams have to work together and share the same priorities. To achieve this at the outset, the generative AI effort must be supported from the highest levels, with goals set by executive champions and an investment in the technology and processes that enable success. This includes, among others, a commitment to build more responsible AI from the beginning—working to identify and mitigate bias, improve explainability, and help keep data private and secure. By taking a people-centric approach, organizations can work to educate their workforce on responsible AI and build more diverse teams to bring more perspectives to the table to improve fairness.

It's important for management to take a wide-scale view while fostering AI initiatives. Executives must be firm in their goals but also flexible in how the organization reaches them. Mistakes are sure to be made. But by staying focused on the long-term outlook and not allowing discouragement, organizations can glean wisdom from every error and apply those learnings to champion an innovative culture throughout the business.

Perhaps the largest cultural change that organizations must undergo is utilizing the opportunity inherent in mistakes. Integrating AI is an iterative process that can only succeed through constant experimentation. Often, these experiments will result in failure. Only by learning from mistakes—and refusing to grind progress to a halt in the name of determining “what went wrong”—can organizations consistently reach the breakthrough successes waiting on the other side.

How Amazon did it

Amazon has been using AI for over 20 years. After a decade of leveraging the technology, our leadership team asked every business leader in the organization—irrespective of whether they ran a research team, a fulfillment center, or an HR organization—to answer the question of how they planned to use AI in their businesses.

“We don’t plan to” wasn’t an acceptable answer in most cases, which forced the leadership, domain experts, and technical experts to collaborate on AI initiatives and let nothing halt their progress—even in instances where tangible benefits were still years down the road.

In addition to hiring external data scientists, Amazon created the [Machine Learning University \(MLU\)](#), which trained many of its developers to use AI more effectively. The company also built

Amazon is using AI to minimize packaging waste, reducing outbound packaging weight by

33% and eliminating
915K TONS

of packaging material worldwide.⁵

tools like [Amazon SageMaker](#), which simplifies model creation and removes the barriers to entry, enabling AI technologies and initiatives to scale more effectively. Additionally, Amazon created a set of pre-built [AI services](#) that provide ready-made intelligence to address common business use cases—without customers having to build their own models. For example, [Amazon Bedrock](#) is a new service that makes generative AI FMs from AI21 Labs, Anthropic, Cohere, Meta, Stability AI, and Amazon accessible via an API. Bedrock is the easiest way for customers to build and scale generative-AI-based applications using FMs. Bedrock offers the ability to access a range of powerful FMs for text and images—including [Amazon Titan FMs](#)—through a scalable, reliable, and secure AWS-managed service.

This is why major organizations—including Intuit, Thomson Reuters, AstraZeneca, Ferrari, Bundesliga, 3M, and BMW—as well as thousands of startups and government agencies around the world, are transforming themselves, their industries, and their missions with generative AI solutions from AWS. We take a democratizing approach to generative AI, and we work to take these technologies out of the realm of research and experiments and extend their availability far beyond a handful of startups and large, well-funded tech companies.

Let’s take a look at some examples of how Amazon is leveraging AI.

First, Amazon uses AI throughout its fulfillment process and leverages a forecast system that can predict demand for nearly every product in its enormous inventory. These prediction models allow Amazon to better deliver on customer expectations of convenience, cost, and delivery speed.

“We forecast millions of products every single day across all of our Amazon sites worldwide,” said Jenny Freshwater, Director of Forecasting at Amazon. “And without machine learning, we would not be able to produce those forecasts.”

The examples go on and on. Amazon also created Alexa, which provides customers with an entirely new way to interact with technology. Additionally, the company developed groundbreaking technology with autonomous flight via Amazon Prime Air drones. Amazon also uses robotics in its fulfillment centers to get packages to customers faster.

Achieving these successes required great investments in technology, research, and talent. But those investments would have gone to waste without the cultural changes that pushed them forward through many failures and unexpected challenges. Every organization must follow suit and foster this same fault-tolerant culture of experimentation and innovation before the AI journey can truly begin.

STEP 2

Make data your differentiator

Success with AI is highly dependent on quality data. Without a proper data strategy in place, progress will slow to a crawl and hamper the effectiveness of the final model. Worse yet, if the model is informed by bad data, the results it generates may be misleading—or even flat-out wrong.


“Machine learning models are highly sensitive to data quality,” Freshwater said. “So we learned—in many cases the hard way—that the time spent on getting data of high quality on the way in paid dividends in our models on the way out.”

The right data strategy for AI should aim to break down silos, enabling IT teams to access and collect the data they need easily, quickly, and securely.

While modern data strategies take many forms, data lakes are becoming an increasingly popular core component of the most efficient models. Data lakes offer more agility and flexibility than traditional data management systems, allowing companies to manage multiple data types from a wide variety of sources and store that data—structured or unstructured—in a centralized repository.

Once stored, data can be analyzed by many types of analytics and AI services faster and more efficiently than with traditional siloed approaches. Data lake architectures also enable multiple groups within the organization to benefit from analyzing a consistent pool of data that spans the entire business.

For help developing a more holistic data strategy that includes data lakes, visit [AWS for Data](#).



How Figure Eight did it

For over a decade, Figure Eight has helped product, data science and ML teams more easily sort through their data for ML projects. Using its enterprise-ready platform, users can quickly prepare and categorize their own data—while Figure Eight annotates, judges and labels that data to create a base for project models.

With AWS, Figure Eight has generated more than 10 billion data labels for real-world AI applications so far. All of which allow teams to focus on improving ML models' accuracy—instead of managing, cleaning and disputing data. Its platform provides accurate training data for likely scenarios by testing labelers before and during data tasks. This activity helps financial institutions keep their transaction models up to date, prevent data drift and better-trained OCR models on financial documents.

“Our team has built a platform that trains, tests, and tunes machine learning models for companies in a variety of industries. For example, one of our previous projects was working with an online marketplace to automatically shorten product titles—which can often be 40-50 words long—so they could then be read aloud for a conversational shopping experience. SageMaker has enabled us to both spend more time focusing on important processes and increase the rate at which we prototype solutions.”

Joan Xiao, Lead Machine Learning Scientist, Figure Eight

STEP 3

Finding the right business problem to address

One mistake software companies often make in their AI journeys is employing discrete data scientists who work in silos to build models as proofs of concepts rather than solve real business problems. This way of working has shown that many software and digital-native companies are missing genuine business value when doing this. And that AI has a lot to offer.

Here are some important questions that organizations should ask before embarking on an AI journey:

1. Is the project important enough to get attention and adoption?
2. Does it solve a real business problem?
3. Do we have the right data to solve the problem?
4. Will the project benefit from AI?
5. Can it eventually be operationalized?

The **AWS AI Use Case Explorer** is a business outcome-centric search and navigation site that enables users to find the right AI use cases, discover relevant customer success stories, and mobilize their teams toward deployment.

In a successful AI journey, organizations create AI teams to address specific business problems. This requires including both technical and domain experts within these teams. While the technical experts will take on the brunt of model creation, they need the field knowledge of domain experts to define precise business challenges and identify the data most important to finding a solution.

This approach is also critical to change management. When technical and domain experts collaborate to create models, employees will feel more confident in making decisions based on the algorithm's logic.

Together, these teams should also work through how to measure success. Freshwater said the companies should make sure that they have “very crisp and clear metrics as you embark on the machine learning journey. Many times, your models are taking over for something existing, and you want to make sure that they’re actually better and that you can measure it.”

For more on measuring the success of AI initiatives, refer to Step 6 in this eBook.

Some organizations have the talent in-house to identify the problems that would be best addressed by AI and to implement the appropriate pilot programs. Organizations can leverage the **AWS Generative AI Innovation Center** to work backwards from business challenges and go step by step through the process of creating AI projects and initiatives.

How Tangent Works did it

Tangent Works specializes in automating machine learning modeling processes, and leverages AWS to provide accessible, cost-effective AI technology for businesses without dedicated data science teams.

By using AWS, Tangent Works efficiently launches services, manages computer and storage needs, and supports rapid growth. Its Tangent Information Modeler (TIM) empowers business users and data scientists to easily create accurate AI models for time-series data. AWS then enables Tangent Works to scale resources dynamically to meet different customers' demands. While tools like Amazon EKS, AWS Fargate, and Amazon RDS for PostgreSQL can then concentrate on facilitating workload management.

TIM offers bespoke AI capabilities, delivering rapid, cost-effective model creation and execution. This technology is widely applied across multiple industries for tasks like demand forecasting, maintenance scheduling, and fraud detection. The simplicity of updating and running models allows customers to keep insights current and adapt quickly to changing conditions.

Working with Siemens Digital Industries Software, Tangent Works integrated its technology into Siemens' industrial IoT platform, enabling users to analyze IoT data and develop data models for improved decision-making. Tangent Works also aims to enhance automated modeling functions within existing AWS machine learning tools and refine its time series anomaly detection technology.

By utilizing AWS, Tangent Works facilitates widespread access to advanced analytics and machine learning capabilities, so organizations can harness insights without the need for extensive in-house data expertise.



STEP 4

Upskilling teams

In parallel with creating a comprehensive data strategy, today's software companies must focus on arming their teams with the right skills to succeed in the era of generative AI. However, businesses across dozens of industries are growing increasingly aware of an expanding skills gap—the separation between technologies and the ability of internal IT and business specialists to take full advantage of them. It's an issue that should set off alarm bells in light of recent research from the World Economic Forum, which found that more than 75 percent of organizations plan to adopt AI in the next five years.⁶

To help individuals train for the AI jobs of the future, AWS released on-demand skills training to support those who want to understand, implement, and begin using generative AI. Amazon has also designed training courses specifically for developers who want to use **Amazon CodeWhisperer**.

While there's no one-size-fits-all solution to the AI skills gap, there are proven methods that can maximize the abilities of existing staff, reducing the need to make large investments in recruiting pretrained expert talent.



⁶ "The Future of Jobs Report 2023," World Economic Forum, April 2023



These methods include:

Defining the skills gap: Before closing the skills gap, an organization must identify the precise differences between what it needs or wants its employees to do and what its employees currently can do.

Understanding how skills are mapped: Because AI initiatives are interdisciplinary efforts, an organization should map the skills needed across data scientists, data engineers, business analysts, application developers, statisticians, and other subject matter experts in the business.

Customizing training for specific needs: If an organization has existing training curriculums that could be useful, it should work to tailor those materials to the business' specific AI needs. Leaders should also investigate pretrained AI services that provide ready-made intelligence for business applications and workflows.

In addition to training, organizations need to align teams to successfully tackle AI problems. They can achieve this by:

Promoting a culture of empowered teams: AI project teams must be cross-functional, with the authority to execute individual objectives and the freedom to organically cross-pollinate with other teams as demands dictate and opportunities arise. To make this kind of teamwork possible, management will need to embrace new structures—letting go of the strictly hierarchical and departmentally siloed organizational models of the past.

Starting with a pilot team: The business should establish a pilot team of engineers, IT and AI practitioners, and line-of-business leaders—and task it with an AI project. “I’d recommend putting a couple of really smart people on trying to figure out what metrics you want to optimize for or predict” Freshwater said.

Enabling organic transformation: Once the pilot project is complete, the business can split up the team, add people to create new teams, and task them with new projects. This process continues, allowing knowledge to spread organically from veteran team members to new recruits and pollinate between teams.

By following this guidance, many organizations are finding that the people they currently have are actually the people they need to close their AI skills gap.

While some recruiting may still be required, investment in the right tools, processes, and management changes can do much of the work to upskill talent for AI success.

How Persistent did it

Persistent Systems is a multinational tech services company which specializes in software development and tech consulting. It has been working with transformer models and AI for the past few years, building capabilities, accelerators, and client engagements.

As part of its goal to increase developer productivity, Persistent has given its 16,000 engineers access to Amazon CodeWhisperer. This allows them to build applications faster and more securely.

Having been trained on billions of lines of code, CodeWhisperer now allows engineers to generate code with comments or existing code—saving time on development and freeing them up to create new solutions.

Persistent Systems is dedicated to bringing Generative AI at scale in a safe, trusted, and cost-effective manner. And focuses on privacy, security, and regulatory requirements to address client concerns and accelerate the adoption of generative AI.

“We have always believed in staying ahead of the disruptive technology curve, such as generative AI. Our proactive approach and deep collaboration with AWS have enabled us to provide game-changing engineering solutions that have shortened time to value for our clients across industries.”

“With the adoption of CodeWhisperer by our engineering teams, we are witnessing valuable productivity gains across the software engineering lifecycle and expect these to improve significantly over time.”

Sandeep Kalra, CEO and Executive Director, Persistent

[Read the full story ›](#)



STEP 5

Scaling beyond pilot projects

After the first few successful pilots, software companies must take the next step on the journey: sustainably scaling AI across the business. This is both a technical and a cultural challenge. There are several ways companies can approach the cultural shift necessary to scale AI. Some might find success by creating a center of excellence that rallies the community and continues to push for new initiatives. Or, like Amazon, organizations can make AI an integral part of yearly planning processes, continuously bringing domain and technical experts together to brainstorm and determine their next steps.

Achieving scalability requires organizations to help their developers use AI. Building models at scale can be labor-intensive, which can slow innovation. With the **Amazon Bedrock** serverless experience, you can get started quickly to privately customize FMs with your own data and easily integrate and deploy them into your applications using the AWS tools and capabilities you already know (including integrations with features like pipelines to manage your FMs at scale). Many organizations are solving scalability with SageMaker to prepare data. As well as build, train, and deploy models, and get them into production faster at lower cost. This enables sustainable expansion of AI initiatives beyond pilot projects.

Many more organizations are scaling through **AWS AI services**, a set of pretrained and managed services that can be used as building blocks to address common use cases, including personalizing recommendations, modernizing their contact centers, improving safety and security, and increasing customer engagement.



How Showpad did it

Showpad is a global leader in sales and revenue-enablement technology. It brings together industry-leading training and coaching software with innovative content solutions, to empower sales and marketing teams and engage buyers.

To meet new expectations, Showpad needed to evolve their seller roles. This would help prospects and customers make sense of new information and put it into context for their businesses and the specific challenges they face.

With a bit of AI innovation throughout the product experience, Showpad introduced Amazon Bedrock and Amazon SageMaker to test the waters and subsequently push new models to production.

What came out of this process was PitchAI, Showpad's latest sales coaching software. It now helps sellers improve their sales pitches with actionable feedback in real time. The AI-powered search enriches search results and reduces sellers' time spent searching for the information. While its AI-powered test questions streamline the process of creating training materials for sellers.

“Showpad is leveraging AI to improve how sales and marketing teams partner together in pursuit of a shared goal—closing more deals. It enables sellers to focus more time on what truly matters—bringing meaningful value to each buyer interaction.”

Tony Grout, Chief Product Officer, Showpad



STEP 6

Measuring the results

When measuring the results of AI efforts, the traditional “project ROI” viewpoint—where a project has defined start and end points, a budget, and a return—is reductive and can be detrimental to the initiative’s success. If the project doesn’t generate a positive return within the given time frame, the business may lose interest and miss out on critical opportunities down the line. Instead, executives and IT alike must measure AI efforts based on what success means for their businesses with regard to the processes being optimized. In addition, they must view AI efforts as long-term investments, acknowledging that a true “return” may not be realized for several years and after countless iterations.

When planning AI initiatives, it’s better to view the process through the lenses of agility, competitive advantage, or risk tolerance rather than “expected” return. An organization will have greater success if it disregards the question of “What will be my return on investment in X months?” in favor of something more like, “If we don’t invest in this now, will we fall behind our competitors in X years when the technology matures?”

While traditional ROI metrics may not be the best approach, the business impact of AI initiatives can still be measured—it just requires a different outlook. AI results can be measured through something resembling a “value tree,” where the main trunk of the tree represents the traditional revenue return and branches extending from the trunk recognize the value of other business outcomes. The specific branches of the value tree will depend on the organization, the industry, and the initiative. But they might be things like time saved through automated processes, new leads, markets, and opportunities identified, customer service improvements, or increases in upsells.

Measuring the success of AI through a more holistic and long-term model will keep organizations focused on the best outcomes for their business future.

How AI21 Labs did it

AI21 Labs (AI21) is a leader in generative AI and large language models (LLMs). Initially, the company released two models: one with 7 billion parameters and another with 178 billion parameters. But its team saw an opportunity to offer customers an LLM of 17 billion parameters to bridge the gap between the existing sizes.

AWS trained the FM in under 20 days using **SageMaker**—saving several weeks of time compared with AI21’s previous training methods. “Because Amazon SageMaker handles node failures, restarts elegantly, and orchestrates large, distributed runs, the team working on pre-training the model can focus on core tasks,” says Dan Padnos, Vice President of platform at AI21.

[Read the case study ›](#)

“We have a really good relationship with the AWS team—(its members) have gone deep into the technical details with us and collaborated on challenging tasks. Throughout the process, the AWS team has been creative and has had awareness about our challenges and goals.”

Dan Padnos, VP of Platform, AI21 Labs



⁷ [“In the News: Morningstar Launches Global AWS DeepRacer Corporate Competition to Accelerate Application of Machine Learning,” AWS for Industries Blog, January 2021](#)

CONCLUSION

Taking the next step with AWS

No matter where your software company is in its AI adoption journey, you can take the next step with AWS solutions built on the most comprehensive cloud platform and optimized for generative AI with high-performance computing (HPC), security, and analytics. Featuring the world's broadest and deepest set of AI services, over 100,000 customers are running their AI workloads on AWS. Generative AI can help you realize new business value within your organization. From reinventing customer experiences to enhancing productivity and accelerating growth, generative AI holds the power to help you not only transform your own business, but your clients' too.

Generative AI with AWS, by the numbers

100,000+

customers using AWS for their AI workloads

20+ years

of building experience at Amazon

Hundreds

of algorithms and models in Amazon SageMaker JumpStart

Empower better customer experiences



Chatbots and virtual assistants: Streamline customer self-service processes and reduce operational costs—leverage generative-AI-powered chatbots, voice bots, and virtual assistants to automate responses to client and customer queries.



Call analytics: Analyze calls from contact centers to extract valuable insights to improve experience and loyalty.



Agent assist: AI and generative AI can support and enhance the capabilities of human agents in tasks such as customer service, problem-solving, and decision-making.

Support better employee productivity



Conversational search: Quickly and easily find accurate information and summarize content through a conversational interface.



Code generation: Accelerate application development with code suggestions based on the developer's comments and code.



Automated report generation: Generate financial reports, summaries, and projections, saving time and reducing errors.

Enhance creativity and content creation



Marketing: Create starting points and ideas for engaging marketing content, such as blog posts, social media updates, and email newsletters, to save time and resources.



Sales: Personalize emails and messages, based on prospect's profile and behavior, to help improve response rates. Generate sales scripts or talking points based on clients' or customers' segment, industry, and product or service.



Product development: Quickly generate multiple design prototypes based on certain inputs and constraints, speeding up the ideation phase. Or optimize existing designs based on user feedback and specified constraints.

Accelerate process optimization



Document processing: Improve business operations by automatically extracting and summarizing data from documents and insights through generative-AI-powered question and answering.



Data augmentation: Generate synthetic data to train ML models when the original dataset is small, imbalanced, or sensitive.



Supply chain optimization: Improve logistics and reduce costs by evaluating and optimizing different supply chain scenarios.

CONCLUSION

Solving the biggest artificial intelligence challenges

Most software companies have made investments and progress on their AI journeys and are exploring the possibilities of generative AI. But many find themselves hitting stopgaps along the way, worried that costs and complexities will grow too high as they progress.

Throughout this eBook, we explored the steps to forge ahead and realize the full power of generative AI.

To recap, let's look at the biggest challenges we identified along the way, along with a brief recommendation of how your business can solve them.

To learn more about how you can overcome obstacles and accelerate your AI journey, visit the [AWS AI Resource Hub](#).

To learn more about how generative AI can boost productivity, build differentiated experiences, and innovate faster for every businesses, visit the [AWS Generative AI Homepage](#).



Challenge	Solution
Discouraging failures	Developing a fault-tolerant culture.
Siloed, unprocessed data	Creating a modern data strategy that includes data lakes.
Finding the right business problems	Building blended teams that include both technical and domain experts.
The AI skills gap	Adopting new organizational models, processes, and team management philosophies.
Sustainably scaling beyond pilot projects	Leveraging end-to-end tools like Amazon Bedrock and SageMaker to build and scale generative AI applications.
Measuring the results	Forgoing traditional ROI metrics in favor of agility, competitive advantage, and risk tolerance using the value tree model.

Let's get started

To learn more about how your software company can overcome obstacles and accelerate your AI journey, visit the [AWS AI Resources Hub](#).

Visit us ›