



6 steps to success with generative AI

A practical guide to make your artificial
intelligence vision a reality



Table of contents

INTRODUCTION

Forging ahead	3
What are 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	10
Step 4: Upskilling teams.....	12
Step 5: Scaling beyond pilot projects.....	15
Step 6: Measuring the results	17

CONCLUSION

Taking the next step with AWS	19
Solving the biggest artificial intelligence challenges	21

INTRODUCTION


Forging ahead

When deployed with the right strategy, generative artificial intelligence (AI) has the potential to change the anatomy of work. By streamlining work activities that absorb 60–70 percent of employee time, organizations can improve productivity down to the individual worker.¹ This can increase agility, streamline processes, boost revenue, and allow for better, faster decision making.

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 the 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 AWS best practices and its experience helping thousands of customers realize their own initiatives.

**Generative AI could add
\$2.6
TRILLION
to
\$4.4
TRILLION
in economic benefits annually**

The impact of generative AI on productivity could add trillions of dollars in value to the global economy, increasing the effect of all AI by 15%–40%¹



What are 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 businesses should begin their 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 will 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 employee productivity

Generative AI is being leveraged for its transformative value to help AWS customers reach new levels of productivity. The latest advancements in generative AI can be used to increase employee productivity with the help of generative AI-powered conversational search, content creation, text summarization, and code generation.

Improving customer experiences

Today's organizations can take advantage of generative AI to improve 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, and personalization. **FOX Sports** improved its customer experience by building an end-to-end ML solution with the **Amazon Machine Learning Solutions Lab** that converts in-game tabular stats into natural-sounding narratives to improve the fan experience.³

Transforming content creation for greater creativity

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

Driving outcomes for real, competitive advantage

Regardless of an organization's field and domain, generative AI should be embedded into every business plan to stay relevant and competitive. **LG AI Research** successfully deployed a 300-billion-parameter FM that uses both images and text data.⁴ The multipurpose model can carry out a range of tasks, helping to boost LG AI Research's competitive advantage across multiple industries.

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 must take place in order to successfully implement, deploy, and scale AI solutions.



³ Wang, H., et al., "Enhance sports narratives with natural language generation using Amazon SageMaker," AWS Machine Learning Blog, May 2021

⁴ "LG AI Research Develops Foundation Model Using Amazon SageMaker," AWS Customer Story, 2023

The generative AI journey

The generative AI journey is not necessarily a straightforward path. Achieving success with generative AI requires more than great technology—it also means ensuring 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 must 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 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 minimizing 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 organizations must undergo is utilizing the opportunity inherent in mistakes. 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.

As new generative AI applications are deployed, companies should provide training to employees on how to use generative AI responsibly and effectively. Companies should also have clear policies guiding appropriate usage, while leaders should encourage experimentation with generative AI while managing risks. When implemented thoughtfully, generative AI can augment human capabilities and unlock innovations.

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 provides 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. Amazon Bedrock also is the easiest way for customers to build and scale generative AI-based applications using FMs. Amazon 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.

Amazon uses AI throughout its fulfillment process,

leveraging 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. And 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

Like with ML, data is an important piece of generative AI. While general large language models (LLMs) can be used “out of the box” for some use cases, when you want to build generative AI applications that are unique to your business needs, your organization’s data is your strategic asset. Data is the difference between a generic generative AI application and one that truly knows your business and your customer in order to deliver a better, more differentiated experience.

While some companies will build and train their own LLMs with vast amounts of data, many more will use their organizational data to fine-tune FMs for their unique business needs or to add context to prompts through Retrieval Augmented Generation (RAG). For example, you can use data from sources such as your data lake, database, and data warehouse to create a chatbot that can provide technical support for your products or a model that provides marketing copy trained on your best-performing ads.

Because of this, success with generative AI requires relevant, high-quality data, which means that you need a strong data strategy in the cloud. According to [McKinsey Digital](#), “...companies that have not yet found ways to effectively harmonize and provide ready access to their data will be unable to fine-tune generative AI to unlock more of its potentially transformative uses.”

The right data strategy for generative AI includes a comprehensive set of services to store and query data at scale, breaks down silos so you have ready access to all of your data to leverage for generative AI applications, and makes sure your data is secured and governed throughout the lifecycle of building generative AI applications.

Discover how a [data foundation built on AWS](#) gives you a strategic advantage when it comes to generative AI.

“We are extremely excited about the prospect of Amazon Bedrock and what it can mean to our business. We prioritize data integrity and security as we create applications that drive innovation for our end customers. As we continue to implement artificial intelligence in everything we do, it is important that we have good data to build our models around. With Amazon Bedrock, we are confident in its ability to create synthetic images and data to train our models in a responsible way that protects sensitive customer data and keeps us compliant with our highest standards, such as HIPAA, PII, and Hightrust.”

[Ashok Shenoy](#), VP of Portfolio Development,
Ricoh USA

STEP 3

Finding the right business problem to address

One mistake organizations 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. With no specific business problem to solve, IT executives will find it increasingly difficult to demonstrate the value of AI projects and initiatives to their business executive counterparts. This can stall or even stop progress on AI initiatives.

Here are some important questions 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.

For more on measuring the success of AI initiatives, refer to [Step 6](#) in this eBook.

Organizations can also 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.

Game-changer: NFL's AI-driven future

In the dynamic world of professional football, the quest to gain a competitive edge has shifted from the weight room to the data room. The National Football League (NFL), in its ongoing pursuit of excellence, has embraced cutting-edge AI and data science techniques to extract meaningful insights from its troves of raw data.

As part of its ongoing commitment to player health and safety, the NFL teamed up with AWS to create the Digital Athlete program. It uses AI to run virtual representations of NFL players through endless simulations of in-game scenarios, revealing new insights into player injury prevention. Through continual iteration, these AI models become more accurate as time goes on. Just as a quarterback does countless reps that generate gradual improvements to their technique, an AI model can input increasingly more video footage and data and tweak its parameters to generate better insights. In both cases, practice and repetition are key. This is leading to advancements such as better risk analysis for injury, more specialized player training and recovery regimens, improvements to protective equipment, and new game rules.

As we've seen through the AWS and NFL's Next Gen Stats (NGS) partnership, deciphering the intricate patterns of player movement, actions, and interactions on the field is essential for both strategic analysis and enhancing the fan experience. Leveraging Radio Frequency Identification (RFID) tags to track player movement, NGS provides real-time location data, speed, and acceleration for every player during every play on every inch of the field. This translates into over 300 million data points stored and processed per season.

When building new stats that uncover hidden dynamics in the game, the NGS team leans on the AWS ML teams to bring novel technology solutions and techniques to pair with their football expertise. To date, AWS and the NGS team have collaborated to build more than 20 different ML-powered statistics.

One example is Pressure Probability, **a new AI-powered stat** that changes how we evaluate pass rushers and offensive linemen. Unlike traditional stats that only note



whether a pass rusher succeeded in sacking the quarterback, pressure probability delves into the details: how close the defender got, how fast they were moving, and even how quickly the quarterback responded. With this granular data, powered by Next Gen Stat player-tracking data, the stat not only quantifies pressure but also captures the context, allowing analysts to dissect the strategies that influence each play.

The technical achievements unveiled by AWS scientists are just the tip of the iceberg. The journey of predicting pressure probabilities identifying blockers and rushers showcases the power of blending technical prowess with deep subject-matter expertise. The lessons learned, the models refined, and the insights derived through this data-driven journey open doors to a new era of football analysis that combines technological innovation and passion for the game.

[Read the full story ›](#)

STEP 4

Upskilling teams

In parallel with creating a comprehensive data strategy, today's organizations 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 buying or borrowing pretrained expert talent.



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.


Organizations also need to align teams to successfully tackle AI problems 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.

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 Morningstar did it

Investment research firm **Morningstar** uses AI to automate data collection processes and expand the number of funds it covers. The company does this by leveraging predictions from an AI model trained to emulate the fund evaluation process for its analysts.

To train its employees and accelerate AI application, Morningstar uses **AWS DeepRacer**—a tool that facilitates hands-on AI training through a fully autonomous 1/18th scale race car driven by reinforcement learning, a 3D racing simulator, and a global racing league. More than 445 Morningstar employees from multiple functions and eight countries—including 35 percent of its technology function—have been engaged in the DeepRacer League.⁷

Morningstar had dozens of AI projects and initiatives in the pipeline. These included a reinforcement learning program that searches for patterns in regulatory filings and an algorithm that identifies and fixes broken links to the websites of financial institutions.

[Read the full story ›](#)

“Our DeepRacer challenge harnesses our employees’ enthusiasm for machine learning and artificial intelligence. It provides hands-on training across the company and accelerates Morningstar’s practical application of machine learning across our investing products, services and processes. The response from our teams has gone well beyond my expectations, and it has been a fun way to unite our global teams, whether in technology or other functions.”

James Rhodes, CTO, Morningstar



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

STEP 5

Scaling beyond pilot projects

After the first few successful pilots, organizations 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 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 **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 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 and build, train, and deploy models to get them into production faster and at lower cost, enabling 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 the Allen Institute did it

Using AWS generative AI tools, the [Allen Institute](#) is mapping the entire human brain. In the process, its researchers are creating the largest open-source database in the world. Named the Brain Knowledge Platform, the project is helping to accelerate progress in treating brain diseases such as Alzheimer's and Parkinson's.

"We need this map to know exactly what goes wrong in disease and precisely where to target therapeutics," says Dr. Ed Lein, Senior Investigator for the Allen Institute. "Also, the generative AI side of things has some unusual possibilities for us. We'll be able to integrate data of many different types to begin to make inferences that human researchers are not capable of doing."

"I think we can imagine the new territory that we're entering with generative AI, as it may be applied to neuroscience. We can begin to understand how the brain actually works—and to harness this knowledge for the good of society."

Dr. Ed Lein, Senior Investigator, Allen Institute for Brain Science

[Watch the video ›](#)



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 business 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 LLMs. Initially, the company released two models: one with seven 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.

“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

[Read the case study ›](#)



CONCLUSION

Taking the next step with AWS

No matter where your organization 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 transform your business.

Generative AI with AWS, by the numbers

100,000+

customers using AWS for their AI workloads

20+ YEARS

of building experience at Amazon

Improve customer experiences



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



Agent assist and conversational analytics: Enhance agent performance to improve first contact resolution and augment tasks such as knowledge search, call summarization, and problem-solving. Managers can extract valuable insights to improve customer experience, monitor agent performance, and boost business performance.



Personalization: Deliver better-personalized experiences and increase customer engagement with individually curated offerings and communications.

Boost employee productivity



Conversational search: Improve employee productivity by quickly and easily finding accurate information and summarizing 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 automatically, saving time and reducing errors.

Enhance creativity and content creation



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



Sales: Generate personalized emails and messages based on prospect's profile and behavior, improving response rates. Generate sales scripts or talking points based on the customer's segment, industry, and product or service.



Product development: Generate multiple design prototypes based on certain inputs and constraints, speeding up the ideation phase or optimizing 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 organizations have made investments and progress in 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 organization 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](#).

[Get started](#) ›

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