

# Driving Retail and Consumer Goods innovation with generative AI

Transform and modernize faster—start your  
generative AI journey with AWS



# Table of contents

<b>Introduction: Generative AI in Retail and Consumer Goods</b> .....	<b>3</b>
<b>How to start your generative AI journey</b> .....	<b>4</b>
<b>Examining the challenges of adopting generative AI</b> .....	<b>5</b>
<b>Generative AI in action: Retail and Consumer Goods use cases</b> .....	<b>6</b>
Use case 1: Customer-centric .....	7
Use case 2: Product-centric .....	8
Use case 3: Employee-centric .....	9
Use case 4: IT-centric.....	10
<b>How to select the right tools and infrastructure to get started</b> .....	<b>11</b>
<b>Put generative AI to work with the AWS Partner Network</b> .....	<b>17</b>
<b>Get started</b> .....	<b>21</b>

## INTRODUCTION

# Generative AI in Retail and Consumer Goods

Retail and Consumer Goods companies of all sizes are getting started with generative artificial intelligence (AI). They, like you, want to capitalize on it and translate the momentum from betas, prototypes, and demos to real-world innovations and productivity gains.

Amazon Web Services (AWS), born from Amazon, the most innovative retailer in the world, is building upon 25 years of experience innovating with AI technologies to offer enterprise-grade generative AI applications and infrastructure. To foster customer success, AWS also partners with leading technology and consulting companies that offer purpose-built AI solutions for Retail and Consumer Goods companies.

In this ebook, you can explore the Retail and Consumer Goods generative AI use cases that are most relevant to your business. Discover how AWS and our partners can deliver everything you need to accelerate generative AI-powered innovation and identify actionable next steps to get started or accelerate your generative AI journey.

# \$400B– \$660B

is the additional value that McKinsey estimates generative AI could deliver in the Retail and Consumer Goods industries annually<sup>1</sup>

In 2024, Retail and Consumer Packaged Goods executives will invest in AI tools for:<sup>2</sup>

**56%**  
marketing

**39.2%**  
customer service and user experience

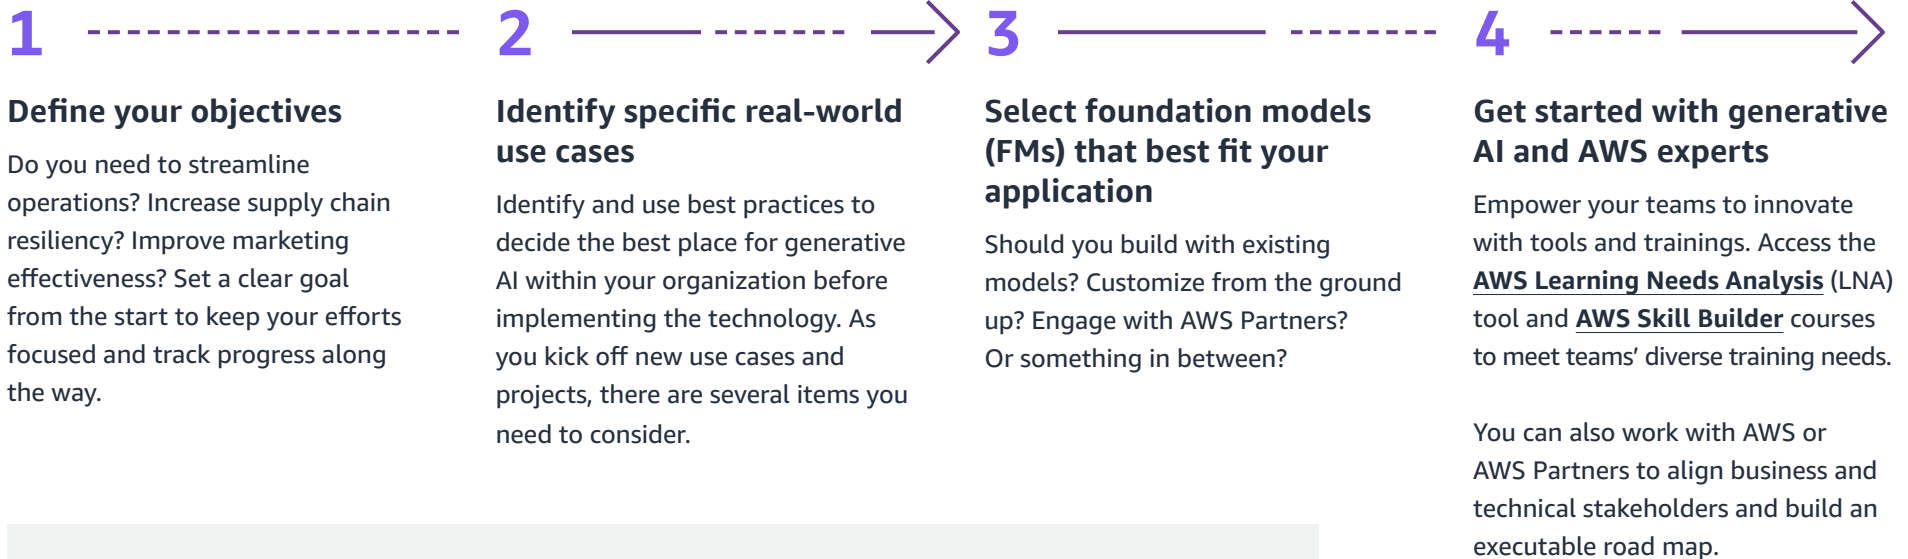
**34.4%**  
supply chain

<sup>1</sup> "The economic potential of generative AI: The next productivity frontier," McKinsey Digital, June 2023

<sup>2</sup> "Innovation in the Retail and CPG Sector Provides a Cause for Optimism as We Head into 2024," Carl Marks Advisors, January 2024

# How to start your generative AI journey

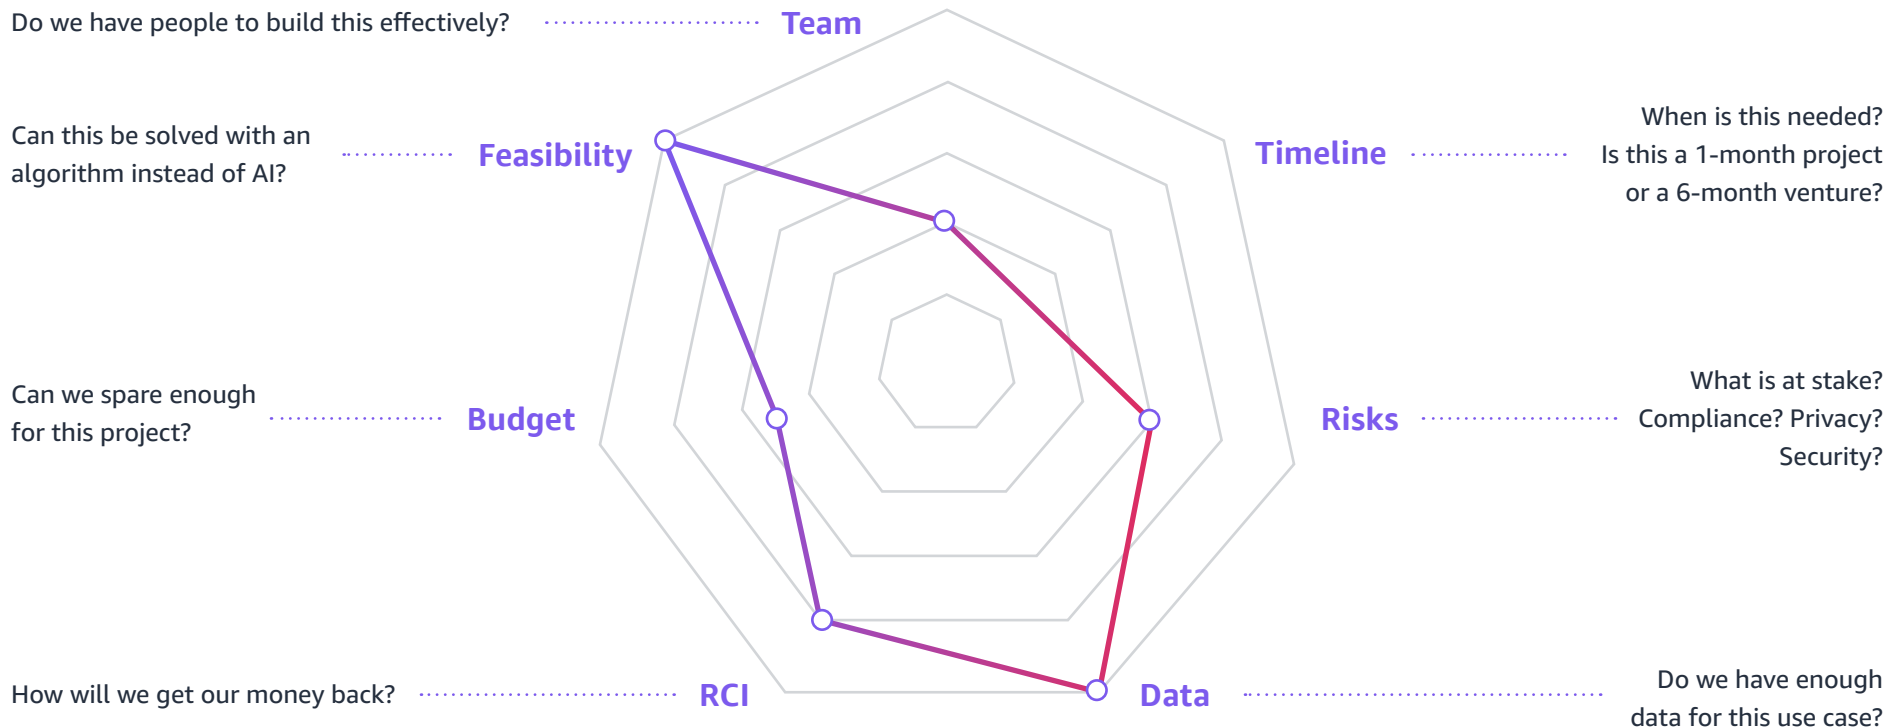
Successfully adopting and unlocking the benefits of generative AI requires the right strategy. The following steps can help your organization get off to a good start:



AWS is democratizing generative AI so that Retail and Consumer Goods companies of any size can reinvent their products, processes, and experiences.

# Examining the challenges of adopting generative AI

The race is on to unlock the business value and competitive advantages of generative AI for the Retail and Consumer Goods industries. However, organizations are seeking guidance as they embark on this journey. There are several details you will need to consider. Make sure to have conversations about both the technology and the impacts, risks, data, budget, and team. Include all these factors in the selection criteria, success measurements, and project planning. Read on to explore Retail and Consumer Goods use cases that illustrate how AWS can help your organization quickly realize the benefits of adopting generative AI technology to keep pace with or surpass the competition.





## USE CASES

# Generative AI in action: Retail and Consumer Goods use cases

AWS collaborates with retailers to envision the final solution to a problem first and then works backwards to identify the tasks needed to achieve their business objectives. Read on to learn more about a variety of use cases.



### Customer-centric ›

#### Contact Center

- Conversation analysis
- Q&A chatbots

#### Marketing

- Recommendations
- Store analytics
- Social campaigns
- Tagging and SEO
- Hyper-personalized ads

#### Shopping

- Conversational search
- Personal stylist
- Virtual try-on and fit
- Voice commerce



### Product-centric ›

- Product copy generation
- Manipulate product images
- Assisted product ideation and prototyping
- Customer review analysis
- Price matching
- Data quality and completion
- Economic risk analysis
- Personalized product pages
- Product quality



### Employee-centric ›

- Report summarization
- Business process automation
- Knowledge management
- Planogram design
- Robotic process automation
- Form generation
- Predictive maintenance and repair advice



### IT-centric ›

- Code generation
- Cloud formation generation
- Text to SQL

## USE CASE 1

# Customer-centric

Generative AI can be used to improve the customer experience, from marketing to shopping to customer support.

## Marketing

Generative AI excels at summarizing data, which can help you analyze customer feedback and social media analytics to develop new product and campaign ideas. It can also be used to generate engaging, hyper-personalized copy and imagery that can be targeted to your ideal customers.

## Shopping

Returns can quickly eat away at your margins. Using generative AI, you can help customers find the right products with AI-generated personal stylists and virtual try-ons. You can also make the shopping experience more seamless with voice commerce and digital shopping assistants.

## Customer support

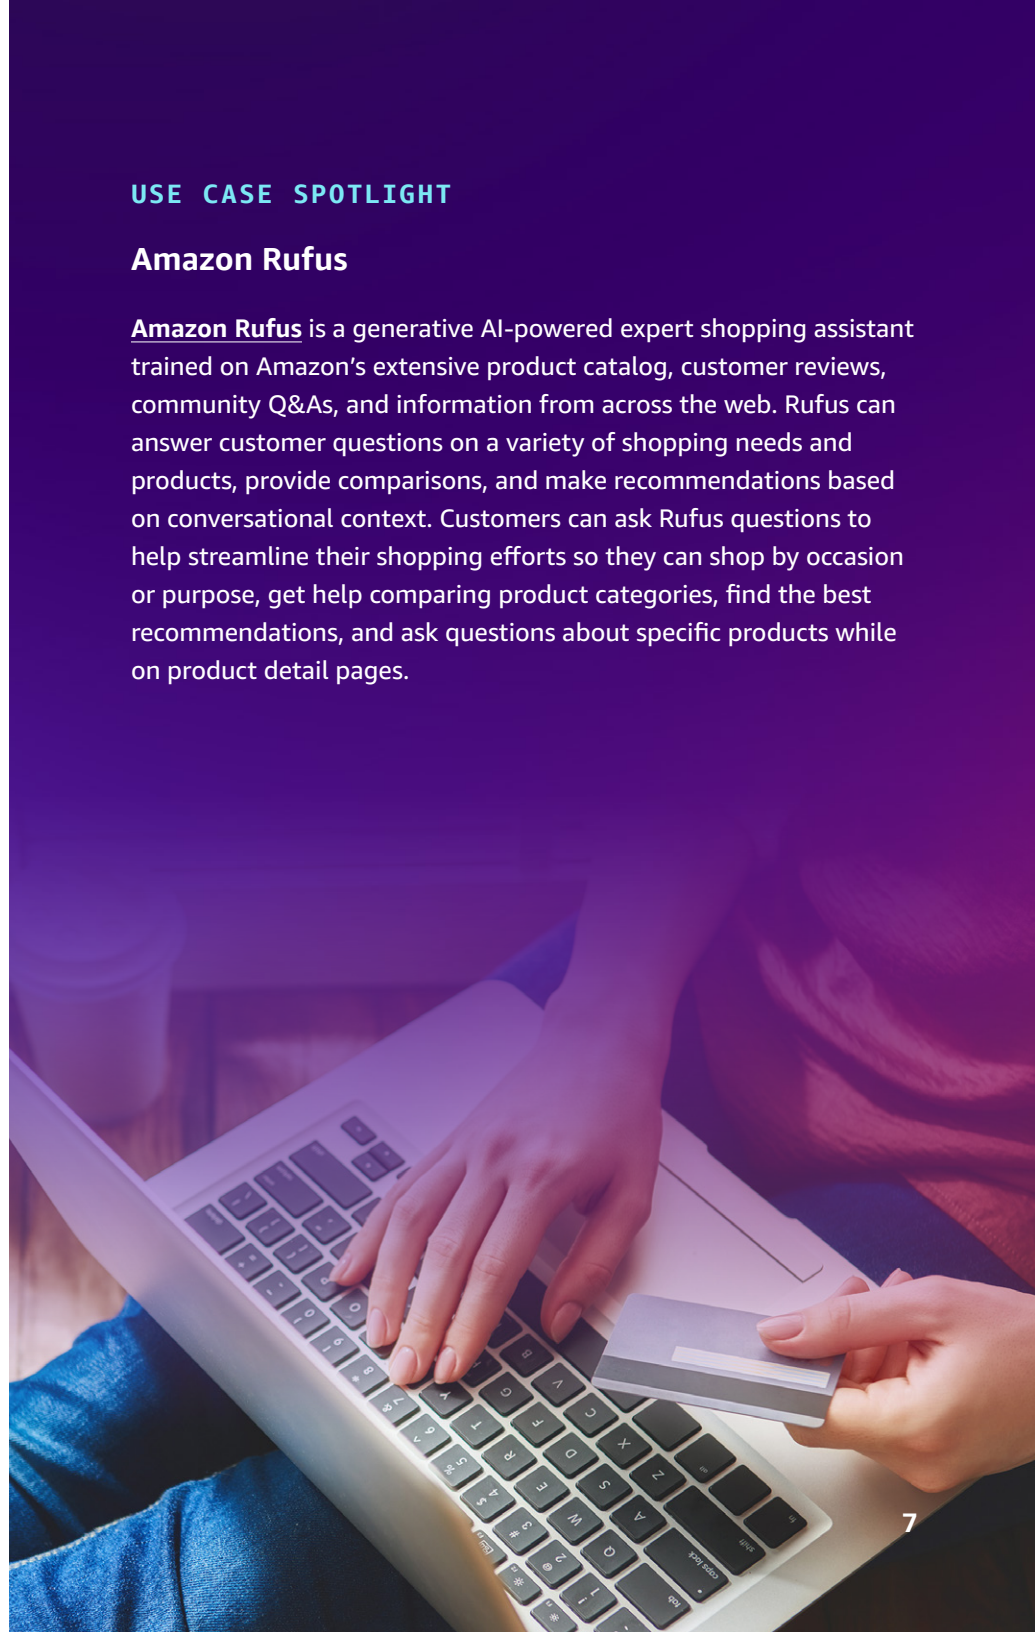
Generative AI can be used to improve the customer experience for marketing, shopping, and customer support. You can use generative AI to deliver suggested responses and actions to customer service agents, providing faster issue resolution and improved customer satisfaction. Or you can create chatbots that use generative AI to help customers find answers using voice or text.

Answer questions quickly with chatbots, analyze contact center calls, or provide agents with recommended responses and actions.

## USE CASE SPOTLIGHT

### Amazon Rufus

**Amazon Rufus** is a generative AI-powered expert shopping assistant trained on Amazon's extensive product catalog, customer reviews, community Q&As, and information from across the web. Rufus can answer customer questions on a variety of shopping needs and products, provide comparisons, and make recommendations based on conversational context. Customers can ask Rufus questions to help streamline their shopping efforts so they can shop by occasion or purpose, get help comparing product categories, find the best recommendations, and ask questions about specific products while on product detail pages.



## USE CASE 2

# Product-centric

The same summarization models that are helpful in analyzing customer sentiment can also help Consumer Goods companies identify new product ideas. Using generative AI, brands can accelerate time to market by generating new product ideas, renderings, and prototypes. Generative AI can also be used to develop new packaging ideas and color and flavor combinations. You can then use it to ensure product quality meets expectations, set pricing, and develop personalized product pages—all of which can help you meet customers' expectations and drive sales.

## USE CASE SPOTLIGHT

### The Very Group

The Very Group collaborated with the AWS Generative AI Innovation Center to build a system that uses Amazon Bedrock, large language models (LLMs), and multi-modal models to function as an intelligent product analyzer and description writer for copywriters. The solution reduced the time employees spent on creating and checking product descriptions across their ecommerce platforms. The time savings has enabled The Very Group to have a higher and more accurate completion rate of getting products to market across the entire product portfolio.

[Learn more >](#)

## USE CASE SPOTLIGHT

### adidas

adidas, one of the largest sports brands in the world, trained a stable diffusion algorithm on 150,000 shoe images at different angles. Now, employees can ask the algorithm to help them generate a running shoe with certain criteria, like a collaboration with a partner or a mash-up of two shoe types, and it will generate ideas that their designers can choose from or use as inspiration to build a new shoe.

[Learn more >](#)



## USE CASE 3

# Employee-centric

By automating mundane, repetitive tasks with generative AI, companies can increase efficiency, improve employee retention, and boost quality. Retail and Consumer Goods companies, especially those with an existing data lake, can use generative AI to provide information to employees quickly.

Employees can ask chatbots questions like:

- What's the current open to buy (OTB) for SKU 1234?
- How can I repair machinery that isn't working?
- What was the best-selling shoe last June?

Providing a natural language interface for employees to ask questions, including follow-up questions, can enhance their ability to access data and empower them to make better decisions. In addition, you can help your employees build, discover, and share actionable insights and narratives in seconds using generative business intelligence (BI).

## USE CASE SPOTLIGHT

### adidas

Developers at adidas recently created a chatbot assistant to help new engineers get up to speed within the organization with the ability to ask questions like “How do I get an AWS account?” and “How do I get to a Kubernetes namespace?” With clearly defined resources already established, adidas was able to feed the data into **Amazon Titan** Embeddings and build an assistant on top using LangChain with **Amazon Bedrock** and Amazon Titan in the background. The chatbot can answer questions quickly and provide additional resources to the engineers. adidas also has a pilot underway with **Amazon Q Developer**, providing coding assistance to engineers to help them work faster and more efficiently.

[Learn more ›](#)

## USE CASE 4

# IT-centric

Companies are using generative AI to help programmers write code faster with fewer errors. Because generative AI is trained on billions of lines of code, it can generate code suggestions ranging from snippets to full functions in near real time based on your comments and existing code. With this service, you can bypass time-consuming coding tasks and accomplish complex coding tasks with unfamiliar frameworks, APIs, or SDKs.

Generative AI can also flag or filter code suggestions that resemble publicly available code. Additionally, you can scan your code to detect hard-to-find vulnerabilities and quickly get code suggestions to remediate them.

**“We were excited to be part of the Amazon Bedrock preview and get our hands on the service. This service quickly became a highly valued addition to our generative AI tool kit, empowering us to focus on the core aspects of our LLM projects while letting it handle the heavy lifting of managing infrastructure. Using Amazon Bedrock, we have developed a generative AI solution that gives the community of adidas engineers the ability to find information and answers from our knowledge base through a single conversational interface, covering everything from getting started to highly technical questions.”**

Markus Rautert, VP of Platform Engineering & Architecture, adidas




# How to select the right tools and infrastructure to get started

Once you have set objectives and narrowed down your use case, you can:


- 1 Easily build and scale generative AI applications with security and privacy
- 2 Benefit from the most performant, low-cost infrastructure for generative AI
- 3 Transform user experiences with generative AI-powered applications
- 4 Leverage data as a differentiator

## AWS Generative AI stack


### APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs



### TOOLS TO BUILD WITH LLMs AND OTHER FMs



### INFRASTRUCTURE FOR FM TRAINING AND INFERENCE





# 1. Easily build and scale generative AI applications with security and privacy

From day one, AWS makes it possible for organizations of all sizes and developers of all skill levels to build and scale generative AI applications with security, privacy, and responsible AI built in. With AWS, customers can access leading FMs, customize with their own data, and use the leading security, access control, and features they trust from AWS.

AWS marquee services include [Amazon Bedrock](#), [Amazon SageMaker](#), [AWS AI Service Cards](#), and [Amazon Simple Storage Service](#) (Amazon S3).

## Amazon Bedrock

Amazon Bedrock is a fully managed service that offers a choice of high-performing FMs from industry-leading AI companies like AI21 Labs, Anthropic, Cohere, Meta, Mistral AI, Stability AI, and Amazon through a single API, along with a broad set of capabilities you need to build generative AI applications with security, privacy, and responsible AI. Using Amazon Bedrock, you can easily experiment with and evaluate top FMs for your use case, privately customize them with your data using techniques such as fine-tuning and retrieval augmented generation (RAG), and build agents that execute tasks using your enterprise systems and data sources. Because Amazon Bedrock is serverless, you don't have to manage any infrastructure and you can securely integrate and deploy generative AI capabilities into your applications using the AWS services you are already familiar with.

**“Thousands of customers have implemented Databricks on AWS, giving them the ability to use MosaicML to pre-train, fine-tune, and serve foundation models for a variety of use cases. AWS Trainium gives us the scale and high performance needed to train our Mosaic MPT models, and at a low cost. As we train our next generation Mosaic MPT models, Trainium2 will make it possible to build models even faster, allowing us to provide our customers unprecedented scale and performance so they can bring their own generative AI applications to market more rapidly.”**

Naveen Rao, VP of Generative AI, Databricks

More than 10,000 organizations worldwide including Comcast, Condé Nast, and over

**60%** OF FORTUNE  
**500 COMPANIES**

rely on Databricks to unify their data, analytics, and AI<sup>3</sup>

## 2. Benefit from the most performant, low-cost infrastructure for generative AI

Whether customers are training their own models, customizing models, or running machine learning (ML) applications, AWS is the best place to train and run inference at scale with infrastructure purpose-built for ML. From the highest performance GPU-based [Amazon Elastic Compute Cloud \(Amazon EC2\) P5 instances](#) to continued investments in our purpose-built accelerators [AWS Trainium](#) and [AWS Inferentia](#), customers get the most performant and low-cost infrastructure for generative AI.

AWS marquee services include [AWS Trainium](#), [AWS Inferentia](#), and [Amazon EC2 P5 instances](#).

### AWS Trainium

Although the use of deep learning and generative AI is accelerating, many development teams are limited by fixed budgets, which puts a cap on the scope and frequency of training needed to improve their models and applications. AWS Trainium-based [Amazon EC2 Trn1 instances](#) solve this challenge by delivering faster time to train while offering up to 50 percent cost-to-train savings over comparable Amazon EC2 instances. AWS Trainium has been optimized for training natural language processing (NLP), computer vision (CV), and recommender models used in a broad set of applications, such as text summarization, code generation, question answering, image and video generation, recommendation, and fraud detection.

**“With millions of listings posted daily, it is important that we continually improve our personalized search and recommendation experiences for our users. To achieve this goal, we are experimenting with Amazon Titan Multimodal Embeddings, with the aim of revolutionizing local commerce through cutting-edge semantic search capabilities. During an initial evaluation with the new multimodal model, we have observed substantial improvement in relevance recall for keyword searches. This advancement will significantly expedite successful matches, benefiting both our buyers and sellers.”**

Melissa Binde, CTO, OfferUp

[Read the customer story >](#)

[OfferUp](#) is one of the largest mobile marketplaces for local buyers and sellers in the US and is changing the way people transact in their communities by providing a uniquely simple and trusted experience.



# 3. Transform user experiences with generative AI-powered applications

At AWS, we build powerful new applications that help our customers boost productivity in the enterprise, streamline coding, simplify BI, and improve efficiency for organizations. With security and privacy built in, easy customization, and seamless data integration, enterprises can quickly take advantage of generative AI adapted to the specific needs of their organization.

AWS marquee services include [Amazon Q](#).

## Generative AI-powered applications from AWS

### Amazon Q

Amazon Q generates code, tests, and debugs and has multistep planning and reasoning capabilities that can transform and implement new code generated from developer requests. Amazon Q also makes it easier for employees to get answers to questions across business data—such as company policies, product information, business results, codebases, employees, and many other topics—by connecting to enterprise data repositories to summarize the data logically, analyze trends, and engage in dialogue about the data.

Products featuring Amazon Q include [Amazon Q Business](#), [Amazon Q Developer](#), [Amazon Q in QuickSight](#), [Amazon Connect](#), and [AWS Supply Chain](#).

## Amazon Q Developer

Amazon Q Developer is trained on billions of lines of code and can generate code suggestions ranging from snippets to full functions in real time based on your comments and existing code. Bypass time-consuming coding tasks and accelerate building with unfamiliar APIs.

## PartyRock

**PartyRock**, an Amazon Bedrock playground, is a generative AI app-building playground that makes it easy and accessible for anyone to experiment hands-on with prompt engineering in an intuitive and fun way. In just a few clicks, you can build entertaining apps to explore the possibilities of generative AI. By building and playing with PartyRock apps, you will learn about the techniques and capabilities needed to make the most of generative AI, including understanding different models' strengths, experimenting with different text-based prompts, and chaining prompts together.

**“Our business is constantly evolving and developing new data needs, which led us to create and update dashboards and reports. QuickSight enables our operations teams to deliver data to users across a variety of use cases, from distribution-center forecasts to reporting Amazon Connect call-center metrics. QuickSight Q has shown us the power of natural-language experiences in to accelerate data work by helping our business users get insights instantly. We are excited to see the additional Generative BI capabilities for authors raise our speed to respond to these changing business needs to a new level. Natural-language experiences like these are changing the way people work.”**

Corey Savory-Venzke, VP of Customer Experience, Traeger Grills

**Traeger Grills** is a leading provider of smokers, grills, and barbecue products.



## 4. Leverage data as a differentiator


Data is the difference between a general generative AI application and one that truly knows your business and your customers. With AWS, it's easy to use your organization's data as a strategic asset—customize generative AI applications and models and build more differentiated experiences using features like the customization capability in Amazon Q Developer. Use other built-in tools, like [Agents for Amazon Bedrock](#), to build applications that know your business, data, and customers.

With the most comprehensive, integrated set of data and AI services for all workloads, use cases, and types of data and the tools to govern that data, AWS is the best place to build a data strategy for your generative AI application.

AWS marquee services include [Amazon Bedrock](#) and [Amazon S3](#).

### Putting your data to work

Generative AI applications require operational databases to support the user experience. Analytics and data lakes are where you accumulate your domain-specific data. These tools can help you explore data and understand how to use it in generative AI. Data integrations are required to source your data, and setting up pipelines enables you to keep up with changing data so it becomes usable for generative AI. Another key component to consider is governance, which includes processes to ensure data quality, privacy and compliance with privacy laws, and security and access controls.



## FEATURED AWS PARTNERS

# Put generative AI to work with the AWS Partner Network

AWS Retail and Consumer Goods Competency Partners help drive pivotal advancements with the services, tools and infrastructure for implementing generative AI to help boost productivity, build differentiated experiences, and accelerate innovation.

The AWS Partners featured on the following pages can help retailers and brands unlock greater business value with AI across multiple solution areas including product development, manufacturing, supply chain, unified commerce, and more.

[AWS Partner solutions ›](#)

FEATURED AWS PARTNERS

# Customer-centric AWS Partner Solutions

## Conversational search



**XGEN<sup>AI</sup>**

## Knowledge management



## Assisted product ideation and prototyping



## System integrators



**slalom**

**Deloitte.**



FEATURED AWS PARTNERS

# Product-centric AWS Partner Solutions

Product content  
generation



Product  
recommendation



Image and video  
production



Demand planning  
and price  
optimization



Data quality and  
completion



FEATURED AWS PARTNERS

# Employee-centric AWS Partner Solutions

## Data insights

**ADA**ASTRA

## Predictive maintenance and repair advice

BRAINBOX **AI**

## Knowledge management

fractal

TREDENCE  
Beyond Possible

# Get started

As you begin, align your business and technology teams to prioritize the most impactful use cases, taking advantage of AWS workshops that facilitate this. Work toward a proof of concept with the help of AWS experts, and empower your developers through training opportunities.

If you've already launched a proof of concept, make sure to measure and track the business value and ROI, plan to monitor and optimize over time as technology advances, and put the right infrastructure in place to scale. Lastly, establish compliance and governance to ensure the technology is being used responsibly.

**Contact an [AWS for Retail](#) or [Consumer Goods](#) expert to learn how to grow your business:**

**[Discover more about AWS for Retail and Consumer Goods](#) ›**

**[Learn about generative AI for Retail and Consumer Goods](#) ›**

**Explore AWS Partners that have deep technical knowledge:**

**[Find AWS Retail Competency Partners](#) ›**

**[Find AWS Consumer Goods Partners](#) ›**

Follow [AWS for Retail and Consumer Goods LinkedIn page](#) to learn more about AWS customer stories, industry use cases, events, and more.