



Building AI and ML powered applications without machine learning expertise

Joseph Sadler

Sr. Solutions Architect
Amazon Web Services

At **AWS**, our goal is to put **machine learning**
in the hands of **business teams**





**Enhance
customer
experience**



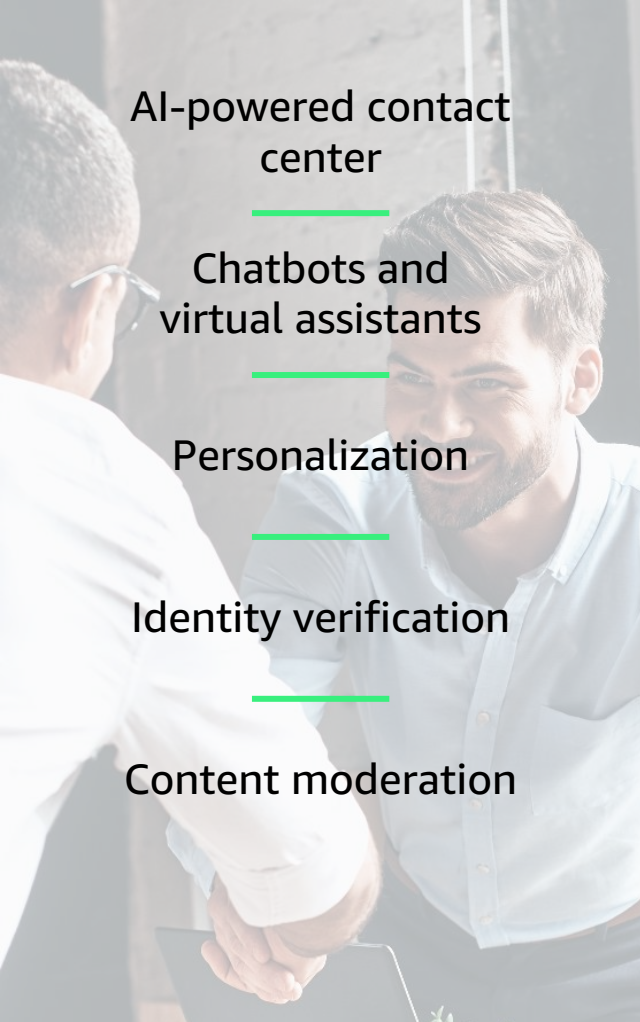
**Augment human
ingenuity**



**Improve business
operations**



**New products
and services**




AI-powered contact
center

Chatbots and
virtual assistants

Personalization

Identity verification

Content moderation



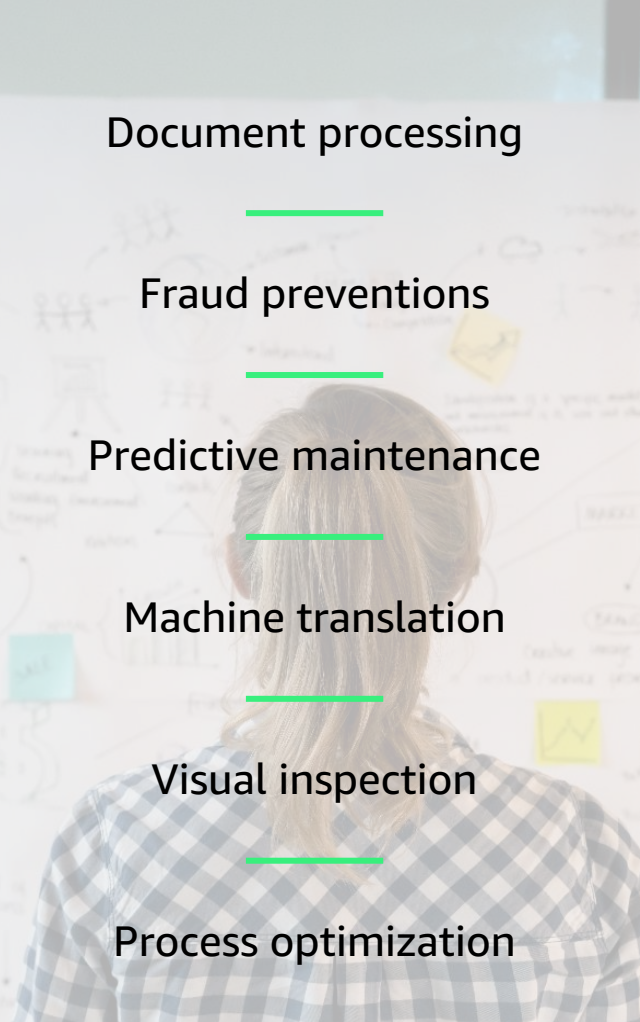
Employee assistants

Content generation

Forecasting

Code generation

Report generation



Document processing

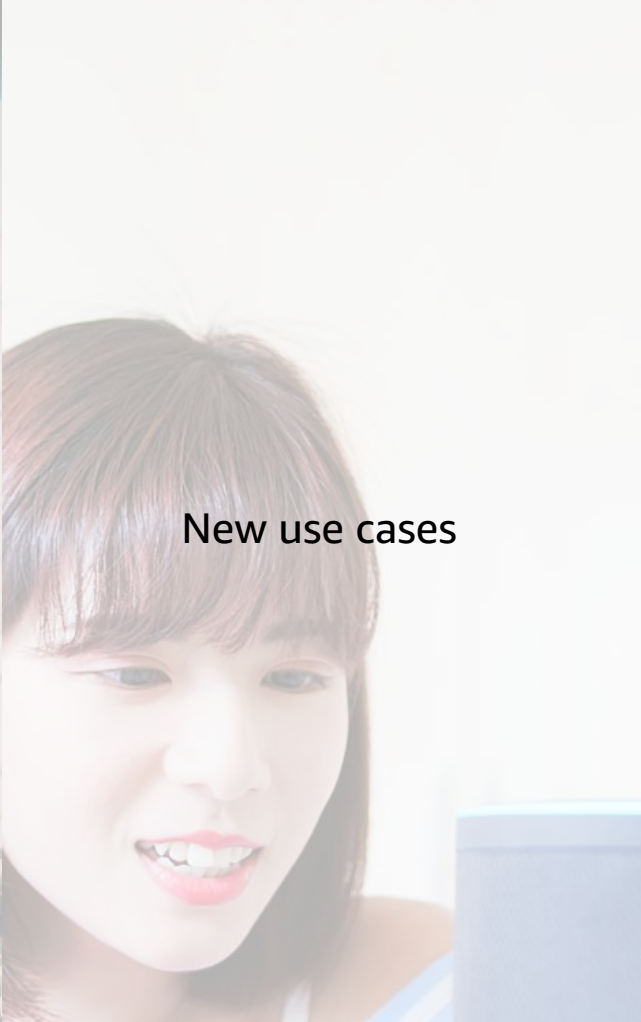
Fraud preventions

Predictive maintenance

Machine translation

Visual inspection

Process optimization



New use cases

**Enhance
customer
experience**

**Augment human
ingenuity**

**Improve business
operations**

**New products
and services**

Low-code/ No-code AWS AI/ML Services

EASILY ADD INTELLIGENCE TO YOUR BUSINESS APPLICATIONS

AI SERVICES

SPECIALIZED

BUSINESS PROCESSES

Amazon Personalize

Amazon Forecast

SEARCH

Amazon Kendra

CONVERSATION

Amazon Lex

Moderation

Amazon Augmented AI

CORE

TEXT

Amazon Translate

Amazon Comprehend

SPEECH

Amazon Polly

Amazon Transcribe

VISION

Amazon Textract

Amazon Rekognition

AWS Panorama

ML Services

AMAZON SAGEMAKER CANVAS

No-code ML for business analysts

Generative AI

AMAZON Bedrock Studio

Build Gen AI applications faster and more securely

Amazon Q

A gen AI-powered assistant

Amazon Q Business

Amazon Q Developer

Amazon Q in QuickSight

Amazon Q in Connect

AI Enhanced Services

AMAZON CONNECT

AI Enhanced Contact center



Example use cases with AI Services

Example Use Cases

Conversational AI

Intelligent Search

Identity Verification

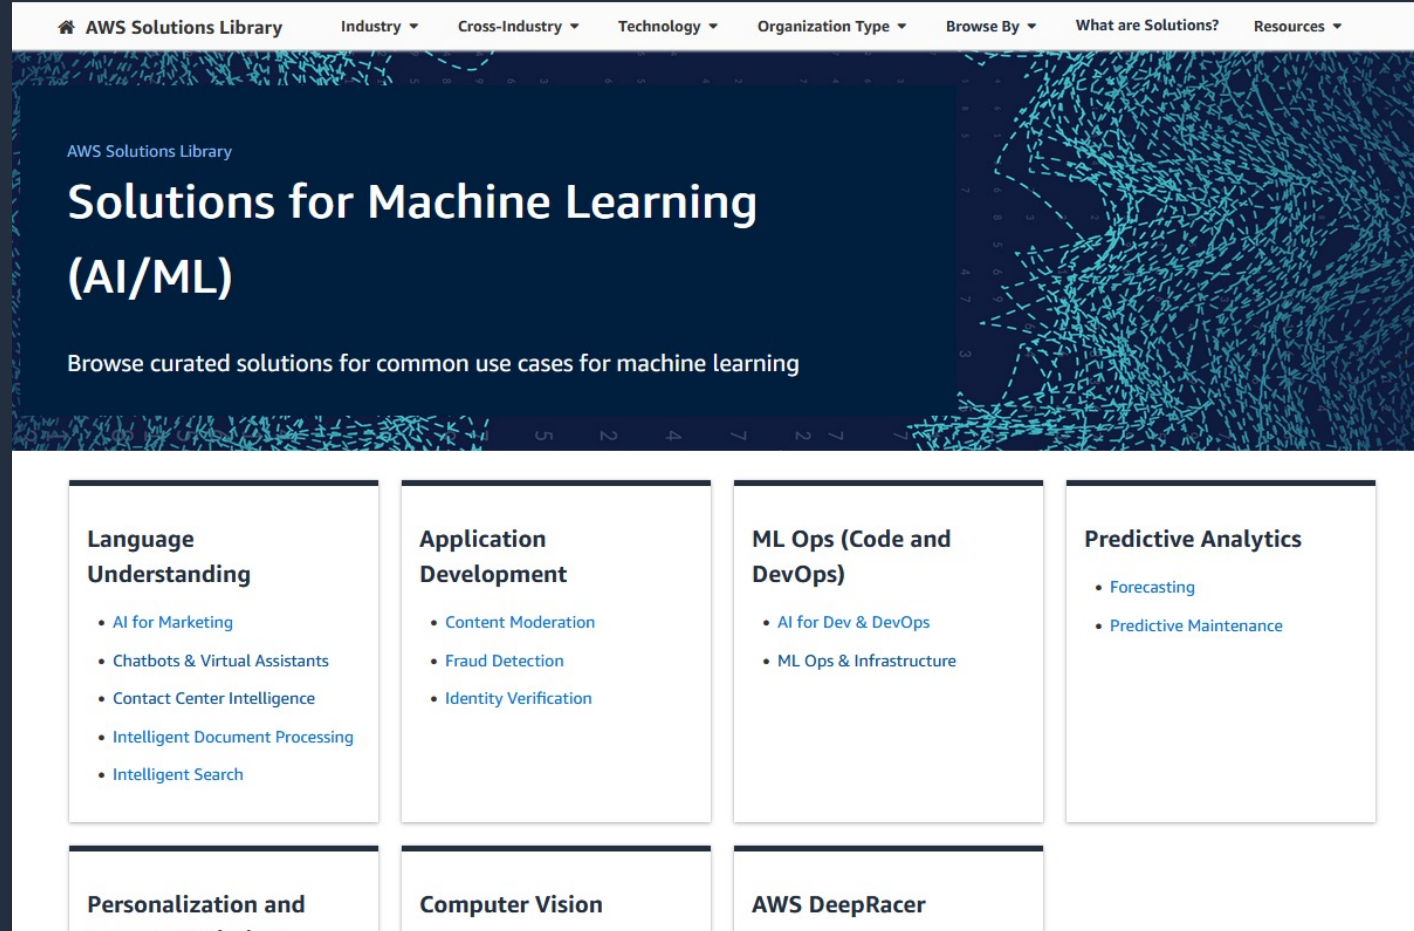
Personalization

Intelligent Document

Processing (IDP)

Forecasting

AWS AI/ML Solutions Library - Solutions



The screenshot shows the AWS Solutions Library website. The header includes navigation links: AWS Solutions Library, Industry, Cross-Industry, Technology, Organization Type, Browse By, What are Solutions?, and Resources. The main heading is "Solutions for Machine Learning (AI/ML)" with a subtext "Browse curated solutions for common use cases for machine learning". Below this, there are four columns of solution categories:

- Language Understanding**
 - AI for Marketing
 - Chatbots & Virtual Assistants
 - Contact Center Intelligence
 - Intelligent Document Processing
 - Intelligent Search
- Application Development**
 - Content Moderation
 - Fraud Detection
 - Identity Verification
- ML Ops (Code and DevOps)**
 - AI for Dev & DevOps
 - ML Ops & Infrastructure
- Predictive Analytics**
 - Forecasting
 - Predictive Maintenance

Below these columns, the following categories are partially visible:

- Personalization and Recommendation
- Computer Vision
- AWS DeepRacer

aws.amazon.com/solutions/ai-ml/



Conversational AI

- Voice and text-based interfaces to enhance end user experience and increase customer satisfaction

Use case categories

- Self-service bots
- Virtual agents and assistants
- Transactional bots
- Proactive help based on usage behavior

Benefits

- Enable new ways of engagement
- Increase customer satisfaction
- Reduce operational costs
- Streamline business processes

Underlying AI services



Amazon
Lex
Conversational AI



Amazon
Polly
Text to speech



Amazon
Kendra
Intelligent
Search



Amazon
Comprehend
Natural language
processing



Amazon
Translate
Machine
translation

Solution

QnABot on AWS



Intelligent search

- Get the most out of your data with Intelligent Search

Use case categories

- Improve customer interactions
- Accelerate research and development
- Minimize regulatory & compliance risks
- Equip employees with the data they need

Benefits

- Find answers faster across unstructured content
- Boost workforce productivity
- Enhance customer experiences
- Centralize access to knowledge

Underlying AI services



Amazon
Kendra
Intelligent
Search

Identity verification

- Enable secure and compliant digital experiences

Use case category

- Customer onboarding & verification
- Online proctoring
- Gig economy verification
- Event & airport check-ins
- Use challenges

Benefits

- Reduce onboarding friction
- Reduce fraud
- Lower costs and overheads

Underlying AI services



Amazon
Rekognition
AI Computer
Vision



Amazon
Textract
Intelligent Text
Extraction

Guidance

[Guidance for Identity Verification](#)



Personalization

- Generate personalized recommendations to increase customer engagement

Use case categories

- Deliver unique homepage experiences
- Help customers discover products faster
- Target customers more accurately
- Highlight new products, content, and promotion offerings

Benefits

- Implement a personalization engine in days, not months—no ML expertise required
- Adapt recommendations in real time
- Increase engagement and revenue through relevancy

AI services



Amazon
Personalize
AI personalization

ML Services



Amazon
SageMaker

Solution

Maintaining Personalized
Experiences
with Machine Learning

Solution

Personalized
recommendations



Intelligent document processing

- Make faster decisions by automatically extracting and analyzing data from documents

Use case categories

- Insurance forms text extraction
- Mortgage applications data extraction
- Medical forms processing
- Financial

Benefits

- Higher accuracy of data
- Faster data processing
- Improve employee productivity
- Cost savings

AI services



Amazon
Comprehend
Natural language
processing



Amazon
Textract
Intelligent Text
Extraction



Amazon
Augmented AI
Procedural
Human Review

Solution

Document understanding solution

Solution

Extract and analyze
data from documents



Forecasting

- Forecast inventory, product demand, financial metrics, and workforce staffing

Use case category

- Retail inventory forecasting
- Supply chain demand planning
- Revenue and financial metrics forecasting
- Workforce planning and staffing

Benefits

- Grow retail sales through fewer stockouts
- Lower costs by reducing wasted inventory
- Increase profitability through improved product allocation
- Optimize workforce productivity & staffing

AI Services



Amazon
Forecast

Guidance

Improving Forecast
Accuracy with ML

Solution

Demand
Forecasting

Use Cases Wrap Up

Conversational AI

Intelligent Search

Identity Verification

Personalization

Intelligent Document

Processing (IDP)

Forecasting

Amazon SageMaker Canvas

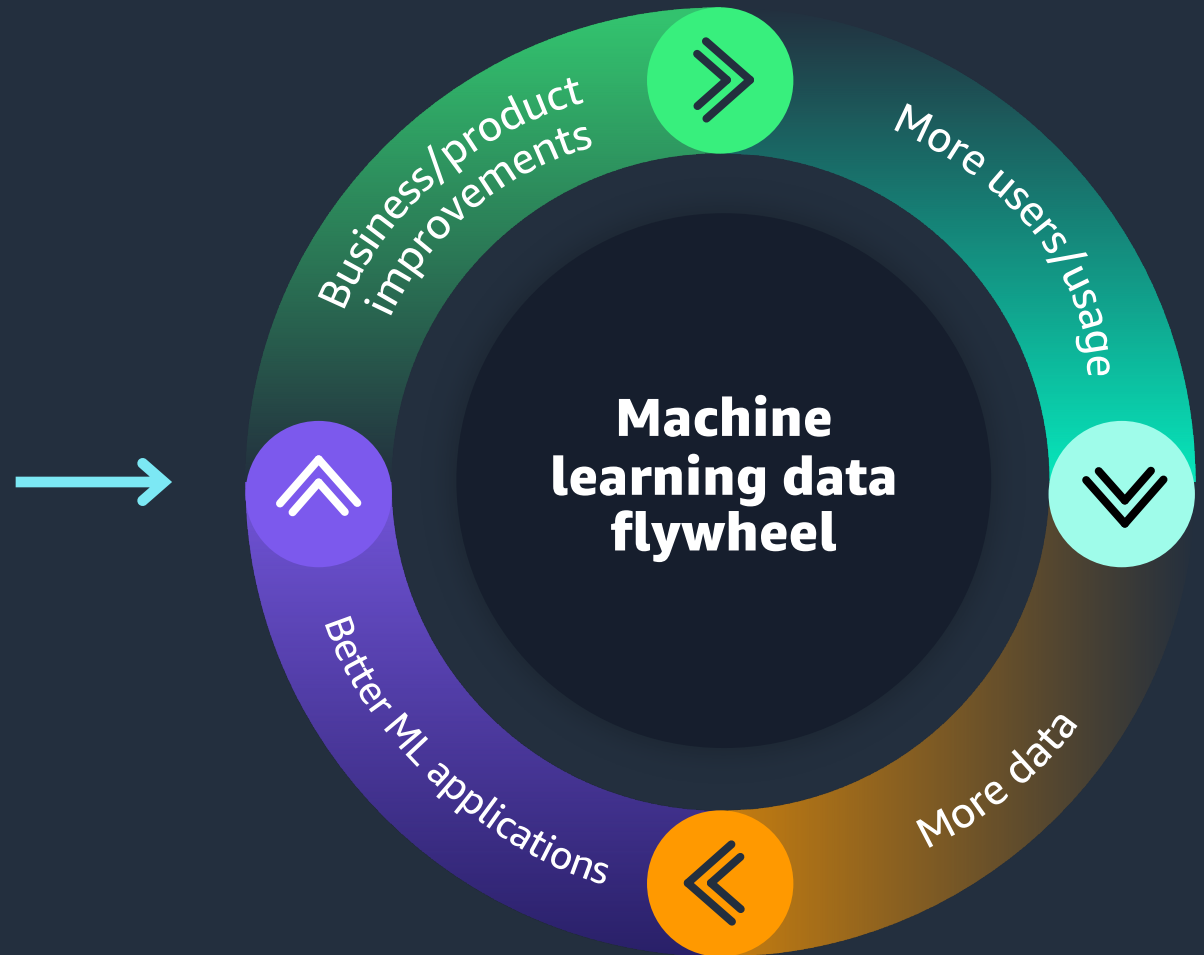


Building Practical ML Applications is hard

The bottleneck

BUILDING BETTER ML APPLICATIONS CAN BE THE BOTTLENECK BECAUSE:

- Line of business teams depend on data science teams to make ML powered decisions
- Data science teams are oversubscribed, while ML needs are only increasing
- High learning curve for technical users to learn how to code
- Friction points in the machine learning journey and prioritizing use cases



Amazon SageMaker Canvas

No-code workspace for
business teams to **build,**
customize, and deploy
ML and Generative AI models



Use ready-to-use models
Pretrained ML models including
Foundation Models



Build custom models
Prepare data, build custom models,
train and deploy models



Collaboration with ML experts
Interoperate with other tools

Ready-to-use ML models powered by **AWS AI Services**

- Foundation Models
- Intelligent document processing
- Natural Language processing
- Computer Vision

The screenshot displays the AWS AI Services console interface. At the top, there is a search bar labeled 'Search use case' and a link 'Can't find the right model? Create a custom model'. Below this, the 'Generative AI-powered foundation models' section is highlighted, featuring a card for 'Generate, extract and summarize content' powered by Amazon Bedrock. The 'Additional ready-to-use models' section follows, with a filter for 'Text', 'Image', and 'Document' data types. This section contains ten model cards arranged in two columns: 'Document queries' (Amazon Textract), 'Identity document analysis' (Amazon Textract), 'Document analysis' (Amazon Textract), 'Expense analysis' (Amazon Textract), 'Sentiment analysis' (Amazon Comprehend), 'Entities extraction' (Amazon Comprehend), 'Language detection' (Amazon Comprehend), 'Personal information detection' (Amazon Comprehend), 'Object detection in images' (Amazon Rekognition), and 'Text detection in images' (Amazon Rekognition). Each card includes a brief description of the model's capabilities and the underlying AWS service it is powered by.

Search use case

Can't find the right model? [Create a custom model](#)

Generative AI-powered foundation models

Our content generation models can help you craft engaging narratives, articles, answer questions, and more, tailored to your needs.

Generate, extract and summarize content New
Powered by Amazon Bedrock and publicly available models

Additional ready-to-use models

Our ready-to-use content extraction models can quickly distill insights from text, image, and document data.

Filter by data type: Text Image Document

- Document queries**
Extract information from structured documents such as paystubs, bank statements, W-2s, and mortgage application forms by asking questions using natural language.
Powered by Amazon Textract
- Identity document analysis**
Extract information from passports, driver licenses, and other identity documentation issued by the US Government.
Powered by Amazon Textract
- Document analysis**
Analyze documents and forms for relationships among detected text.
Powered by Amazon Textract
- Expense analysis**
Extract information from invoices and receipts, such as date, number, item prices, total amount, and payment terms.
Powered by Amazon Textract
- Sentiment analysis**
Detect sentiment in lines of text, which can be positive, negative, neutral, or mixed.
Powered by Amazon Comprehend
- Entities extraction**
Extract entities, which are real-world objects such as people, places, and commercial items, or units such as dates and quantities, from text.
Powered by Amazon Comprehend
- Language detection**
Determine the dominant language in text such as English, French or German.
Powered by Amazon Comprehend
- Personal information detection**
Detect personal information that could be used to identify an individual, such as addresses, bank account numbers, and phone numbers, from text.
Powered by Amazon Comprehend
- Object detection in images**
Detect objects, concepts, scenes, and actions in your images.
Powered by Amazon Rekognition
- Text detection in images**
Detect text in your images.
Powered by Amazon Rekognition



Custom Models - Comprehensive ML capabilities

Prepare data, build custom models, train and deploy models

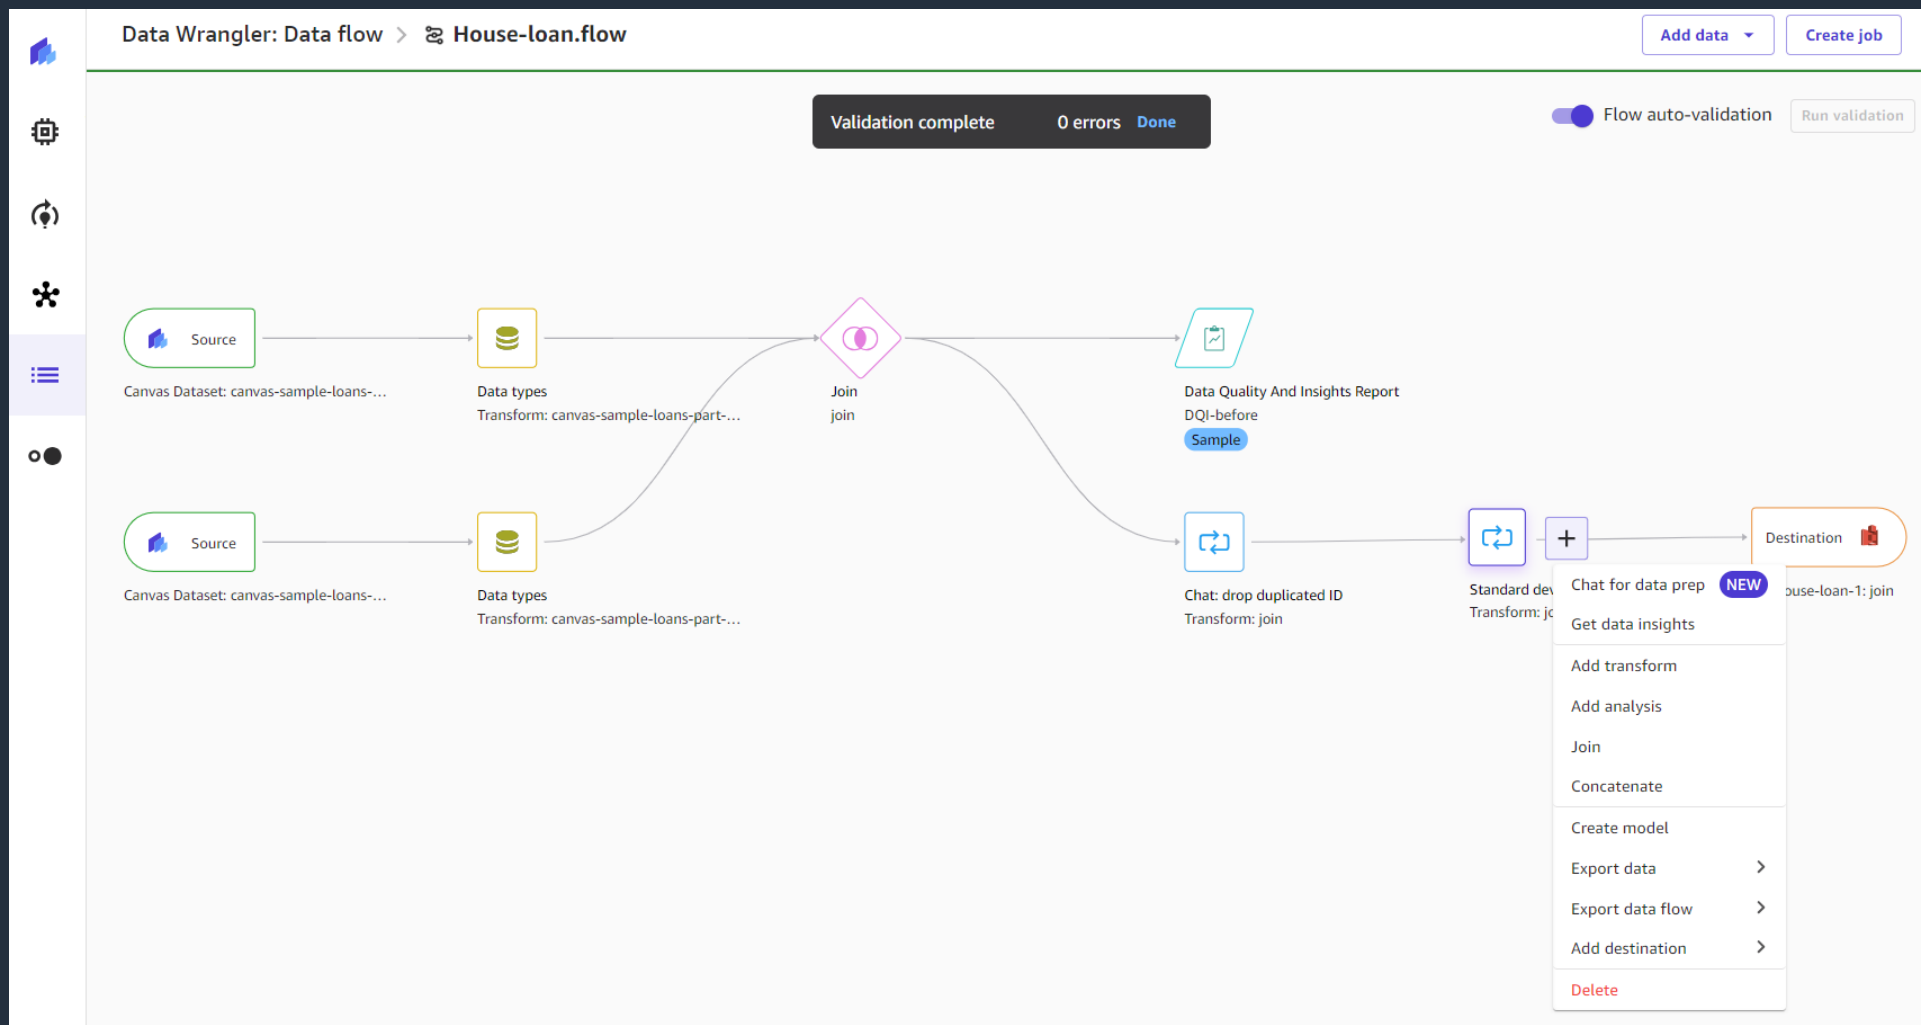
Prepare Data

- 50+ data connectors
- Rich data insights powered by ML
- Built-in visualizations
- 300+ build-in transforms

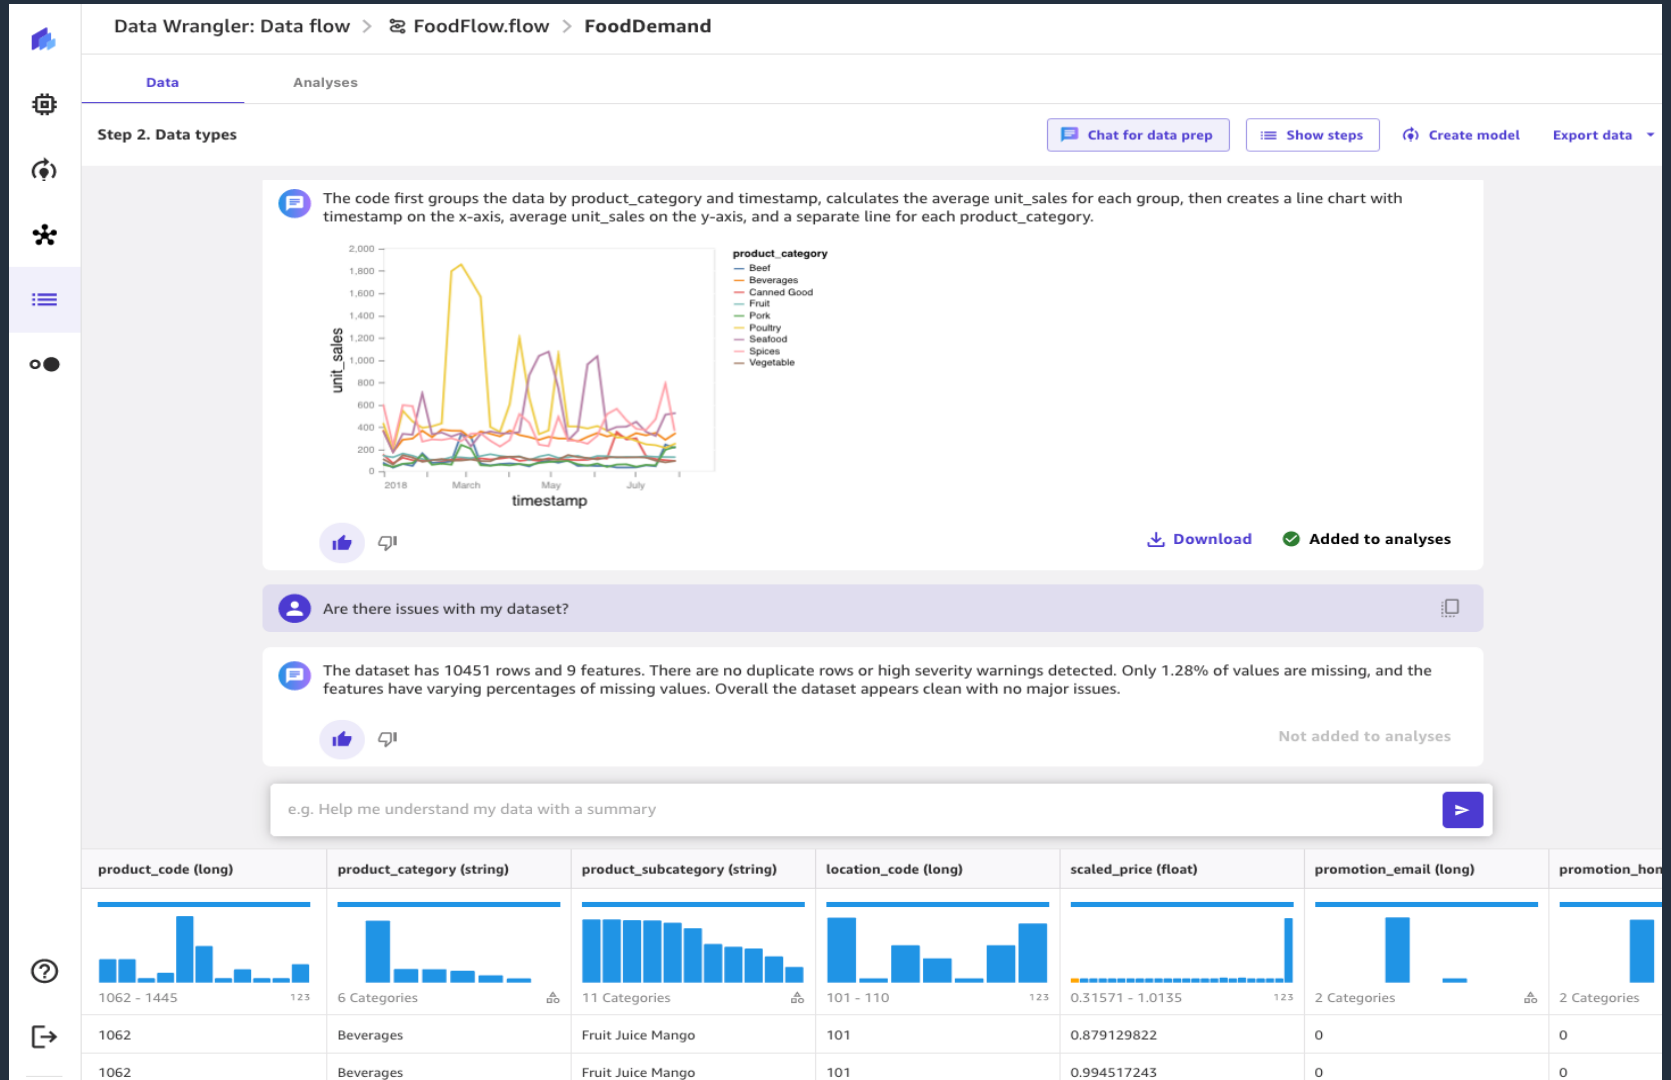
The screenshot displays the AWS SageMaker Canvas interface, which is used for preparing data and building machine learning models. The interface is divided into several sections:

- Data Wrangler: Data flow**: Shows the data source as `canvas-sample-loans-part-1.csv`.
- Data quality and insights report: DQI-before**: Provides a summary of data quality, including target column (`loan_status`), type (`Classification`), dataset, and date (November 20, 2023 at 3:17 PM PST).
- Summary**: Displays the model being built, `Diabetic Readmission model`, Version 1, with a target column of `readmitted` and a 3+ category prediction.
- My models / Lab 1 - Churn / Version 2**: Shows the model being built, `Churn`, with a target column of `Churn` and a 2 category prediction. The model type is `SageMaker Canvas`.
- ChurnData**: Displays a table of data with columns: `Vmail_Plan`, `Vmail_Me...`, `State`, `Phone`, and `Night_Mi`. The table shows 5,000 rows of data.
- Add transform**: Provides options to manage columns, change data type, use custom formula, replace missing values, replace outlier values, manage rows, drop duplicate rows, and drop rows by formula.

Data Preparation Flow



Data prep using natural language



Build and evaluate custom models

- Choose model type
- Analyze model metrics
- Model leaderboard

The screenshot displays the AWS SageMaker console interface. A 'Create new model' dialog is open, showing the 'Analyze' tab for a model named 'Customer Churn Version 2'. The model's status is 'Ready', and its accuracy is 97.303%. The dialog also shows a 'Model leaderboard' table with 10 rows of model performance metrics.

Create new model

My models > Customer Churn > Version 2

Model status

Accuracy 97.303% F1 0.973 Optimization metric

The model predicts the correct Churn 97.303% of the time.

Model leaderboard

	Model name	F1	Accuracy	AUC	Balanced Accuracy	Precision	Recall	Log Loss	Inference latency (s)
1	Night_Calls								
2	Eve_Mins	97.303%	97.303%	0.991	97.303%	97.206%	97.400%	0.200	0.514
3	CustServ_Mins	91.929%	91.808%	0.970	91.810%	90.504%	93.400%	0.251	0.222
4	Eve_Charge	97.114%	97.103%	0.993	97.103%	96.634%	97.600%	0.102	0.119
5	Vmail_Plans	94.466%	94.406%	0.986	94.407%	93.359%	95.600%	0.155	0.111
6	Day_Mins	94.466%	94.406%	0.986	94.407%	93.359%	95.600%	0.155	0.107
7	Day_Calls	94.466%	94.406%	0.986	94.407%	93.359%	95.600%	0.155	0.108
8	Night_Mins	97.114%	97.103%	0.993	97.103%	96.634%	97.600%	0.102	0.121
9	Vmail_Messages	97.114%	97.103%	0.993	97.103%	96.634%	97.600%	0.102	0.121
10	Intl_Calls	94.466%	94.406%	0.986	94.407%	93.359%	95.600%	0.155	0.107

Generate highly accurate predictions

- In-app predictions & what-if analysis
- Automate predictions
- One-click model deployment
- Share predictions to Amazon QuickSight

The screenshot displays the Amazon SageMaker console interface for a model named 'Customer Churn'. The top navigation bar shows 'My models > Customer Churn > Version 2'. The main content area has tabs for 'Select', 'Build', 'Analyze', 'Predict', and 'Deploy', with 'Predict' currently selected. Under the 'Predict' tab, there are options for 'Batch prediction' and 'Single prediction', and a section for 'Predict target values' with a text input field and a 'Modify values to predict Churn in real time.' button.

Overlaid on the console are two panels. The first panel, titled 'Automate batch prediction', contains the text: 'Select a dataset to generate predictions for. Every time the selected dataset is updated (either manually or automatically), a new automatic batch prediction job runs.' Below this is a section for 'Prediction frequency'. The second panel, titled 'Churn Prediction', shows a large 'True.' result. Below this, a legend indicates 'New prediction' (blue square) and 'Average prediction' (grey square). Two horizontal bar charts are shown: one for 'False.' at 29.577% and another for 'True.' at 70.423%.


At the bottom, a 'Details' panel for the deployment 'canvas-customer-churn-prediction-model' is visible. It includes fields for 'Deployment name', 'Created' (10/21/23 09:10 AM), 'Instance type' (ml.m5.xlarge), and 'Deployment URL'. The 'Status' is 'In service' and the 'Instance count' is '1'. The 'Deployment type' is 'Real-time' and the 'Inference response content' is 'predicted_label, probability, probabilities, labels'. A 'Download prediction' button is at the bottom right.

Amazon Bedrock Studio



Bedrock Studio

BUILD GEN AI APPLICATIONS FASTER AND MORE SECURELY



Easy to use playground

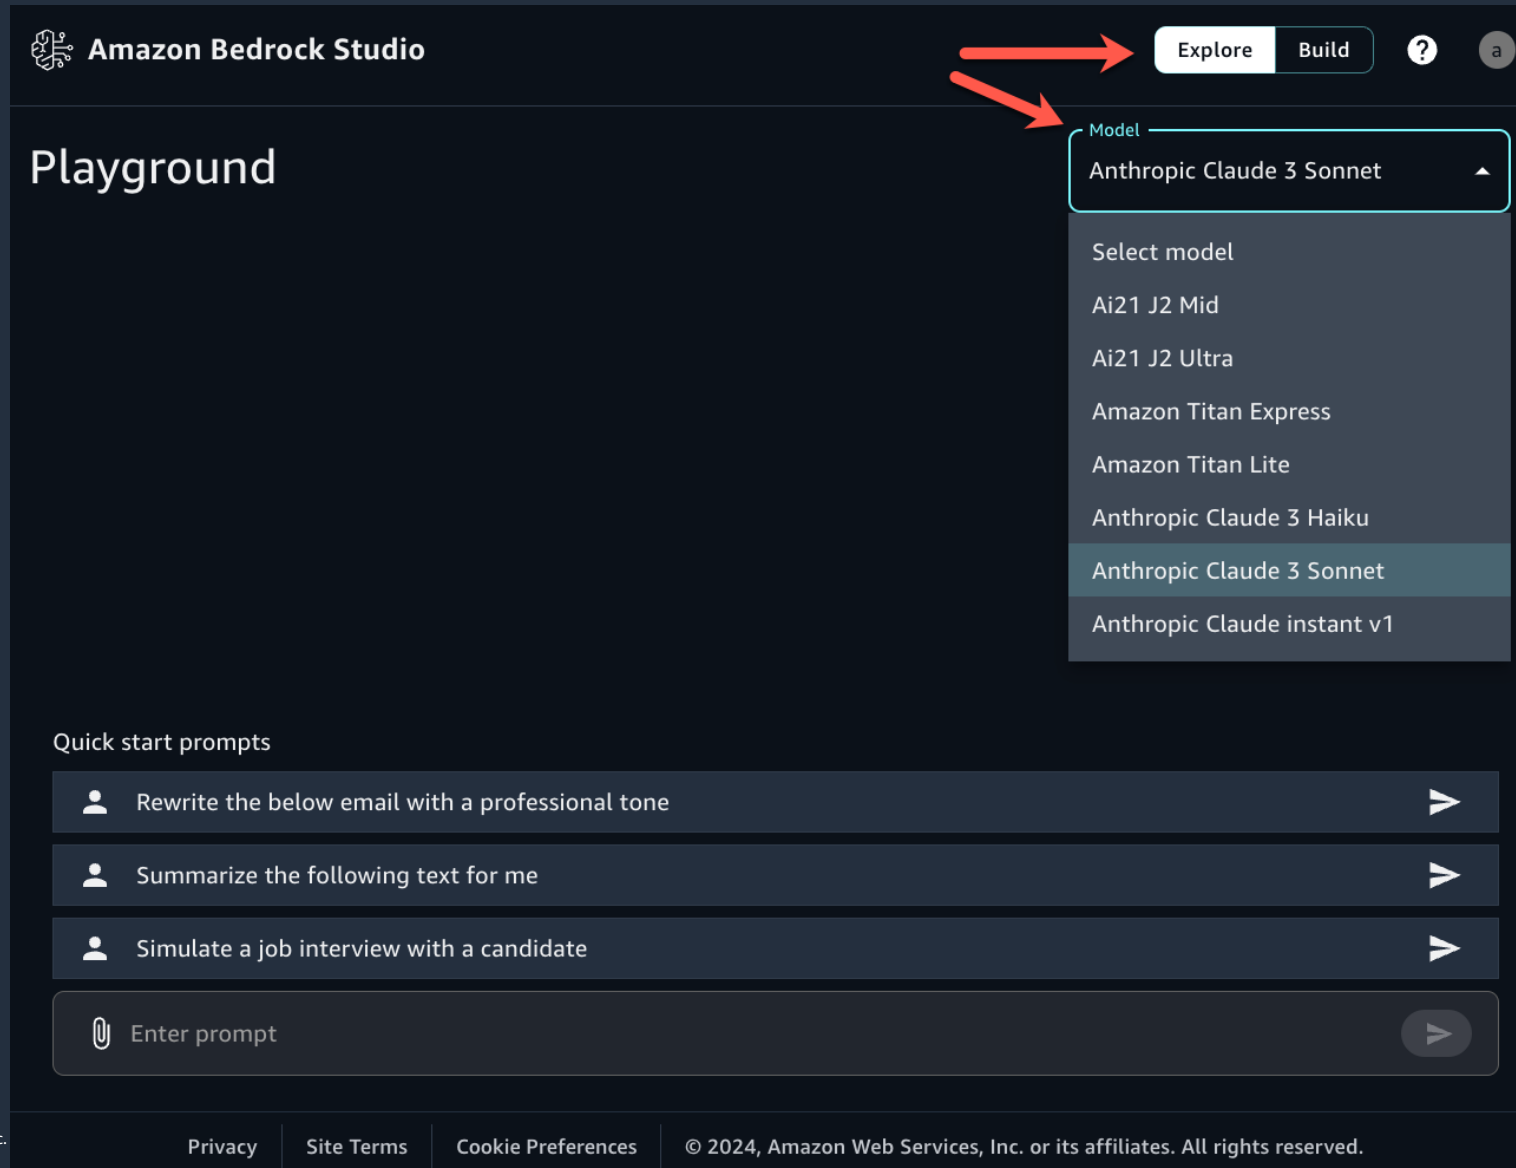


Projects based
collaboration



Easy access with
corporate SSO

Start building Gen AI apps in a playground mode



Amazon Bedrock Studio




Explore Build ? a


Playground

Model

- Anthropic Claude 3 Sonnet
- Select model
- Ai21 J2 Mid
- Ai21 J2 Ultra
- Amazon Titan Express
- Amazon Titan Lite
- Anthropic Claude 3 Haiku
- Anthropic Claude 3 Sonnet
- Anthropic Claude instant v1

Quick start prompts

-  Rewrite the below email with a professional tone
-  Summarize the following text for me
-  Simulate a job interview with a candidate

 Enter prompt

aws

© 2024, Amazon Web Services, Inc.

Privacy Site Terms Cookie Preferences © 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Experiment with model configurations

The screenshot displays the Amazon Bedrock Studio interface. At the top, the header includes the Amazon Bedrock Studio logo, a project dropdown menu set to 'antje-project', and navigation buttons for 'Explore' and 'Build'. A red arrow points to the 'Build' button. Below the header, the main workspace is divided into several sections. On the left, a sidebar shows the project name 'antje-app' with a 'Saved' status. The main area is split into two columns. The left column contains a 'System prompt & examples' section with a text area for the system prompt: 'You are a developer advocate, helping developers build generative AI applications on AWS by answering questions and demonstrating new Amazon Bedrock capabilities.' Below this is an 'Examples' section with a '+ Add examples' button. The right column features a 'Preview' section with a 'Reset' button and a chat window showing the message: 'Hi! I am your developer advocate chat assistant.' At the bottom of the left column, a 'Parameters' section is highlighted with a red rounded rectangle. It contains three sliders: 'Temperature' set to 0.5, 'Top P' set to 0.5, and 'Top K' set to 200. At the bottom right, there is a 'Quick start prompts' section with three buttons: 'Create a list of 10 blog titles with...', 'Summarize the following text...', and 'Write a blog post about...'. Below these is a text input field labeled 'Enter prompt' with a send button.

Amazon Bedrock Studio

Project: antje-project

Explore Build

antje-app

Saved

System prompt & examples

System prompt

You are a developer advocate, helping developers build generative AI applications on AWS by answering questions and demonstrating new Amazon Bedrock capabilities.

Examples

+ Add examples

Parameters

Temperature

0.5

Top P

0.5

Top K

200

Preview

Reset

Hi! I am your developer advocate chat assistant.

Quick start prompts

Create a list of 10 blog titles with...

Summarize the following text...

Write a blog post about...

Enter prompt

Securely bring your own data

The screenshot displays the Amazon Bedrock Studio interface. The top navigation bar includes the Amazon Bedrock Studio logo, a project dropdown menu set to 'antje-project', and buttons for 'Explore' and 'Build'. The left sidebar shows a list of applications, with 'antje-app' selected. The main content area is divided into a 'System prompt' section, a 'Preview' section, and a 'Data' section. The 'Data' section is highlighted with a red box and contains two radio buttons: 'Use single file only' and 'Use Knowledge Base'. The 'Use Knowledge Base' option is selected. Below the radio buttons is a dropdown menu labeled 'Select Knowledge Base' with 'antje-gaia' selected. A red arrow points from the 'Data' section to the 'Create Knowledge Base' dialog box on the right. The dialog box has a title 'Create Knowledge Base' and a close button. It contains a 'Knowledge Base name' field with the placeholder 'Add Knowledge Base name', a 'Knowledge Base description' field with the placeholder 'Add Knowledge Base description here', and an 'Add data sources' section with a 'Click to upload' button and a list of supported file formats: .docx, .doc, .pdf, .txt, .md, .html, .csv, .xls, .xlsx. The dialog also includes a 'Cancel' button and a 'Create' button.

Amazon Bedrock Studio

Project: antje-project

antje-app

System prompt

You are a developer advocate, helping developers build generative AI applications on AWS by answering questions and demonstrating new Amazon Bedrock capabilities.

Examples

+ Add examples

Parameters

Data

☐ Use single file only

☒ Use Knowledge Base

Select Knowledge Base

antje-gaia

+ Create new Knowledge Base

Guardrails

Guardrails are not available while data sources or functions are applied

Functions

UI

Create Knowledge Base

Knowledge Base name

Add Knowledge Base name

Valid characters are a-z, A-Z, 0-9, _ (underscore) and - (hyphen). Up to 100 characters

Knowledge Base description

Add Knowledge Base description here

The description can have between 1 and 150 characters.

Add data sources

Click to upload or drag and drop

.docx, .doc, .pdf, .txt, .md, .html, .csv, .xls, .xlsx (max. 50 MB per file, 50 files)

Embeddings model

Cancel Create

Create Guardrails

The screenshot displays the Amazon Bedrock Studio interface. On the left, the 'antje-app' configuration is visible, with the 'Guardrails' section highlighted by a red rectangle. The 'Guardrails' section shows a dropdown menu with 'None selected' and a '+ Create new guardrail' button. A red arrow points from this section to the 'Create guardrail' modal on the right. The modal contains the following fields and options:

- Guardrail name:** A text input field with the placeholder 'Add guardrail name'. Below it, a note states: 'Provide a name up to 50 characters. Valid characters are a-z, A-Z, 0-9, _ (underscore) and - (hyphen). No spacing allowed.'
- Guardrail description:** A text input field with the placeholder 'Add guardrail description here'. Below it, a note states: 'The description can have up to 100 characters.'
- Content filters:** A section with a title 'Content filters ⓘ' and a description: 'Filter content that violates your usage policies by adjusting the strength of filters that block harmful user prompts and model responses.'
- Enable content filters:** A checkbox that is checked.
- Filters for prompts:** A section with a title 'Filters for prompts ↺ Reset' and a slider for 'Hate' ranging from 'None' to 'High'.
- Buttons:** 'Use advanced filters', 'Cancel', and 'Create'.

Save and share apps

Amazon Bedrock Studio

Project

EduGov Assistant - Project

Explore

Build

?

EduGov Assistant - Project

Delete

Share

This is the project for the EduGov Assistant

Apps (3)

Last updated

Show all

+ Create app

Application 1

Last updated 6 minutes ago

Open

Application 2

Last updated about 2 hours ago

Open

EduGov Assistant

Last updated about 3 hours ago

Open

Components (3)

Last updated

Show all

Type

Show all

+ Create component

servicenow-kb-articles

Knowledge Base . Last updated 3 days ago

Unauthorized-System-Access-Blocker

Guardrail . Last updated 3 days ago

MyTestFunction


Function . Last updated 1 minute ago

Privacy

Site Terms

Cookie Preferences

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.



© 2024, Amazon Web Services, Inc. or its affiliates.

8

Amazon Q



Amazon Q, a generative AI powered assistant from AWS

AMAZON Q DEVELOPER



AMAZON Q BUSINESS

Embedded

Amazon Q
In Connect

Amazon Q
In QuickSight

Amazon Q in
AWS Supply Chain

Benefits

In-built privacy and security

Customizable to your business



Amazon Q: Secure and private by design

YOUR DATA IS YOUR DIFFERENTIATOR



Respects existing governance identities, roles, and permissions,
personalizing interactions accordingly



When you sign up for Amazon Q Developer Pro or Amazon Q Business, **we don't use content to improve underlying models for others**

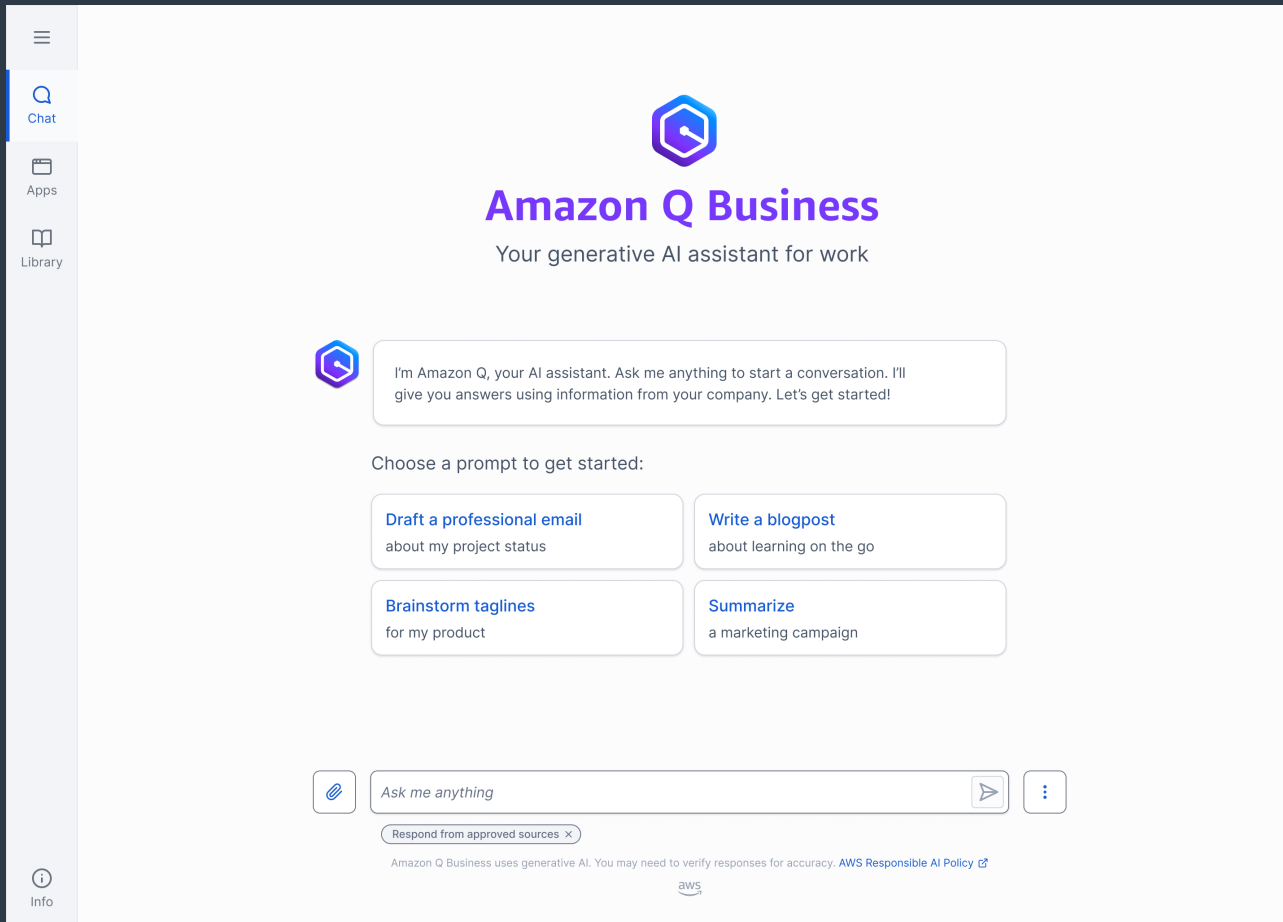
Amazon Q Business

- The generative AI assistant that empowers employees with your organization's knowledge and data



Amazon Q Business Overview

BOOST YOUR WORKFORCE PRODUCTIVITY WITH GENERATIVE AI



Delivers quick, accurate, and relevant answers to your business questions, securely, and privately

Execute actions using out-of-the-box or custom plugins

Respects existing access control based on user permissions

Connects to over 40 popular enterprise applications and document repositories

Enables administrators to easily apply guardrails to customize and control responses

Streamlines daily tasks with user-created lightweight applications



Amazon Q Apps

EASILY BUILD AND SHARE CUSTOMIZED, SECURE, LIGHTWEIGHT APPLICATIONS

The screenshot displays the Amazon Q Apps interface for an application named 'Outreach Assistant'. The interface has a light blue header with a 'Customize' button in the top right. On the left, a sidebar contains icons for 'Chat', 'Apps' (which is highlighted), and 'Library'. The main content area features the app's logo, a title 'Outreach Assistant', and a subtitle 'Generates tailored customer outreach emails'. Below this, there are three text input fields: 'Company' (with placeholder 'Enter company name'), 'Product' (with placeholder 'Enter product name'), and 'Account manager' (with placeholder 'Enter account manger name'). Under the 'Company' field is a file upload section titled 'File upload' with the filename 'Ideal email.docx' and a dashed box for dragging and dropping files, with a note 'Text based files are supported. File size limit is 10MB.' To the right of the file upload is an 'Outlook connector' section titled 'Email message' with a status 'Generating as soon as Company, Product, and Account manager are populated'. At the bottom, there is a 'Jira plugin' section titled 'Related Jira stories' and a 'Run' button with a play icon. A status bar at the very bottom indicates 'Generating as soon as Email message is'.

Easily convert an Amazon Q conversation into a lightweight application or describe what you want to build using natural language

Employees can share their applications with other employees or publish them to the organization's catalog for others to use as is, or modify them to their particular needs

Leverages your enterprise data through Amazon Q 40+ connectors

Inherits Amazon Q security and governance controls, including user authentication and access controls

Creating an app

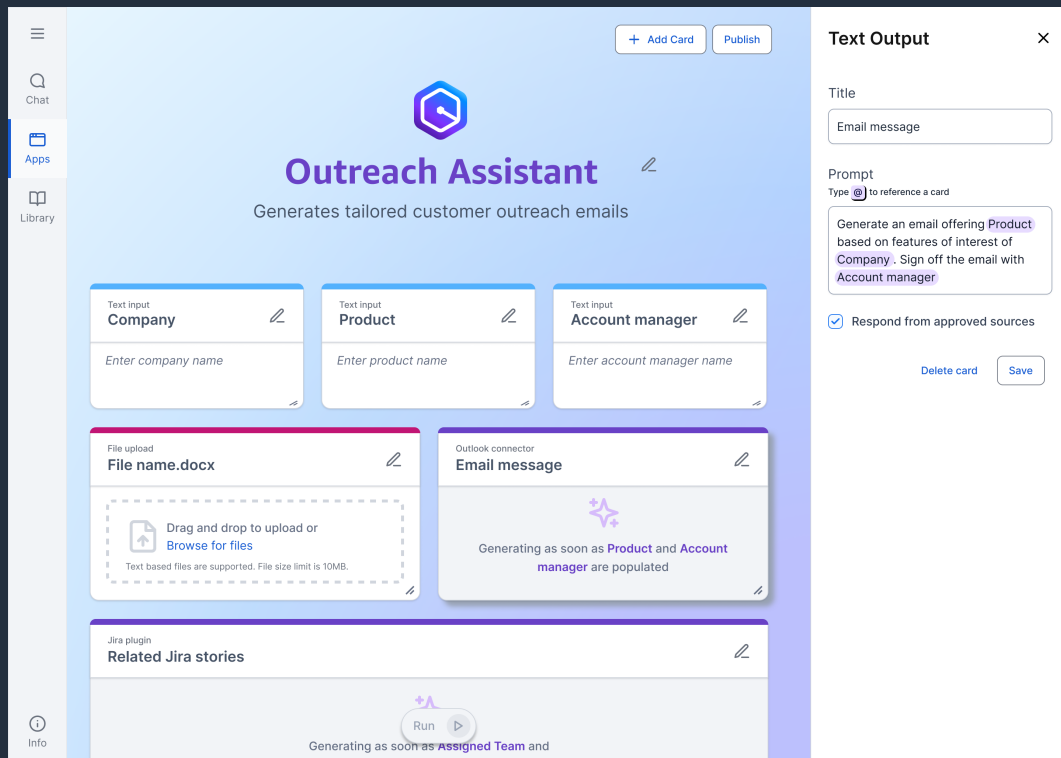
The screenshot shows the Amazon Q Business interface. On the left is a sidebar with 'Chat', 'Apps', and 'Library' options. The main area displays a chat conversation. At the top, there's a 'Create Amazon Q App' button. The chat history includes a message about management and supply chain management, followed by a list of services: Cybersecurity solutions, antivirus software, identity and access management, and cloud computing services. Below this is a 'Sources' dropdown. The next message is from a user asking to generate an email. This is followed by a response from Amazon Q listing key benefits of cloud services: Cost Savings, Reliability, and a prompt for a personalized demo. At the bottom, there's a prompt to 'Try turning this conversation into an Amazon Q App!' with a 'Create Amazon Q App' button. The chat input field at the bottom says 'Ask me anything' and has a 'Respond from approved sources' toggle.

From a conversation

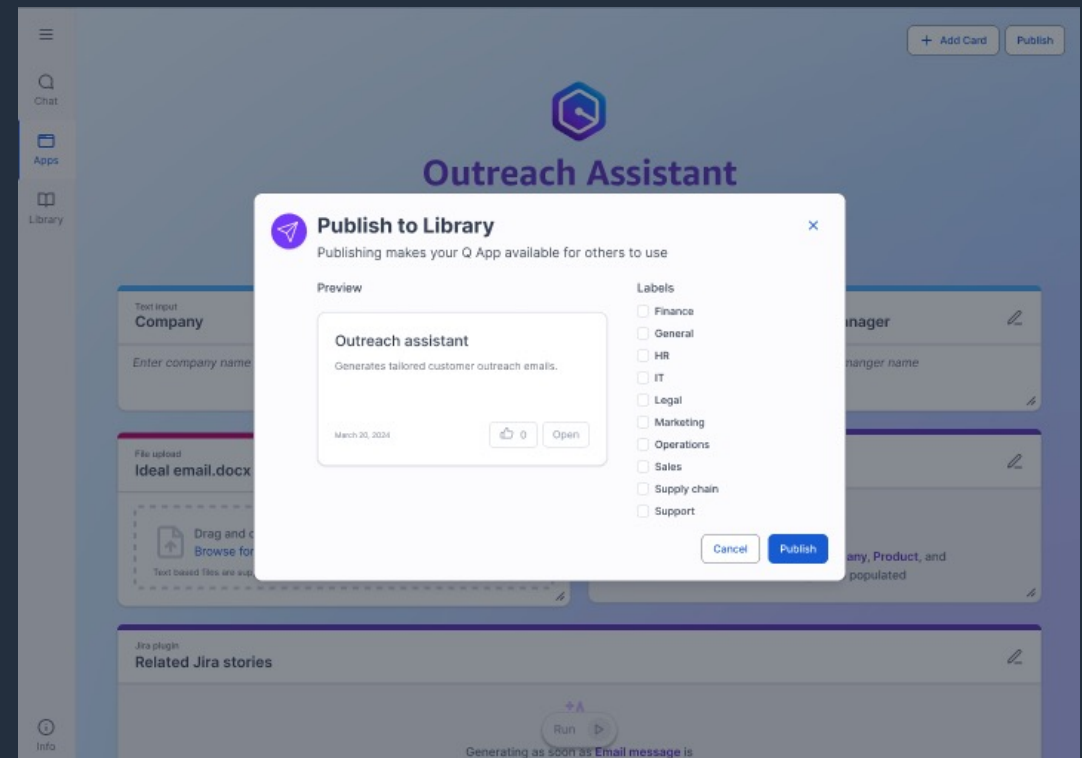
The screenshot shows the Amazon Q App Creator interface. It features the Amazon Q logo and the title 'Amazon Q App Creator' with the subtitle 'Your generative AI productivity app generator'. The main area contains a text input field with a prompt: 'Tired of repetitive tasks? Tell me what you need done and I'll create a custom app tailored for your needs.' Below this is a paragraph explaining how to use the 'sparkle' icon to turn a chat conversation into an Amazon Q App. A text box shows an example prompt: 'You are an expert at writing customer stories for the marketing website. Given a customer name, their business challenge, product used as a solution, and the resulting business impact, you craft a 1000 word customer success story narrative. In this story you cover mission of the customer and structure the story based on past examples.' Below the text box is a 'Character count: 0/10000' indicator and 'Skip this step' and 'Generate' buttons. At the bottom, there's a section 'Try out an example:' with four buttons: 'Content Creator' (Crafts targeted marketing content), 'Interview Question Generator' (Forms questions from a job description), 'Meeting Notes Summarizer' (Summarizes discussion and action items), and 'Grammar Corrector' (Corrects grammar, spelling, and tone).

Describing it in natural language

Using, customizing, and publishing the app



Using or customizing the app



Publishing the app in the library


Discovering apps in the library

Chat

Apps

Library

Info



Amazon Q App Library

Get started with a collection of published Q Apps

All Q Apps

Finance

General

HR

IT

Legal

Marketing

Operations

Sales

Supply chain

Support

Outreach assistant

Generates tailored customer outreach emails

Marketing

Sales

Published on
March 20, 2024

132

Open

Transformer troubleshooter

References the transformer troubleshooter guide

Operations

Support

Published on
March 20, 2024

130

Open

Content creator

Crafts targeted marketing content

Marketing

Published on
March 20, 2024

129

Open

HR Assistant

Advises on HR related topics

HR

Support

Published on
March 20, 2024

125

Open

Newsletter Generator

Generates tailored newsletters

General

Operations

Published on
March 20, 2024

123

Open

Debugging Assistant

Assists and advises in debugging

IT

Published on
March 20, 2024

122

Open

Task Manager

Prioritizes tasks

Incident Processing Planner

Processes incidents

Promo Doc Assistant

Generates content for promotion documents

© 2024, Amazon Web Services, Inc. or its affiliates.

Amazon Q Developer

Reimagine the experience across the entire software development lifecycle



© 2024, Amazon Web Services, Inc. or its affiliates.

Build better and more secure software

AMAZON Q DEVELOPER



Amazon Q Developer assists developers and IT professionals with coding, testing, and upgrading applications, to diagnosing errors, performing security scanning and fixes, and optimizing AWS resources. Amazon Q has advanced, multistep planning and reasoning capabilities that can transform and implement new features generated from developer requests, and help build reliable, secure applications, faster



Amazon Q supports developers across the SDLC

Plan

- Amazon Q Developer in Console (best practices, WAF, EC2 instance optimization, etc.)
- Business-specific application
- Explain code via **conversational coding**

Create

- In-line coding companion in IDE and CLI
- Feature development
- **Conversational coding**

Test & secure

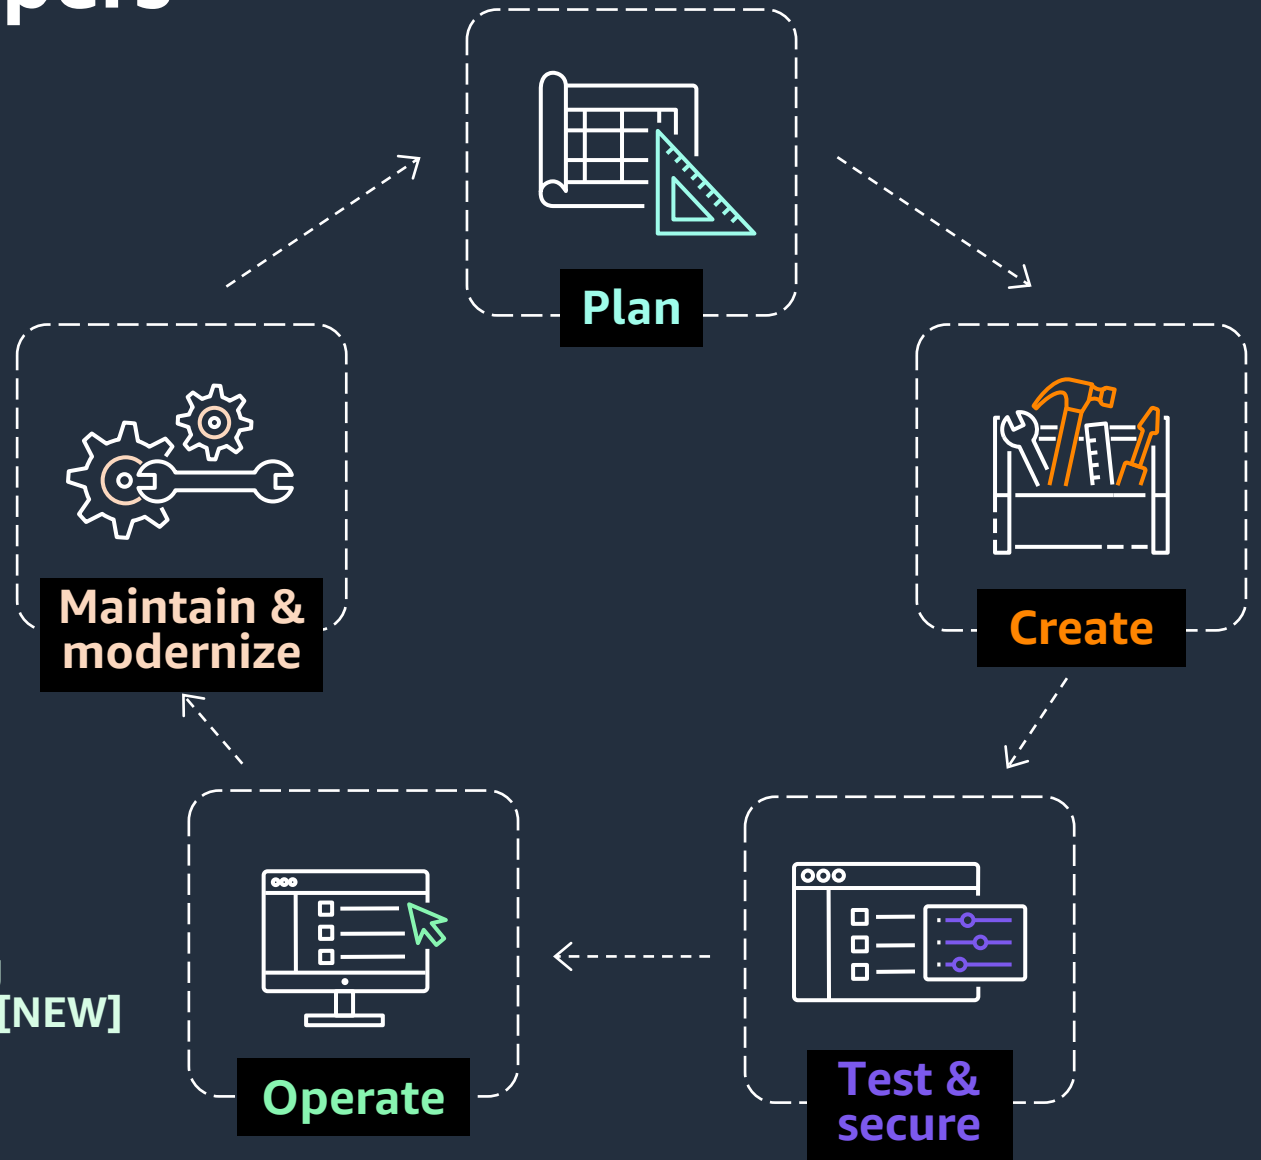
- Unit test generation
- Security scanning & remediation

Operate

- Troubleshoot errors (S3, Lambda, EC2, EKS)
- VPC Reachability Analyzer
- Debug and optimize code via **conversational coding**
- **Helps you optimize your AWS resources and costs [NEW]**

Maintain & modernize

- Update code with Q Code Transformation




Amazon Q in QuickSight

A new generative AI assistant for BI



© 2024, Amazon Web Services, Inc. or its affiliates.



Amazon QuickSight

UNIFIED BI SERVICE
AT HYPERSCALE



Unified BI for all your analytics needs



Pay-as-you-go pricing



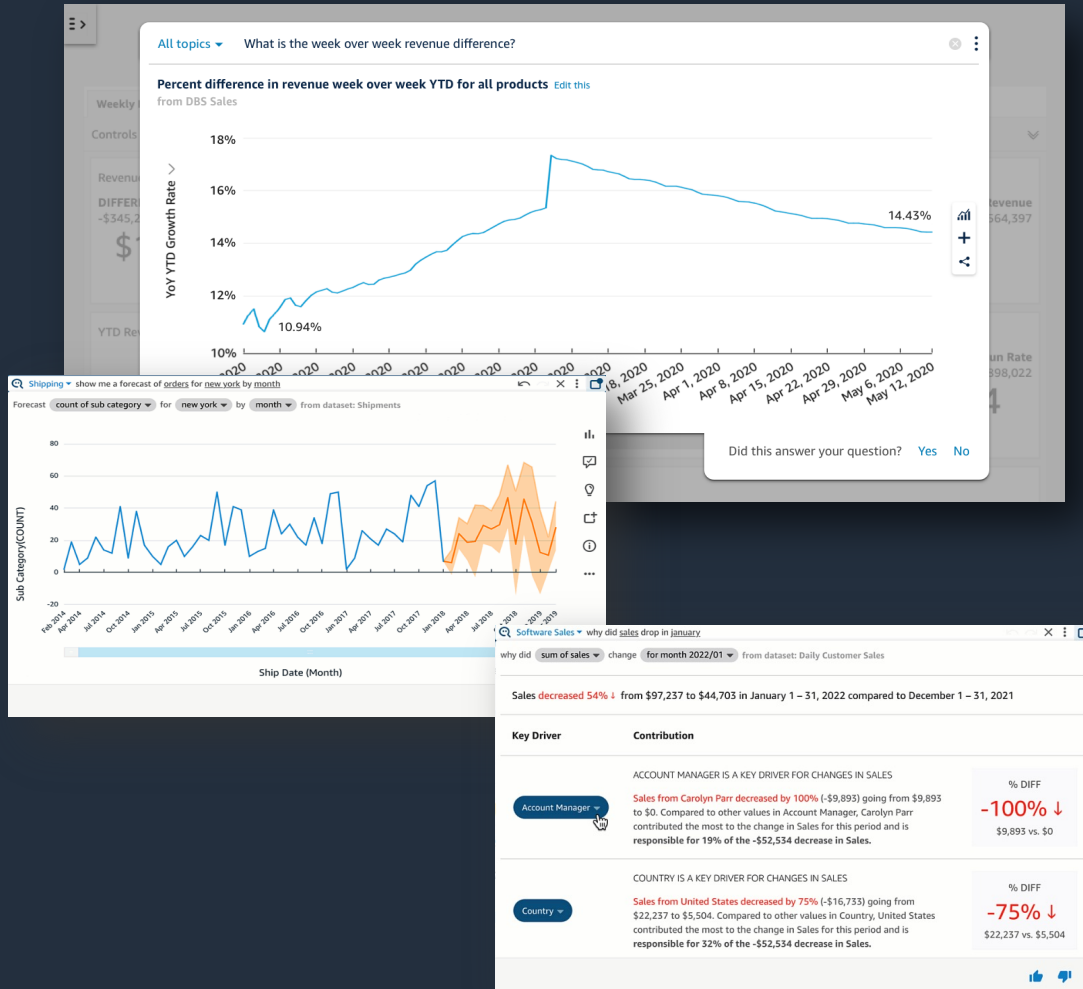
Consistent high performance with auto scaling



Augmented analytics with generative AI capabilities

Foster a data-driven culture with Q

SELF-SERVICE INSIGHTS FOR ANYONE USING ML-POWERED NLQ



- Ask questions in natural language
- ML-models interpret user question and intent to generate visualization
- AI-enhanced automated data preparation accelerates time-to-value

Forecast

- See what's likely to happen
- See future trajectories for up to 3 measures simultaneously

Ask 'Why'?

- Identify key drivers to changes in the data with contribution analysis
- Quantify contribution by each driver



Amazon Connect



Generative AI for immediate CX business value



CUSTOMER

Human-like conversations
Generate comprehensive FAQs



AGENT

Generate real-time answers
Personalize every interaction



SUPERVISOR

Accurate, concise summaries
Auto-fill agent evaluations



ADMINISTRATOR


Auto-generated data mapping



Agent assist



Amazon Q in Connect uses generative AI to assist agents with recommended responses and actions



- Generate real-time responses that agents can deliver directly to customers, along with actions to take, based on context from the live conversation and relevant information from your organizations content (e.g., knowledge articles, FAQs, documents)
- Agents can also search across organization content to receive specific recommendations, like actions to take
- Source knowledge articles and documents are displayed so agents can access more detail if needed

Manager assist



Generative AI-powered post-contact summaries

NEW FEATURE IN AMAZON CONNECT CONTACT LENS




- Concisely summarize important information from customer conversations in a structured, easy-to-read format
- Quickly review customer interactions, understand context, address follow-up tasks, and provide precise feedback to agents without having to read through transcripts or listen to calls

Customer self-service



Generative AI-powered customer self-service experiences

MAKING SELF-SERVICE EXPERIENCES EASIER TO BUILD WITH GENERATIVE AI



- Enhance self-service experiences with conversational question and answer abilities that build on knowledge stores
- Build more effective bots using natural language prompts to describe the use case and automatically generate utterances to improve intent classification accuracy
- Understand more complex customer responses to enable more natural and human-like conversations



AI Use Case Explorer

Easily find the most relevant AI use cases with related content and guidance to make them real

aiexplorer.aws.amazon.com

The screenshot shows the AWS AI Use Case Explorer website. At the top is the AWS logo and a navigation bar with links: About AWS, Contact Us, Support, English, My Account, Sign In, and a 'Create an AWS Account' button. Below this is a secondary navigation bar with links: Products, Solutions, Pricing, Documentation, Learn, Partner Network, AWS Marketplace, Customer Enablement, Events, and Explore More. The main header area is yellow and contains the title 'AI Use Case Explorer' and the subtitle 'Explore the art of the possible in AI across industries and business functions'. Below the subtitle are two buttons: 'Get started with generative AI' and 'Connect with a specialist'. The main content area is white and features the section 'Top generative AI use cases' with three items: 'Improve customer experiences', 'Boost employee productivity', and 'Accelerate process optimization', each with a plus sign to its right. Below this is a section titled 'Explore AI Use Cases' with a 'Clear all filters' link and a 'Use Case Category' dropdown menu. The dropdown menu is open, showing three options: 'Customer experience', 'Employee productivity', and 'Process optimization', each with an unchecked checkbox. To the right of the dropdown menu are three small images: a man and a woman in a meeting, a woman working on a laptop, and a man working on a laptop.



Explore Main features



Explore use cases

Explore curated list of AI use cases for your organization

100+ use cases and sub use cases



Discover success stories

Discover how organizations across the world are using AI to drive business outcomes

400+ customer success stories



Mobilize your team

Follow our expert-curated action plan to realize the power of AI

50+ resources (ebooks, videos, demos)

aiexplorer.aws.amazon.com





Thank you!

Joseph Sadler

jsaws@amazon.com

Please Provide Your Feedback



Step 1: Select Artificial intelligence and machine learning

Step 2: Select Build AI and ML powered applications without machine learning expertise

Learning Day Content

<https://sanfrancisco2024.awslearningday.com/>

