

AWS RE:INVENT

re:Cap

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

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS's current product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS's products or services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied. AWS's responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Join a re:Invent re:Cap community event!



Join one of the free-to-attend **Community re:Cap events** organized by volunteers from **AWS User Groups** around the world.

Dive into technical discussions, gain insights from community experts, and network with your peers!

[Check upcoming events!](#)

Upcoming Community re:Cap events

Date	Country	City	Link to Event
Dec 16	Germany	Darmstadt	AWS User Group Bergstraße & AWS User Group Frankfurt
Dec 16	Germany	Munich	AWS User Group Munich
Dec 16	Canada	Vancouver	AWS User Group Vancouver
Dec 17	Canada	Montreal	AWS User Group Montréal
Dec 17	Italy	Milan	AWS User Group Milan
Dec 17	Japan	Tokyo & Virtual	JAWS-UG IoT
Dec 18	Germany	Berlin	AWS UG Berlin & AWS Women's UG Berlin
Dec 19	Korea	Seoul	AWS User Group Korea (KRUG)
Dec 19	Taiwan	Taipei	AWS User Group Taiwan
Dec 27	Bangladesh	Virtual	AWS User Group Bangladesh

Do more with AWS Builder ID

- Connect with fellow builders
- Get advice on technical challenges
- Access 600+ free courses
- Try Amazon Q Developer
- Be first to explore new AWS offerings

Start by claiming your
community alias



What is AWS re:Invent?

AWS re:Invent is a learning conference hosted by AWS for the global cloud computing community.

The in-person event features keynote announcements, training and certification opportunities, access to more than 2,000 technical sessions, the Expo, after-hours events, and so much more.



AWS re:Invent 2024

60,000

attendees

400,000

people watching
keynotes online

1,900

sessions

3,500

speakers

546

updates and announcements
(*before and during re:Invent*)



Housekeeping and Deck Legend

Generally Available (GA)

GA

Public Preview (PP)

PP

Preview (P)

P

Coming Soon (CS)

CS

Each announcement contains a date when it happened, announcement status, information is it a pre-re:Invent or re:Invent release, is it Open Source related, and regional availability.

Color coding: **pre-re:Invent**, **re:Invent**, **Service Name**, **Open Source**

AWS re:Invent 2024

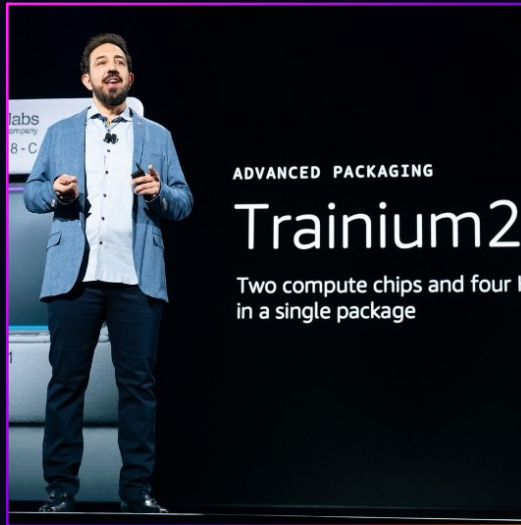
Dec 1st – Dec 6th

What is **pre-re:Invent**?

It's an *unofficial* name of the period before the event itself, during which we observe an increased number of announcements coming from AWS Service teams.

For AWS re:Invent 2024 re:Cap,
we decided to take the date of **October 30th**
as a starting date for this period.

AWS re:Invent 2024 Keynotes



Peter DeSantis

**SVP, UTILITY COMPUTING
AWS**

[Link to video](#)



Matt Garman

**CEO
AWS**

[Link to video](#)



Dr. Swami Sivasubramanian

**VP, AI AND DATA
AWS**

[Link to video](#)

AWS re:Invent 2024 Keynotes



**Dr. Ruba
Borno**

**VP, GLOBAL SPECIALISTS
AND PARTNERS**

AWS

[Link to video](#)



**Dr. Werner
Vogels**

VP AND CTO

AMAZON

[Link to video](#)

AWS re:Invent 2024: Launch Summary

AWS re:Invent 2024 Launch Summary (1/5)

Analytics

[Announcing scenarios analysis capability of Amazon Q in QuickSight \(preview\)](#)

[Announcing Amazon S3 Metadata \(Preview\) – Easiest and fastest way to manage your metadata](#)

[Amazon DynamoDB zero-ETL integration with Amazon SageMaker Lakehouse](#)

[AWS expands data connectivity for Amazon SageMaker Lakehouse and AWS Glue](#)

[Amazon S3 Access Grants now integrate with AWS Glue](#)

[Amazon SageMaker Lakehouse integrated access controls now available in Amazon Athena federated queries](#)

[Introducing AWS Glue 5.0](#)

[Introducing the next generation of Amazon SageMaker](#)

[AWS announces Amazon SageMaker Lakehouse](#)

[Amazon SageMaker Lakehouse and Amazon Redshift support for zero-ETL integrations from eight applications](#)

[AWS Glue Data catalog now automates generating statistics for new tables](#)

Analytics

[Introducing Amazon SageMaker Data and AI Governance](#)

[Announcing the preview of Amazon SageMaker Unified Studio](#)

[Data Lineage is now generally available in Amazon DataZone and next generation of Amazon SageMaker](#)

[Amazon Q in QuickSight unifies insights from structured and unstructured data](#)

[Amazon Q Business now provides insights from your databases and data warehouses \(preview\)](#)

[Amazon OpenSearch Service zero-ETL integration with Amazon Security Lake](#)

[AWS Clean Rooms now supports multiple clouds and data sources](#)

Application Integration

[Amazon SageMaker Lakehouse and Amazon Redshift support for zero-ETL integrations from eight applications](#)

[Amazon Bedrock now supports multi-agent collaboration](#)

[Amazon EventBridge and AWS Step Functions announce integration with private APIs](#)

Business Applications

[Amazon Q Business now provides insights from your databases and data warehouses \(preview\)](#)

[Amazon Connect Contact Lens now supports external voice](#)

[Amazon Connect now supports external voice transfers](#)

[Amazon Connect Contact Lens now automatically categorizes your contacts using generative AI](#)

[Amazon Connect launches AI guardrails for Amazon Q in Connect](#)

[Amazon Connect launches new intraday forecast dashboards](#)

[Amazon Connect launches AI assistant for customer segments and trigger-based campaigns](#)

[Amazon Connect now supports WhatsApp Business messaging](#)

[Amazon Connect Contact Lens now automates agent performance evaluations using generative AI](#)

[Amazon Connect launches generative AI-powered self-service with Amazon Q in Connect](#)

[AWS announces Salesforce Contact Center with Amazon Connect \(Preview\)](#)

AWS re:Invent 2024 Launch Summary (2/5)

Compute

[Amazon EC2 Trn2 instances are generally available](#)

[Announcing Amazon Elastic VMware Service \(Preview\)](#)

[Deploy GROW with SAP on AWS from AWS Marketplace](#)

[Amazon EC2 P5en instances, optimized for generative AI and HPC, are generally available](#)

[Announcing Amazon EC2 I8g instances](#)

[Amazon Web Services announces declarative policies](#)

[Amazon CloudWatch Container Insights launches enhanced observability for Amazon ECS](#)

[Introducing Amazon EC2 next generation high density Storage Optimized I7ie instances](#)

[AWS Marketplace now offers EC2 Image Builder components from independent software vendors](#)

[Amazon EC2 introduces Allowed AMIs to enhance AMI governance](#)

[Announcing Amazon EKS Auto Mode](#)

[Announcing Amazon EKS Hybrid Nodes](#)

[AWS simplifies the use of third-party block storage arrays with AWS Outposts](#)

Containers

[VPC Lattice now includes TCP support with VPC Resources](#)

[Announcing Amazon EKS Auto Mode](#)

[Announcing Amazon EKS Hybrid Nodes](#)

Customer Engagement

[Introducing the AWS Digital Sovereignty Competency](#)

[Amazon Connect now provides the ability to record audio during IVR and other automated interactions](#)

[Amazon Connect Contact Lens now supports external voice](#)

[Amazon Connect now supports external voice transfers](#)

[Amazon Connect launches AI guardrails for Amazon Q in Connect](#)

[Amazon Connect Contact Lens launches built-in dashboards to analyze conversational AI bot performance](#)

[Amazon Connect launches simplified conversational AI bot creation](#)

[Amazon Connect now supports WhatsApp Business messaging](#)

[Amazon Connect launches generative AI-powered self-service with Amazon Q in Connect](#)

[Amazon Connect now makes it easier to collect sensitive customer data within chats](#)

Databases

[Announcing Amazon Aurora DSQL \(Preview\)](#)

[Amazon DynamoDB global tables previews multi-Region strong consistency](#)

[AWS DMS Schema Conversion now uses generative AI](#)

[Announcing the general availability of Amazon MemoryDB Multi-Region](#)

[Oracle Database@AWS is now in limited preview](#)

[Amazon Bedrock Knowledge Bases now supports GraphRAG \(preview\)](#)

[Amazon DynamoDB zero-ETL integration with Amazon SageMaker Lakehouse](#)

AWS re:Invent 2024 Launch Summary (3/5)

Databases *(continued)*

[Amazon Q Business now provides insights from your databases and data warehouses \(preview\)](#)

[Announcing Amazon EC2 I8g instances](#)

[AWS DMS Schema Conversion now uses generative AI](#)

Developer Tools

[Amazon Q Developer adds operational investigation capability \(Preview\)](#)

[Amazon Q Developer now provides transformation capabilities for .NET porting \(Preview\)](#)

[Amazon Q Developer can now automate code reviews](#)

[Announcing GitLab Duo with Amazon Q \(Preview\)](#)

[Amazon Q Developer can now generate documentation within your source code](#)

[Amazon Q Developer transformation capabilities for mainframe modernization are now available \(Preview\)](#)

Front End Web & Mobile

[Storage Browser for Amazon S3 is now generally available](#)

[Amazon Connect now supports WhatsApp Business messaging](#)

Machine Learning

[Amazon Bedrock Knowledge Bases now supports structured data retrieval](#)

[Task governance is now generally available for Amazon SageMaker HyperPod](#)

[AWS announces Amazon SageMaker Partner AI Apps](#)

[Amazon Bedrock Guardrails supports multimodal toxicity detection for image content \(Preview\)](#)

[Announcing new AWS AI Service Cards to advance responsible generative AI](#)

[Amazon SageMaker HyperPod now provides flexible training plans](#)

[Amazon Bedrock Marketplace brings over 100 models to Amazon Bedrock](#)

[Amazon Bedrock announces preview of prompt caching](#)

[Amazon Q Developer can now guide SageMaker Canvas users through ML development](#)

[Amazon Bedrock Data Automation now available in preview](#)

[Amazon Bedrock Knowledge Bases now supports GraphRAG \(preview\)](#)

[Announcing Amazon SageMaker HyperPod recipes](#)

[Amazon Bedrock Knowledge Bases now processes multimodal data](#)

[Announcing GenAI Index in Amazon Kendra](#)

[Amazon Q Developer adds operational investigation capability \(Preview\)](#)

[Amazon Q Developer can now automate code reviews](#)

[Amazon Bedrock Model Distillation is now available in preview](#)

[Amazon Q Business introduces over 50 actions for popular business applications and platforms](#)

[Amazon SageMaker Lakehouse and Amazon Redshift support for zero-ETL integrations from eight applications](#)

[Amazon Bedrock now supports multi-agent collaboration](#)

[Amazon Bedrock Guardrails now supports Automated Reasoning checks \(Preview\)](#)

AWS re:Invent 2024 Launch Summary (4/5)

Machine Learning *(continued)*

[Introducing Amazon SageMaker Data and AI Governance](#)

[Data Lineage is now generally available in Amazon DataZone and next generation of Amazon SageMaker](#)

[Amazon Q in QuickSight unifies insights from structured and unstructured data](#)

[Amazon Q Developer can now generate documentation within your source code](#)

[Announcing Amazon Bedrock IDE in preview as part of Amazon SageMaker Unified Studio](#)

[Announcing Amazon Nova foundation models available today in Amazon Bedrock](#)

[Amazon Q Business now provides insights from your databases and data warehouses \(preview\)](#)

[Amazon Q Developer transformation capabilities for mainframe modernization are now available \(Preview\)](#)

[Amazon EC2 P5en instances, optimized for generative AI and HPC, are generally available](#)

[Introducing latency-optimized inference for foundation models in Amazon Bedrock](#)

[Amazon Bedrock Knowledge Bases now supports RAG evaluation \(Preview\)](#)

[Amazon Bedrock Model Evaluation now includes LLM-as-a-judge \(Preview\)](#)

[Amazon Bedrock Knowledge Bases now provides auto-generated query filters for improved retrieval](#)

[Amazon Q Business adds support to extract insights from visual elements within documents](#)

[Amazon Bedrock Knowledge Bases now supports streaming responses](#)

[Amazon Bedrock now supports Rerank API to improve accuracy of RAG applications](#)

[Amazon Bedrock Knowledge Bases now supports custom connectors and ingestion of streaming data](#)

[PartyRock improves app discovery and announces upcoming free daily use](#)

[Announcing Amazon Q Developer transformation capabilities for VMware \(Preview\)](#)

[The Amazon Q index enhances software vendors' AI experiences](#)

Customer Engagement

[Amazon Web Services announces declarative policies](#)

[Amazon CloudWatch and Amazon OpenSearch Service launch an integrated analytics experience](#)

[Amazon CloudWatch adds network performance monitoring for AWS workloads using flow monitors](#)

[Amazon CloudWatch Container Insights launches enhanced observability for Amazon ECS](#)

[AWS Control Tower launches managed controls using declarative policies](#)

[AWS announces Amazon CloudWatch Database Insights](#)

Migration and Transfer

[Announcing AWS Transfer Family web apps](#)

[Announcing Amazon Q Developer transformation capabilities for VMware \(Preview\)](#)

Networking & Content Delivery

[AWS announces access to VPC resources over AWS PrivateLink](#)

AWS re:Invent 2024 Launch Summary (5/5)

Nova

(Generative AI)

[Announcing Amazon Nova foundation models available today in Amazon Bedrock](#)

Partner

[Buy with AWS accelerates solution discovery and procurement on AWS Partner websites](#)

[Respond and recovery more quickly with AWS Security Incident Response Partners](#)

[Start collaborating on multi-partner opportunities with Partner Connections \(Preview\)](#)

[Introducing the AWS Digital Sovereignty Competency](#)

[New streamlined deployment experience for Databricks on AWS](#)

[Introducing the AWS Consumer Goods Competency](#)

[AWS Marketplace now offers EC2 Image Builder components from independent software vendors](#)

Security, Identity & Compliance

[AWS announces AWS Security Incident Response for general availability](#)

[Amazon Bedrock Guardrails supports multimodal toxicity detection for image content \(Preview\)](#)

[Introducing the Amazon Security Lake Ready Specialization](#)

[Respond and recovery more quickly with AWS Security Incident Response Partners](#)

[Introducing the AWS Digital Sovereignty Competency](#)

[AWS Security Competency Update: New AI Security Category](#)

[Amazon Bedrock Guardrails now supports Automated Reasoning checks \(Preview\)](#)

[Amazon Web Services announces declarative policies](#)

[Amazon OpenSearch Service zero-ETL integration with Amazon Security Lake](#)

[AWS Verified Access now supports secure access to resources over non-HTTP\(S\) protocols \(Preview\)](#)

[Amazon GuardDuty introduces GuardDuty Extended Threat Detection](#)

[AWS announces AWS Security Incident Response for general availability](#)

Serverless

[Amazon OpenSearch Service zero-ETL integration with Amazon Security Lake](#)

Storage

[Amazon S3 Access Grants now integrate with AWS Glue](#)

[Announcing Amazon S3 Metadata \(Preview\) – Easiest and fastest way to manage your metadata](#)

[Announcing Amazon S3 Tables – Fully managed Apache Iceberg tables optimized for analytics workloads](#)

[Announcing Amazon EC2 I8g instances](#)

[Amazon S3 adds new default data integrity protections](#)

[Storage Browser for Amazon S3 is now generally available](#)

[Introducing Amazon EC2 next generation high density Storage Optimized I7ie instances](#)

[Announcing Amazon FSx Intelligent-Tiering, a new storage class for FSx](#)

[Announcing AWS Transfer Family web apps](#)

[Amazon S3 launches storage classes for AWS Dedicated Local Zones](#)

Ready?

Fasten your seatbelts and let's re:Cap!

Visual update for AWS Management Console

BASED ON CLOUDSCAPE DESIGN SYSTEM (CLOUDSCAPE.DESIGN)

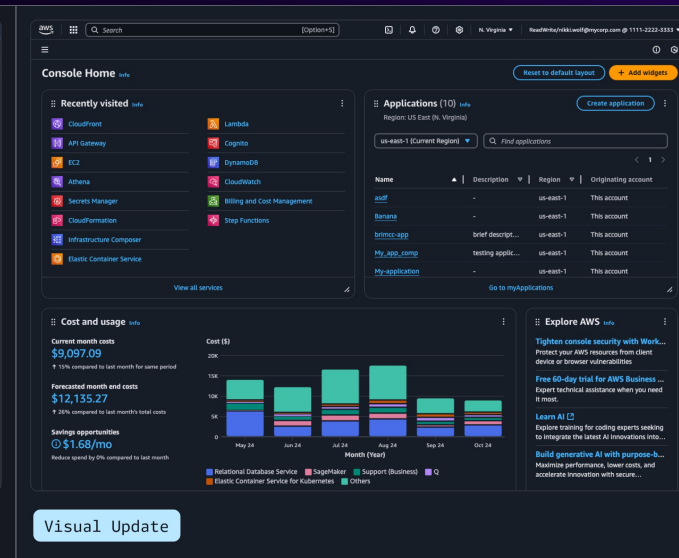
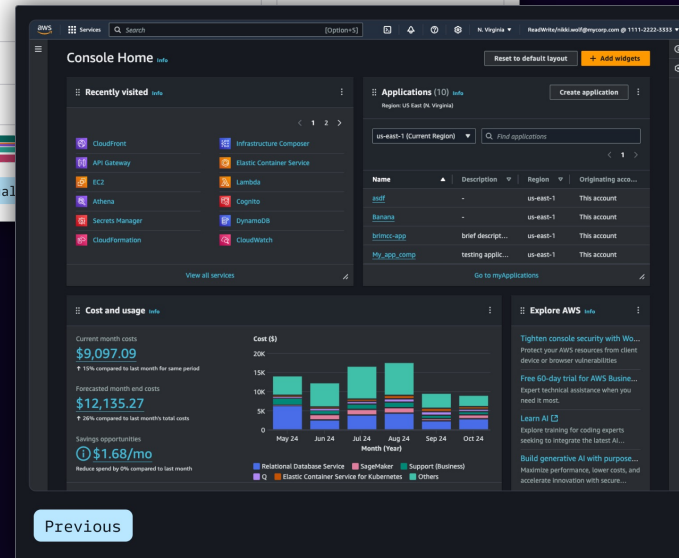
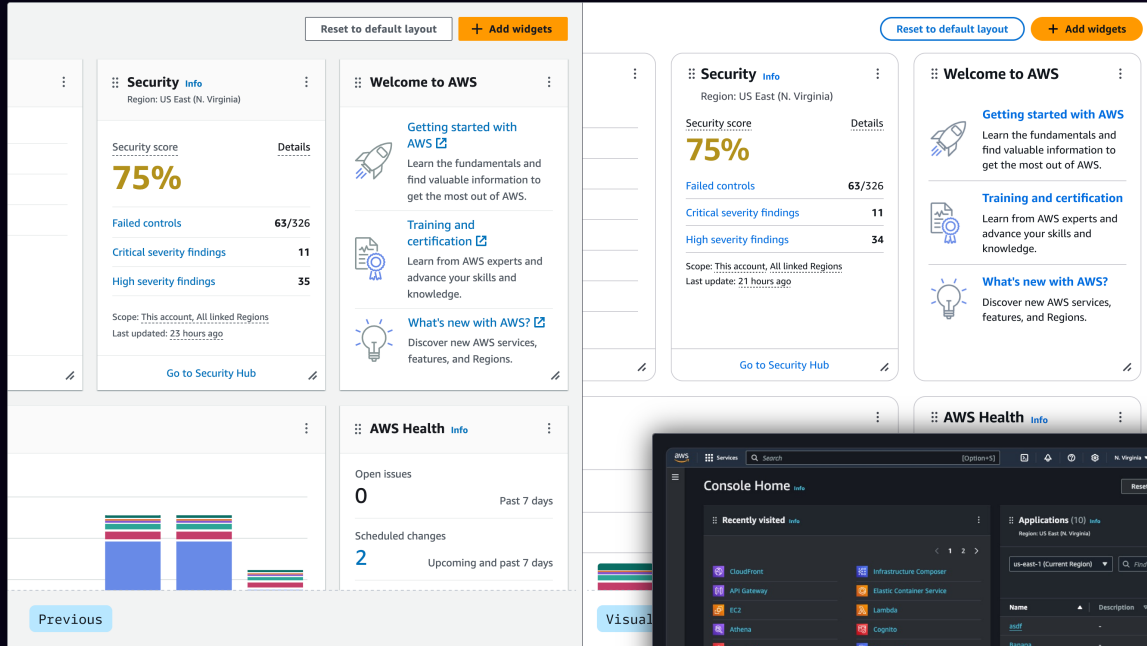


PP

Nov, 21st

pre-re:Invent

All Regions



[AWS News Blog Post](#)



Updates for PartyRock

NEW CAPABILITIES AND FREE DAILY USAGE

PartyRock

An Amazon Bedrock Playground

Everyone can
build AI apps
with
PartyRock

Your App. Your Rules.

Powered by Amazon Bedrock



Billing and Support

How much does PartyRock cost?

For a limited time, AWS offers new PartyRock users a free trial without the need to provide a credit card or sign up for an AWS account, so that you can begin learning fundamental skills without the worry of incurring costs. Starting 2025, a new and improved daily free usage model will replace the current free trial grant.

GA

Dec, 1st

re:Invent

All Regions

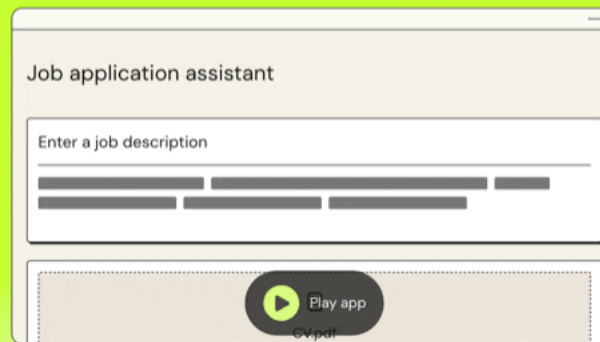
PartyRock

Log in

Get started for free

Get more
done with
generative AI

Start for free



[AWS News Blog Post](#)



AWS Education Equity Initiative

Dec, 4th
re:Invent

EMPOWERING ORGANIZATIONS TO CREATE LEARNING SOLUTIONS FOR UNDERSERVED LEARNERS



GA



C O
D E



[AWS News Blog Post](#)



Partners

7 announcements (0 *pre-re:Invent*)



Buy with AWS

ACCELERATE SOLUTION DISCOVERY AND PROCUREMENT VIA AWS MARKETPLACE AND PARTNERS



GA

Dec, 4th

re:Invent

WIZ Platform Solutions Pricing Resources Customers Company Sign in Get a demo

Partners

aws

Connect Wiz with Amazon Web Services (AWS)

Agentless full stack coverage of your AWS workloads in minutes.

available in
aws marketplace

Buy with AWS Try with AWS

Purchase this listing from Wiz in AWS Marketplace using your AWS account. In AWS Marketplace, you can quickly launch pre-configured software with just a few clicks. AWS handles billing and payments, and charges appear on your AWS bill.



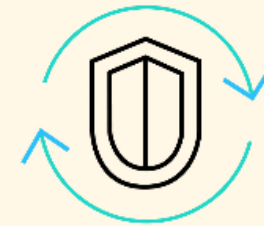
[AWS News Blog Post](#)



New capabilities available for AWS Partners

SOVEREIGNTY, CONSUMER GOODS, AI SECURITY, AMAZON SECURITY LAKE READY

- **New** Competencies:
 - Consumer Goods
 - Digital Sovereignty
- **New** Categories:
 - AI Security (*Security* competency)
- **New** Specializations:
 - Amazon Security Lake Ready
 - AWS Security Incident Response



AWS Security Incident Response

What is AWS Security Incident Response?

AWS Security Incident Response helps you respond when it matters most. The service combines the power of automated monitoring and investigation, accelerated communication and coordination, and direct 24/7 access to the AWS Customer Incident Response Team (CIRT) to quickly prepare for, respond to, and recover from security events.

GA

Dec, 1st

Dec, 4th

re:Invent

Partner Connections

START COLLABORATING ON MULTI-PARTNER OPPORTUNITIES



PP

Dec, 4th

re:Invent

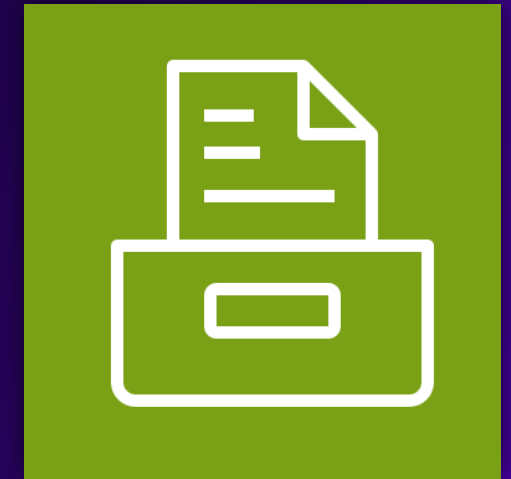


[AWS Partner Network \(APN\)
Blog Post](#)



Storage, Migration and Transfer Services

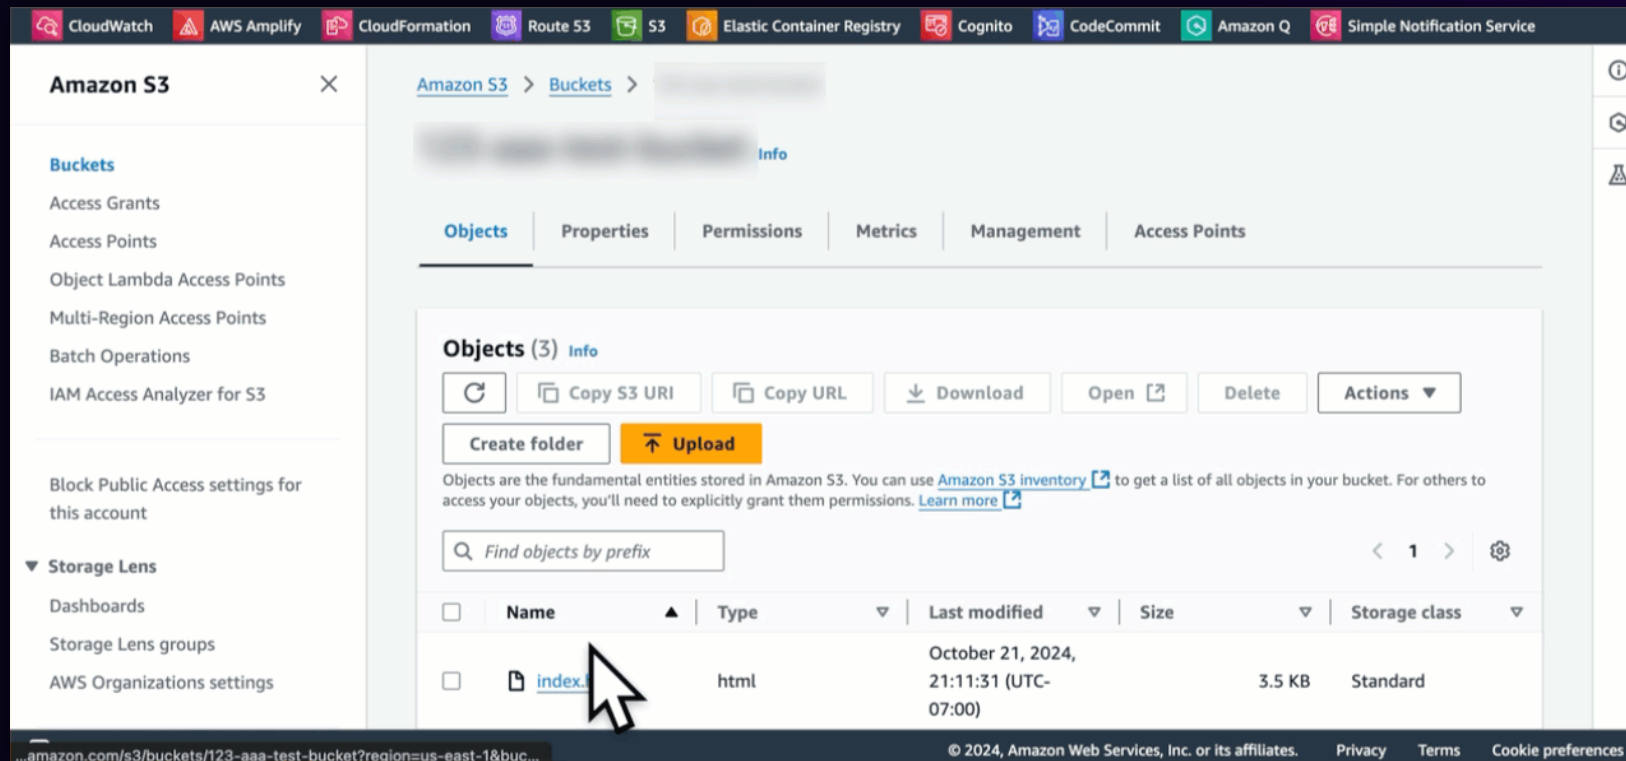
19 announcements (11 *pre-re:Invent*)



Enhance Amazon S3 Static Website Hosting

WITH AWS AMPLIFY HOSTING

From now it is even easier to deploy static websites with content stored in your S3 buckets and serve over a content delivery network (CDN) with just a few clicks powered by *AWS Amplify Hosting*.



GA

Oct, 30th

pre-re:Invent

All Regions



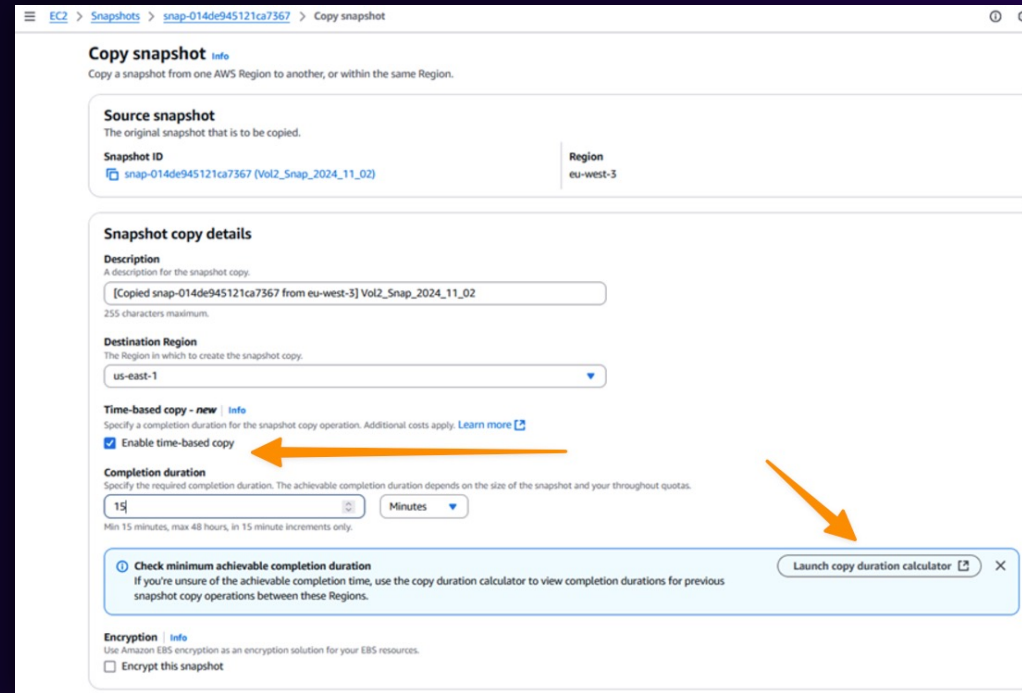
[AWS News Blog Post](#)



Time-based snapshot copy for Amazon EBS

You can now specify a desired **completion duration** (15 min - 48 h) when you copy an *Amazon Elastic Block Store (EBS)* snapshot within or between AWS Regions and/or accounts.

This will help you to meet time-based compliance and business requirements for critical workloads e.g., disaster recovery for meeting recovery point objective (*RPO*) or development/testing needs.



EC2 > Snapshots > snap-014de945121ca7367 > Copy snapshot

Copy snapshot info

Copy a snapshot from one AWS Region to another, or within the same Region.

Source snapshot
The original snapshot that is to be copied.

Snapshot ID: snap-014de945121ca7367 (Vol2_Snap_2024_11_02) Region: eu-west-3

Snapshot copy details

Description
A description for the snapshot copy.
[Copied snap-014de945121ca7367 from eu-west-3] Vol2_Snap_2024_11_02
255 characters maximum.

Destination Region
The Region in which to create the snapshot copy.
us-east-1

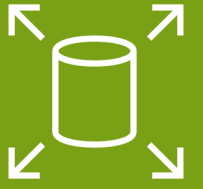
Time-based copy - new info
Specify a completion duration for the snapshot copy operation. Additional costs apply. [Learn more](#)

Enable time-based copy

Completion duration
Specify the required completion duration. The achievable completion duration depends on the size of the snapshot and your throughput quotas.
15 Minutes
Min 15 minutes, max 48 hours, in 15 minute increments only.

Check minimum achievable completion duration
If you're unsure of the achievable completion time, use the copy duration calculator to view completion durations for previous snapshot copy operations between these Regions. [Launch copy duration calculator](#) ✕

Encryption info
Use Amazon EBS encryption as an encryption solution for your EBS resources.
 Encrypt this snapshot



GA

Nov, 26th

pre-re:Invent

All Regions



[AWS News Blog Post](#)



Amazon FSx for Lustre increases throughput

UP TO 12X FOR EITHER INTER-NODE CONNECTIVITY OR WITHIN GPU INSTANCES

This announcement focuses on adding support for *Elastic Fabric Adapter (EFA)* and *NVIDIA GPUDirect Storage (GDS)* in the newest version of Amazon FSx for Lustre (from 2.15).

It means, that you can benefit much higher throughput for inter-node communication via *EFA* or much higher throughput after creating a direct path between local/remote storage and GPU memory via *GDS*.

For both, it provides 12x higher per-client throughput (up to *1200 Gbps*) in comparison to older versions.

FSx

GA

Nov, 27th

pre-re:Invent

All Regions

Specify file system details

File system details

File system name - optional [Info](#)
test-efa
Maximum of 256 Unicode letters, whitespace, and numbers, plus + - = _ : /

Deployment and storage type [Info](#)
Select a deployment type and storage type to fit your workload requirements

- Persistent, SSD
 - with EFA enabled. ←
- Persistent, HDD
 - with SSD cache
- Scratch, SSD

Throughput per unit of storage [Info](#)
Throughput (MB/s) per unit of storage (TiB)

- 125 MB/s/TiB
- 250 MB/s/TiB
- 500 MB/s/TiB
- 1000 MB/s/TiB

Storage capacity [Info](#)
4.8 TiB
Supported sizes for 1,000 MB/s/TiB EFA enabled filesystems: 4.8 TiB or increments of 4.8 TiB

Throughput capacity [Info](#)
Throughput capacity = Storage capacity (TiB) * Per unit storage throughput (MB/s)
4800 MB/s

Metadata Configuration [Info](#)
Amazon FSx provides 12K Metadata IOPS per 24 TiB of data storage capacity. You can also provision additional Metadata IOPS as needed.

- Automatic
6000 IOPS
- User-provisioned

Data compression type [Info](#)
Data compression reduces the physical disk space needed to store file data. Select LZ4 to enable data compression

NONE

Lustre version [Info](#)
Lustre version 2.15 is recommended for all new file systems.
2.15



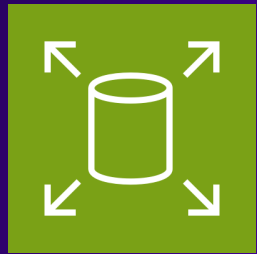
[AWS News Blog Post](#)



Data integrity protections in Amazon S3

UPDATING DEFAULT BEHAVIOR, AWS SDK UPDATE, AND NEW ALGORITHM (CRC64NVME)

With this announcement, a default behavior for calculating data integrity when uploading new objects to *Amazon Simple Storage Service (S3)* will change. **With the use of the newest AWS SDK, by default, checksum will be calculated on client-side in addition to server-side.**



GA

Dec, 1st

re:Invent

All Regions

Using supported checksum algorithms

⚠ Important

Starting on December 1, the `CRC-64NVME` checksum algorithm and improved checksum integrity features will become available in the coming weeks. These new checksum protections will be available in all AWS Regions, including the AWS GovCloud (US) and China Regions.

With Amazon S3, you can choose a checksum algorithm to validate your data during uploads. The specified checksum algorithm is then stored with your object and can be used to validate data integrity during downloads. You can choose one of the following Secure Hash Algorithms (SHA) or Cyclic Redundancy Check (CRC) checksum algorithms to calculate the checksum value:

- `CRC64-NVME` (Recommended)
- `CRC-32`
- `CRC-32C`
- `SHA-1`
- `SHA-256`



[AWS News Blog Post](#)

Storage Browser for Amazon S3



It's an **open source UI component** distributed with [AWS Amplify UI](#), that you can add to your web applications to enable end users to interact with your data stored in Amazon S3.

With this front-end component, authorized end users can browse, upload, download, copy, and delete data from Amazon S3 based on their specific permissions, which you control using AWS identity and security services or custom managed solutions.

GA

Dec, 1st

re:Invent

Open Source

Storage Browser for Amazon S3

The StorageBrowser component gives your end users a simple, visual interface for working with your data stored in Amazon S3.

[React source](#) [Feedback](#)

Theming

The `StorageBrowser` component is built on Amplify UI components so if you already have an Amplify UI theme it will just work with the Storage Browser out of the box. The components used in the Storage Browser are: `<Button>`, `<Breadcrumbs>`, `<Menu>`, `<Pagination>`, `<SearchField>`, `<Checkbox>`, `<TextField>`, and `<Message>`.

Theme object

You can use the `createTheme()` function to theme the `StorageBrowser` component

Home

Filter folders and files

Folder	Bucket	Permissions
some/path/	test	Read/Write
test-bucket	test-bucket	Read
test/	test	Read/Write

```
import * as React from 'react';
import { StorageBrowser } from './MockStorageBrowser';

export default function Example() {
  return <StorageBrowser />;
}
```

Home

1 >

Folder	Bucket	Permissions
some/path/	test	Read/Write
test-bucket	test-bucket	Read
test/	test	Read/Write

```
import * as React from 'react';
import {
  ThemeStyle,
  createTheme,
  defineComponentTheme,
} from '@aws-amplify/ui-react/server';
import { View } from '@aws-amplify/ui-react';
```

Copy



[AWS News Blog Post](#)

AWS Transfer Family **Web Apps** for S3



Fully-managed, no-code web applications that allow authenticated users to operate on objects like they are well-known files and directories for specific *Amazon S3* buckets without a need for custom clients, scripts, and *similar*.

GA

Dec, 1st

re:Invent

9 Regions

The screenshot shows the 'JeffCo File Transfer Portal' interface. At the top left is the 'JeffCo' logo and a user profile picture. At the top right, the user is identified as 'Jeff Barr'. Below the header, the breadcrumb 'Home / jeffco' is visible. The main content area is titled 'jeffco' and features a search bar with the text 'Search current folder', a 'Submit' button, and a checkbox for 'Include subfolders'. To the right of the search bar are navigation icons: a left arrow, the number '1', a right arrow, a refresh icon, and a menu icon. Below this is a table listing files and folders:

<input type="checkbox"/>	Name ^	Type ^	Last Modified ^	Size ^	
<input type="checkbox"/>	Presentations/	Folder			
<input type="checkbox"/>	Reports/	Folder			
<input type="checkbox"/>	Uploads/	Folder			
<input type="checkbox"/>	FAQ_v1.docx	docx	10/23/2024, 3:22:08 PM	13.1 kB	↓
<input type="checkbox"/>	QuotableTweetable.pptx	pptx	10/23/2024, 3:22:09 PM	123.6 kB	↓



[AWS News Blog Post](#)

Amazon FSx Intelligent-Tiering

AVAILABLE FOR AMAZON FSx FOR OPENZFS OFFERING UP TO 85% LESS THAN SSD STORAGE

FSx
ZFS

GA

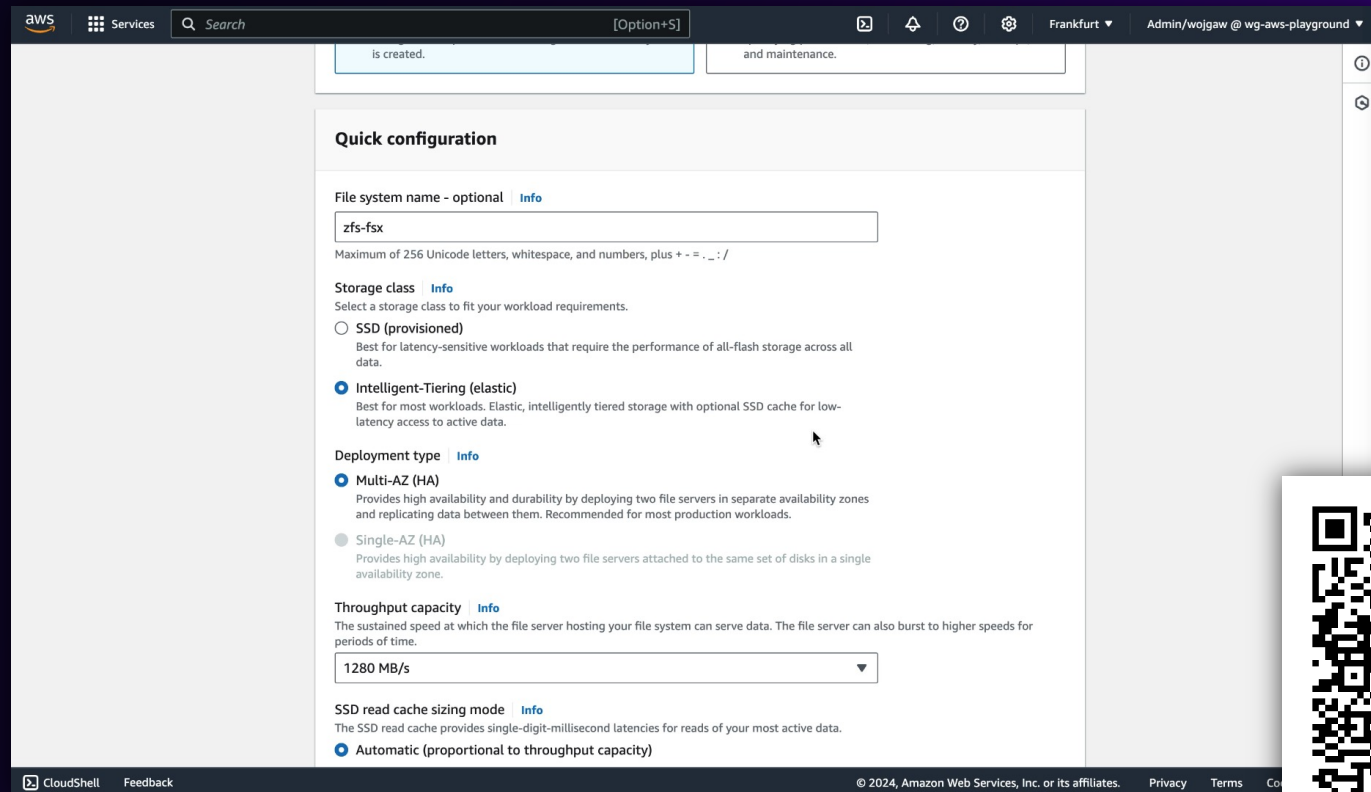
Dec, 1st

re:Invent

10 Regions

Amazon FSx offers now a new Intelligent-Tiering storage class that automatically optimize storage costs as data patterns change, without any upfront commitments.

This storage class delivers low-cost storage and **costs up to 85% less** than the FSx SSD storage class.



[AWS News Blog Post](#)



AWS Data Transfer Terminal

SECURE PHYSICAL LOCATIONS FOR UPLOADING DATA FROM YOUR STORAGE DEVICES TO CLOUD

AWS Data Transfer Terminal is a completely new offering announced this year.

It is a secure physical location which you can reserve and bring your storage devices (or AWS Snowball) to upload data faster to AWS Cloud.

The screenshot displays the AWS Data Transfer Terminal interface for a team named 'ChannyTeam'. The interface includes a navigation sidebar on the left with options for 'Data Transfer Terminal', 'Transfer Teams', 'Facilities', and 'Documentation'. The main content area shows the 'ChannyTeam' details, including a 'Transfer Team summary' with the name 'ChannyTeam' and a description 'The pilot team to test upload data.'. Below this, there are tabs for 'Reservations' and 'Personnel'. The 'Reservations' tab is active, showing a table of 'Upcoming reservations' with columns for ID, Facility, Status, Process Owner, and Starts at. Two reservations are listed: one reserved for November 29th, 2024, and one requested for December 4th, 2024. A 'Create Reservation' button is visible in the top right of the reservations section.

ID	Facility	Status	Process Owner	Starts at
rsv-3aep5e07ya	QTS. JFK Piscataway. NJ	Reserved	Channy	November 29th, 2024, 09:00
rsv-m4udf4ta92	QTS. JFK Piscataway. NJ	Requested	Channy	December 4th, 2024, 13:00

GA

Dec, 1st

re:Invent

2 Locations

(LA, NY)



[AWS News Blog Post](#)



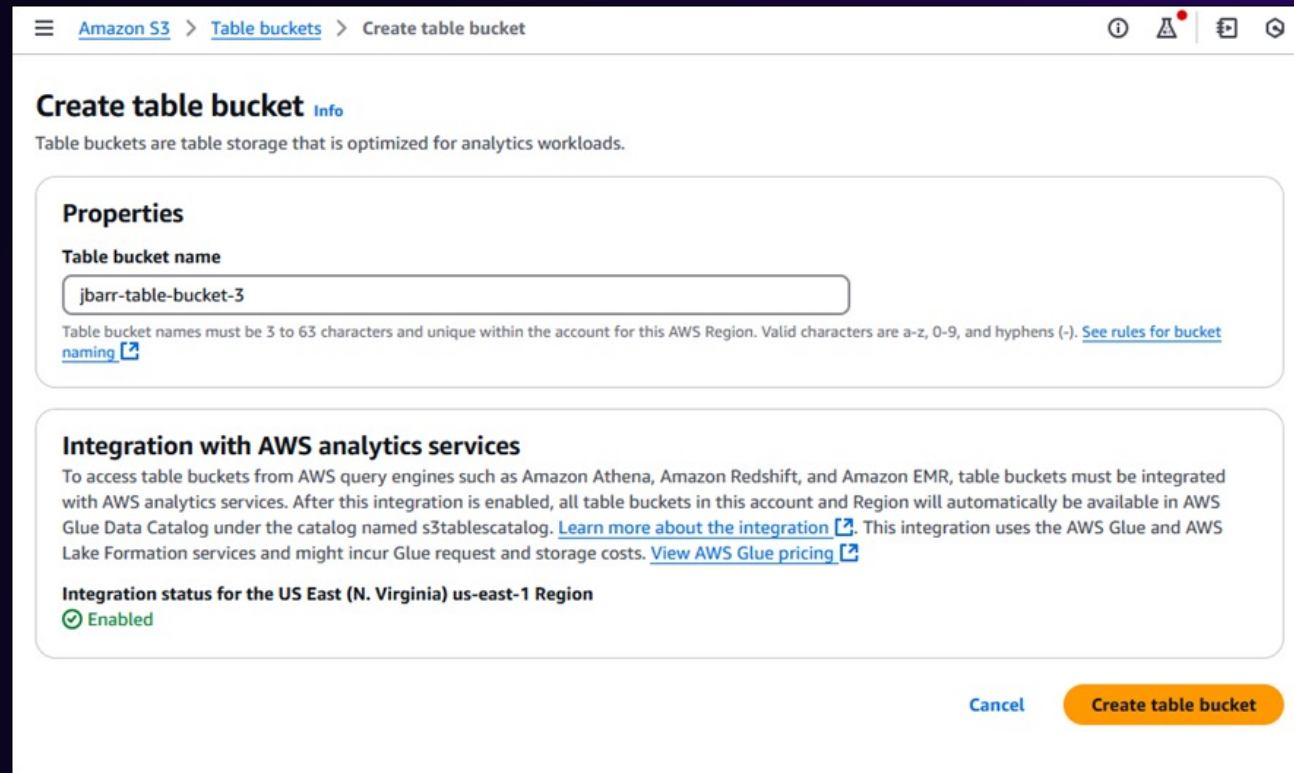
Amazon S3 Tables

STORAGE OPTIMIZED FOR ANALYTICS WORKLOADS POWERED BY APACHE ICEBERG

This announcement adds a 3rd type of S3 bucket:

S3 Tables

They act as an analytics warehouse that stores Apache Iceberg tables (using Apache Parquet format as storage layer) delivering the same S3 bucket capabilities at the same time being optimized for analytics workloads.



Amazon S3 > Table buckets > Create table bucket

Create table bucket Info

Table buckets are table storage that is optimized for analytics workloads.

Properties

Table bucket name

Table bucket names must be 3 to 63 characters and unique within the account for this AWS Region. Valid characters are a-z, 0-9, and hyphens (-). [See rules for bucket naming](#)

Integration with AWS analytics services

To access table buckets from AWS query engines such as Amazon Athena, Amazon Redshift, and Amazon EMR, table buckets must be integrated with AWS analytics services. After this integration is enabled, all table buckets in this account and Region will automatically be available in AWS Glue Data Catalog under the catalog named s3tablescatalog. [Learn more about the integration](#). This integration uses the AWS Glue and AWS Lake Formation services and might incur Glue request and storage costs. [View AWS Glue pricing](#)

Integration status for the US East (N. Virginia) us-east-1 Region

✔ Enabled

Cancel Create table bucket



GA

Dec, 3rd

re:Invent

3 Regions

Open Source



[AWS News Blog Post](#)



Amazon S3 Metadata

EASIEST AND FASTEST WAY TO MANAGE AND QUERY YOUR OBJECTS METADATA

Fully managed automated way for capturing objects metadata across your S3 buckets.

It is an opt-in mechanism on per bucket level, that stores metadata when objects are added and modified.

It stores that in an Apache Iceberg table with schema containing over 20 fields and being accessible for AWS analytics services and open source tools like Apache Spark.



PP

Dec, 3rd

re:Invent

3 Regions

Open Source

The screenshot shows the AWS console interface for the 'jbarr-data-bucket-1' bucket. The 'Metadata - New' tab is selected, displaying a 'Metadata - New' configuration page. The page includes a 'Delete' button and a 'Create metadata configuration' button. Below this, there are two main sections: 'Getting started with metadata' and 'Overview'. The 'Getting started with metadata' section contains two steps: 'Step 1 - Create a metadata configuration for your bucket' and 'Step 2 - Analyze and build with your metadata tables'. The 'Overview' section provides details about the metadata table, including its status (Active), ARN, destination table bucket, namespace, and metadata table name. A 'Go to Athena query editor' button is also visible.


jbarr-data-bucket-1

Objects | **Metadata - New** | Properties | Permissions | Metrics | Management | Access Points

Metadata - New [Info](#) [Refresh](#) [Delete](#) [Create metadata configuration](#)


Accelerate data discovery with automatically generated, near real-time metadata for objects in this bucket, stored in a fully managed Apache Iceberg table. You can query this metadata table to identify and prepare data for use in business analytics and machine learning, identify AI-generated data, retrieve content, and more.

▼ Getting started with metadata



Step 1 - Create a metadata configuration for your bucket

Create a configuration to generate and store object metadata in an S3 managed metadata table. As long as the configuration is active on the bucket, S3 continuously updates the metadata table to reflect the latest changes.



Step 2 - Analyze and build with your metadata tables

Your metadata table is instantly queryable by any AWS query engine that supports Apache Iceberg, such as Amazon Athena, Amazon QuickSight, and Apache Spark. Use the wide range of AWS services to analyze and build with your metadata. [Learn more](#)

Overview

[Go to Athena query editor](#)

You can interact with your metadata table by using any AWS query engine that supports Apache Iceberg. Before you can view the generated metadata table, your destination table bucket must be integrated with AWS Glue Data Catalog. [Learn more](#)

Status Active	Destination table bucket jbarr-table-bucket-1
Metadata table ARN arn:aws:s3tables:us-east-2: :bucket/jbarr-table-bucket-1/table/cbc78a9a-7a36-4278-a231-f2a2b9363855	Namespace aws_s3_metadata
	Metadata table name jbarr_data_bucket_1_table



[AWS News Blog Post](#)



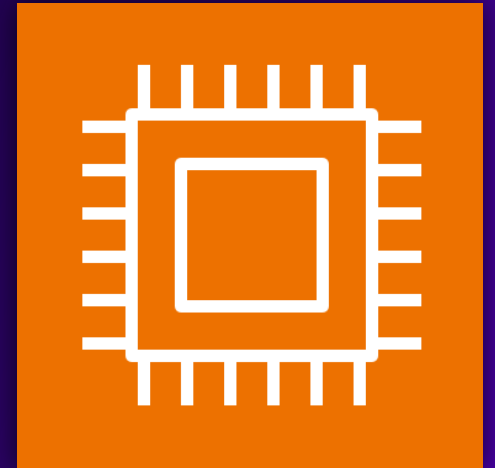
Other Announcements



1. Amazon S3 now supports up to 1M buckets per AWS Account (pre-re:Invent)
2. Amazon S3 adds new functionality for conditional writes (pre-re:Invent)
3. **Amazon S3 Express One Zone** updates:
 - a) Support for the ability to append data to an object (pre-re:Invent)
 - b) Three additional regions (pre-re:Invent)
 - c) Support for conditional deletes (pre-re:Invent)
 - d) Support for S3 Lifecycle Expirations (pre-re:Invent)
4. AWS Database Migration Service (DMS) supports data masking (pre-re:Invent)
5. Cross-account replication for Amazon EFS (pre-re:Invent)
6. AWS Database Migration Service automates schema conversion using AI (re:Invent)

Compute and Containers

16 announcements (5 *pre-re:Invent*)



Predictive scaling for Amazon ECS

PROACTIVELY SCALE CLUSTERS REDUCING OVERPROVISIONING WHILE IMPROVING AVAILABILITY

Predictive Scaling is a new scaling policy within Amazon ECS Service Auto Scaling, that is designed to anticipate demand surges by using machine learning (ML) algorithms.

The goal here is to scale your services ahead of surge, reducing overprovisioning while improving application responsiveness and availability.



GA

Nov, 21st

pre-re:Invent

All Regions

Scaling policies		Scheduled actions	Recommendations		
Scaling policies (5) Info		Actions	Create scaling policy		
<input type="text" value="Find a scaling policy"/>		Refresh			
Name ▲	Type ▼	Status ▼	Scaling mechanism	Target metric	Actions
<input type="radio"/> ps-cpu-10	Predictive Scaling	Forecast only	Add capacity based on forecast	ECSServiceCPUUtilization (10%) View chart	View
<input type="radio"/> ps-cpu-25	Predictive Scaling	Forecast only	Add capacity based on forecast	ECSServiceCPUUtilization (25%) View chart	View
<input type="radio"/> ps-cpu-55	Predictive Scaling	Forecast only	Add capacity based on forecast	ECSServiceCPUUtilization (55%) View chart	View
<input checked="" type="radio"/> ps-cpu-60	Predictive Scaling	Forecast and scale	Add capacity based on forecast	ECSServiceCPUUtilization (60%) View chart	View
<input type="radio"/> tt-cpu-60	Target Tracking	Active	Maintain target	ECSServiceAverageCPUUtilization (60%)	View



Future-dated Capacity Reservations

UP TO 120 DAYS IN ADVANCE RESERVATION FOR AMAZON EC2 ON-DEMAND INSTANCES

With this announcement, **EC2 On-Demand Capacity Reservations** supports scheduling reservation up to 120 days in advance, with no upfront charges for scheduling the reservation.

After reviewing and approving the reservation, it will be scheduled to activate and be ready for the requested at a given date.

Capacity Reservation details

You can now request future-dated Capacity Reservations. With future-dated Capacity Reservations, the capacity becomes available on a future date that you specify [Learn more](#)

Capacity Reservation starts - new | [Info](#)
Specify when the Capacity Reservation should become available for use.

Immediately
The Capacity becomes available for use immediately after you create the Capacity Reservation

At a future date
The Capacity Reservation will become available for use at a date and time that you specify. Supported with C, I, M, R, and T instance families only. Minimum 100 vCPUs.

Start date (UTC-07:00)
2024/12/02 09:00:00
Specify a date between 5 and 120 days in the future.

Commitment duration | [Info](#)
Specify the minimum duration for which you commit to keeping the Capacity Reservation after it becomes active in your account. Use the recommended value to improve EC2's ability to support your request.
14 days
Minimum: 1 day, Recommended: 14 days

Capacity Reservation ends | [Info](#)
Specify when the Capacity Reservation should end.

At a specific time
The Capacity Reservation will end automatically at a date and time that you specify.

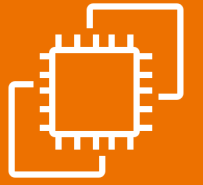
Manually
The Capacity Reservation remains active in your account until you manually cancel it.

Capacity Reservation continues until you manually cancel it
The Capacity Reservation has no end date. It will remain active in your account and continue to be billed until you manually cancel it.

Instance eligibility | [Info](#)
Indicate the criteria for instances that can run in this Capacity Reservation.

Open: Any instance with matching attributes.
All instances with matching attributes (instance type, platform, and Availability Zone) automatically run in the capacity.

Targeted: Only instances that target the Capacity Reservation.
When you launch instances, you must specify the reservation ID or the reservation resource group ARN associated with this reservation.



GA

Nov, 25th

pre-re:Invent

All Regions



[AWS News Blog Post](#)



New instances for Amazon EC2

UPDATES FOR STORAGE AND GPU-POWERED EC2 INSTANCE FAMILIES

- [i8g](#): powered by AWS Graviton4 and 3rd gen AWS Nitro SSDs (2 regions)
- [i7ie](#): with Intel Xeon and low-latency NVMe storage (6 regions)
- [p5en](#): with NVIDIA H200 Tensor Core GPUs and EFA v3 networking (3 regions)
- [p6](#): with NVIDIA Blackwell GPUs (*coming soon*)

Capacity Blocks requirements

Capacity attributes

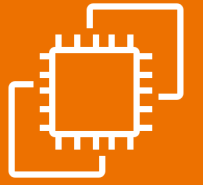
Instance type
p5en.48xlarge
192 vCPU 2 TiB Memory 1128 GB HBM3 GPU

Platform
Linux/UNIX

Platforms other than Linux/UNIX incur a usage-based charge while your instances are running.
[Learn more](#)

Total capacity
16 instance(s)
The instance purchase quantity must be 1, 2, 4, 8, 16, 32, or 64 instances.

Duration
The number of days for which you want to purchase Capacity Block.
10 day(s)
The duration must be within 1-14 days, 21 days, or 28 days.



GA

CS

Dec, 1st

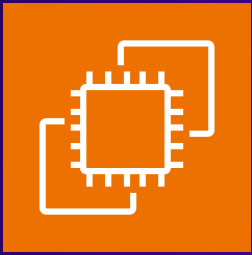
Dec, 2nd

re:Invent

Limited
Regions

Amazon EKS Auto Mode

AMAZON ELASTIC KUBERNETES SERVICE (EKS): PROVISIONING TO MAINTENANCE WITH 1-CLICK



GA

Dec, 1st

re:Invent

All Regions

Amazon EKS Auto Mode is streamlining Kubernetes cluster management for compute, storage, and networking – from provisioning to maintenance with a single click.

You can create new clusters with Auto Mode or enable it for existing clusters.

Configure cluster

Configuration options - new [Info](#)
Choose how you would like to configure the cluster.

Quick configuration (with EKS Auto Mode) - new
Quickly create a cluster with production-grade default settings. The configuration uses EKS Auto Mode to automate infrastructure tasks like creating nodes and provisioning storage.

Custom configuration
To change default settings prior to creation, choose this option. This configuration gives the option to use EKS Auto Mode and customize the cluster's configuration.

Cluster configuration

Name
Use the auto-generated name or enter a unique name for this cluster. This property cannot be changed after the cluster is created.
 [🔄](#)
The cluster name should begin with letter or digit and can have any of the following characters: the set of Unicode letters, digits, hyphens and underscores. Maximum length of 100.

Kubernetes version [Info](#)
Select Kubernetes version for this cluster.

Cluster IAM role [Info](#)
Select the Cluster IAM role to allow the Kubernetes control plane to manage AWS resources on your behalf. This cannot be changed after the cluster is created. To create a new custom role, follow the instructions in the [Amazon EKS User Guide](#).
 [🔄](#) [Create recommended role](#)

Node IAM role [Info](#)
Nodes need an EC2 Instance IAM Role to launch and register with a cluster. To create a new custom role, follow the instructions in the [Amazon EKS User Guide](#).
 [🔄](#) [Create recommended role](#)

VPC [Info](#)
Select a VPC to use for your EKS cluster resources.
 [🔄](#) [Create VPC](#)

Subnets [Info](#)
Choose the subnets in your VPC where the control plane may place elastic network interfaces (ENIs) to facilitate communication with your cluster. To create a new subnet, go to the corresponding page in the [VPC console](#).
 [🔄](#) [Clear selected subnets](#)

us-west-2a 172.31.16.0/20 us-west-2c 172.31.0.0/20

[▶ View quick configuration defaults](#)

[Cancel](#) [Create](#)

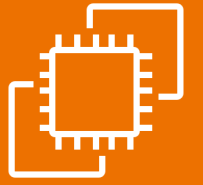


[AWS News Blog Post](#)



Amazon EKS Hybrid Nodes

ATTACH ON-PREMISES AND EDGE INFRASTRUCTURE AS NODES TO EKS CLUSTERS IN THE CLOUD

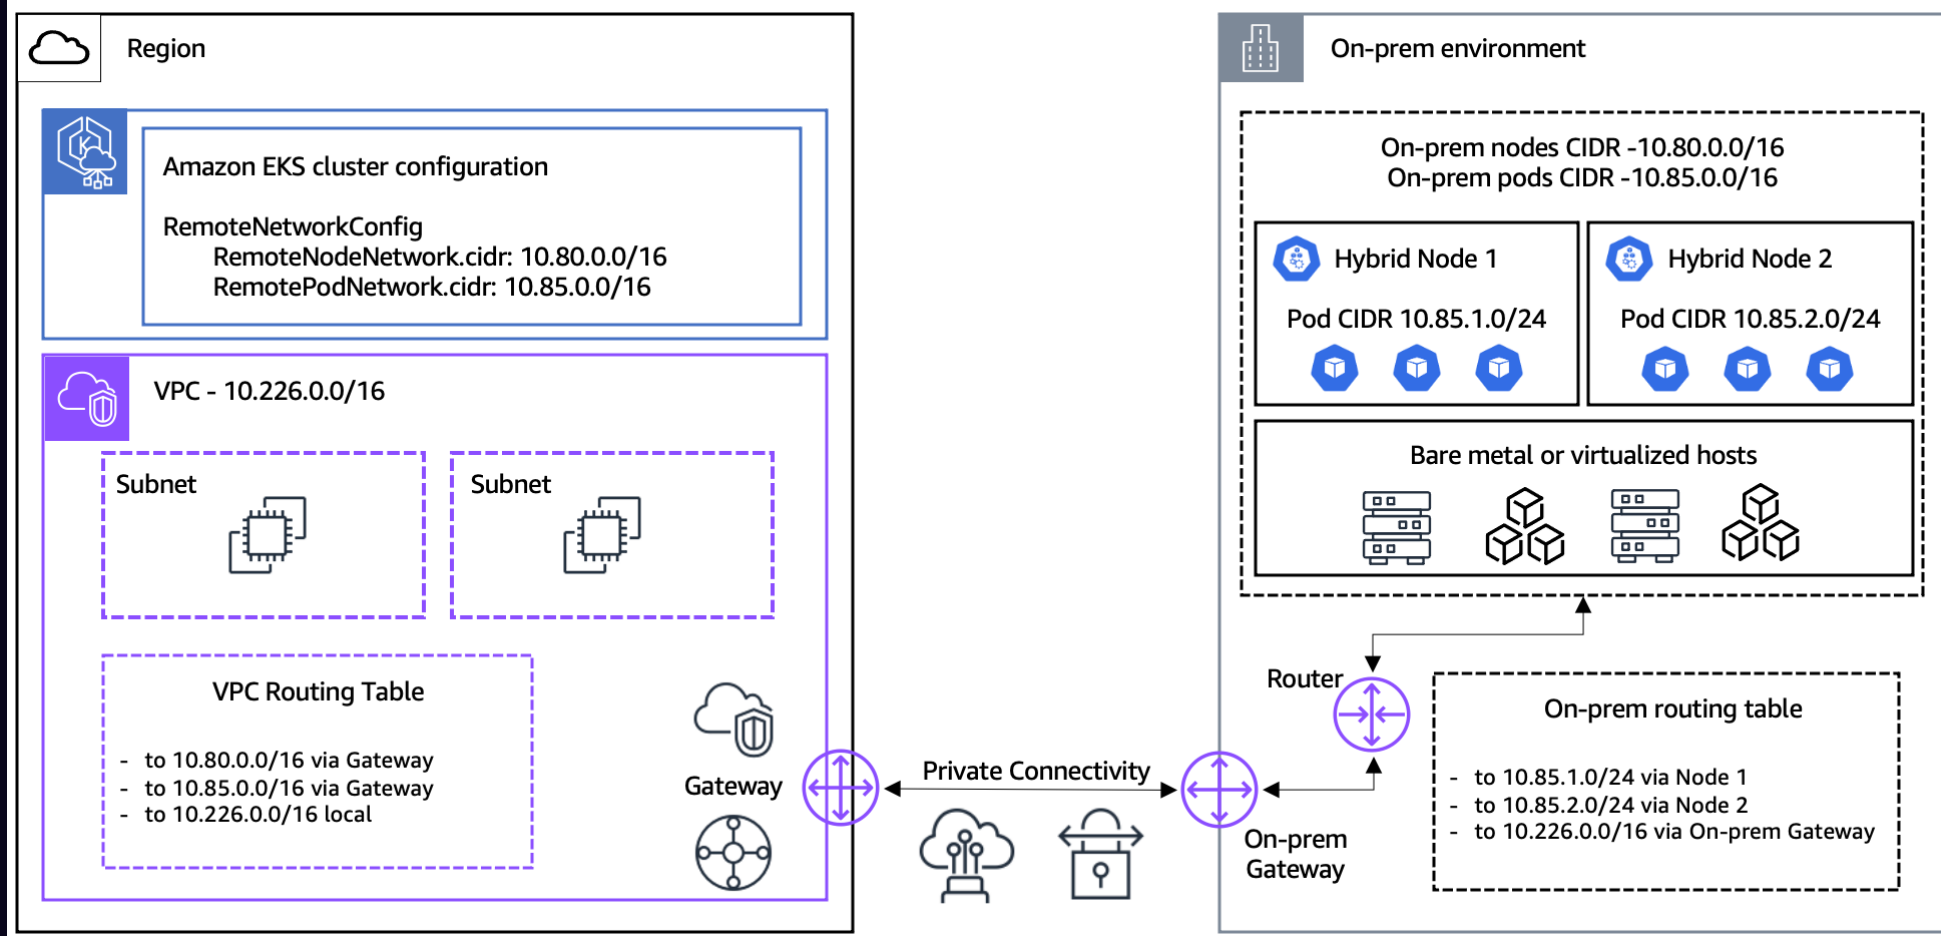


GA

Dec, 1st

re:Invent

All Regions



Amazon Elastic VMware Service (EVS)

NATIVE AWS SERVICE TO RUN VMWARE CLOUD FOUNDATION (VCF) WITHIN AMAZON VPC

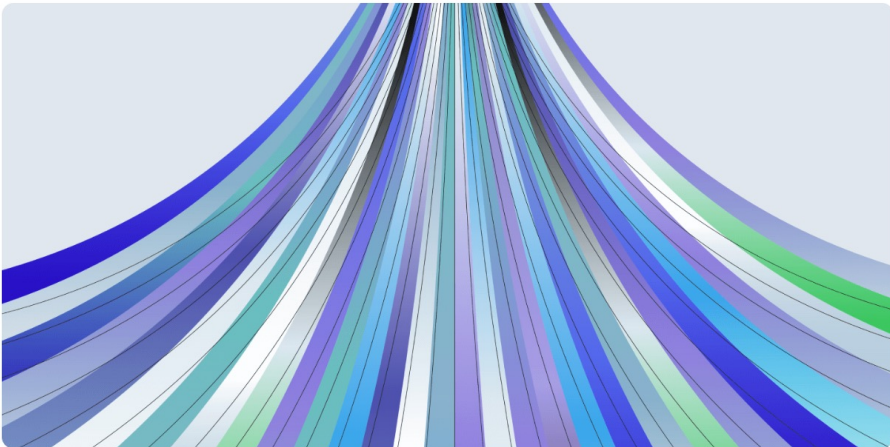
With Amazon Elastic VMware Service (EVS) preview, AWS is strengthening its commitment to provide customers a native pathway to migrate and modernize their VMware-based workloads on AWS to achieve the scalability, agility, and cost benefits of the cloud.

AWS is committed to delivering VMware customers the most choice and flexibility when it comes to running their VMware-based workloads in the cloud, and continue to expand our migration and modernization pathways, programs, and partnerships to support their needs.



© 2025, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Run VCF on AWS Migrate to Amazon EC2 Modernize on AWS



Relocate your workloads to Amazon Elastic VMware Service

Amazon Elastic VMware Service (Amazon EVS) is the fastest path to migrate and operate VMware workloads on AWS. Get the scale, resilience, and performance of AWS together with familiar VMware Cloud Foundation (VCF) software and tools. Choose from self-managed or managed options to meet your business needs for control and customization of your VMware architecture. Embrace AWS native and managed services for analytics, storage, databases, serverless compute, generative AI, and more to drive business agility, innovation, and transformation.

[Learn more about Amazon EVS \(Preview\)](#)

P

Dec, 1st

re:Invent

Private
Preview



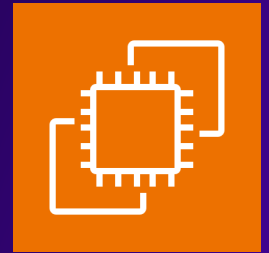
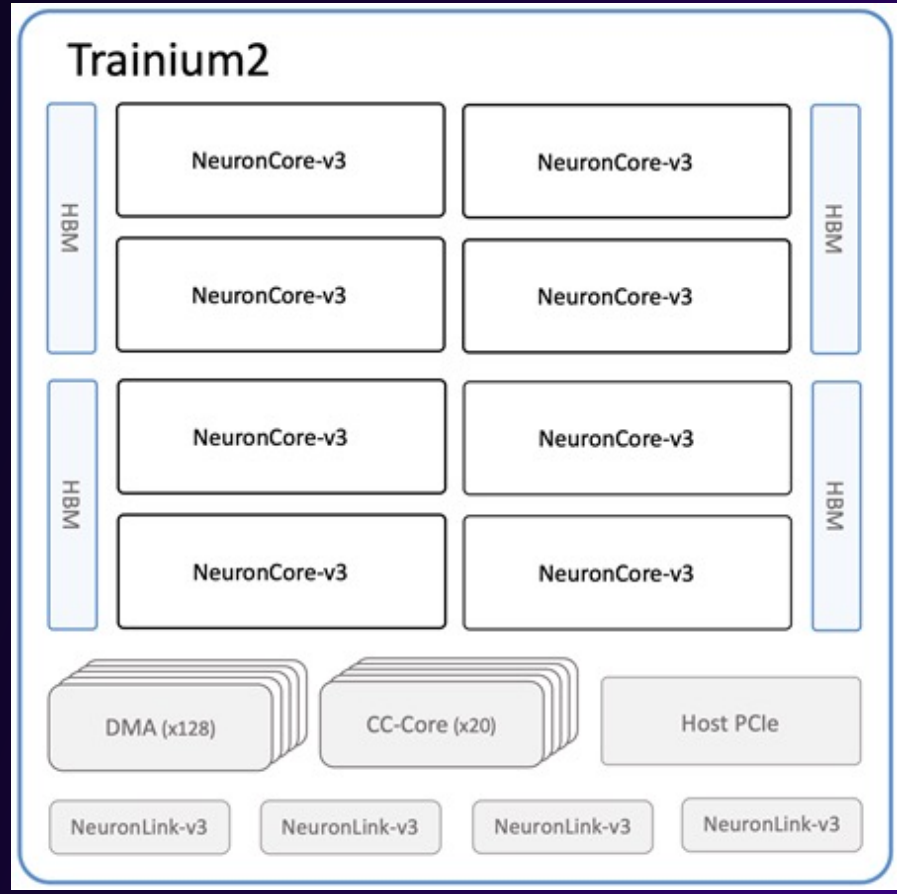
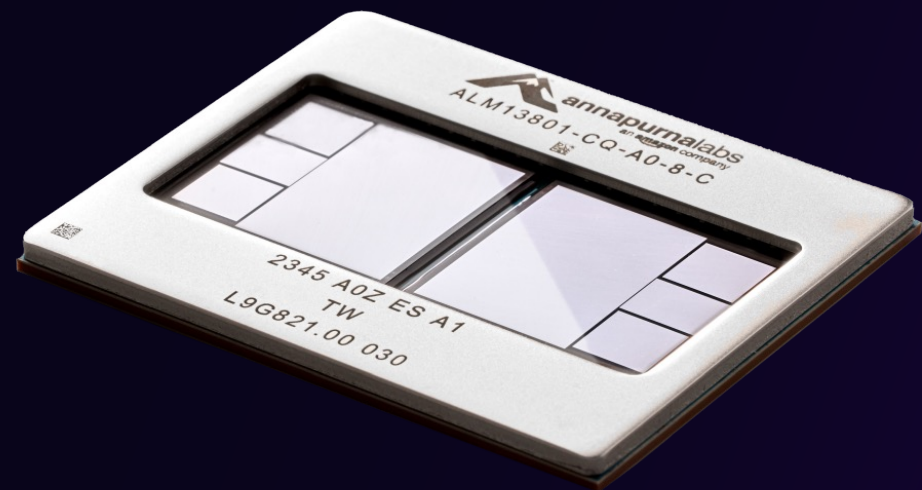
[AWS M&M Blog Post](#)

AWS Trainium2 (trn2)

GENERAL AVAILABILITY OF TRN2 INSTANCES AND PREVIEW OF TRN2 ULTRASERVERS

AWS Trainium2 instances are 4x faster, offer 4x more memory bandwidth and 3x more memory capacity than first generation of AWS Trainium.

This makes EC2 instances with those processors the most powerful compute options for ML training and inference – with emphasis on Trn2 UltraServers featuring 64 trn2 chips.



GA

P

Dec, 3rd

re:Invent

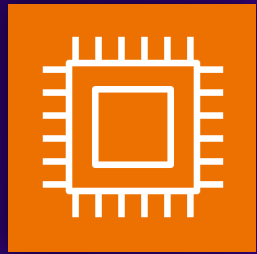
1 Region



[AWS News Blog Post](#)



Other Announcements



1. Service versioning and deployment history for Amazon ECS services (pre-re:Invent)
2. Other Amazon Elastic Kubernetes Service (EKS) updates:
 - a) Amazon EKS simplifies providing IAM permissions to EKS add-ons (pre-re:Invent)
 - b) Amazon EKS enhances Kubernetes control plane monitoring (pre-re:Invent)
3. Container Insights with enhanced observability now available in Amazon ECS (re:Invent)
4. AWS shares data center power usage effectiveness numbers for the first time (re:Invent)

Databases

14 announcements (10 *pre-re:Invent*)

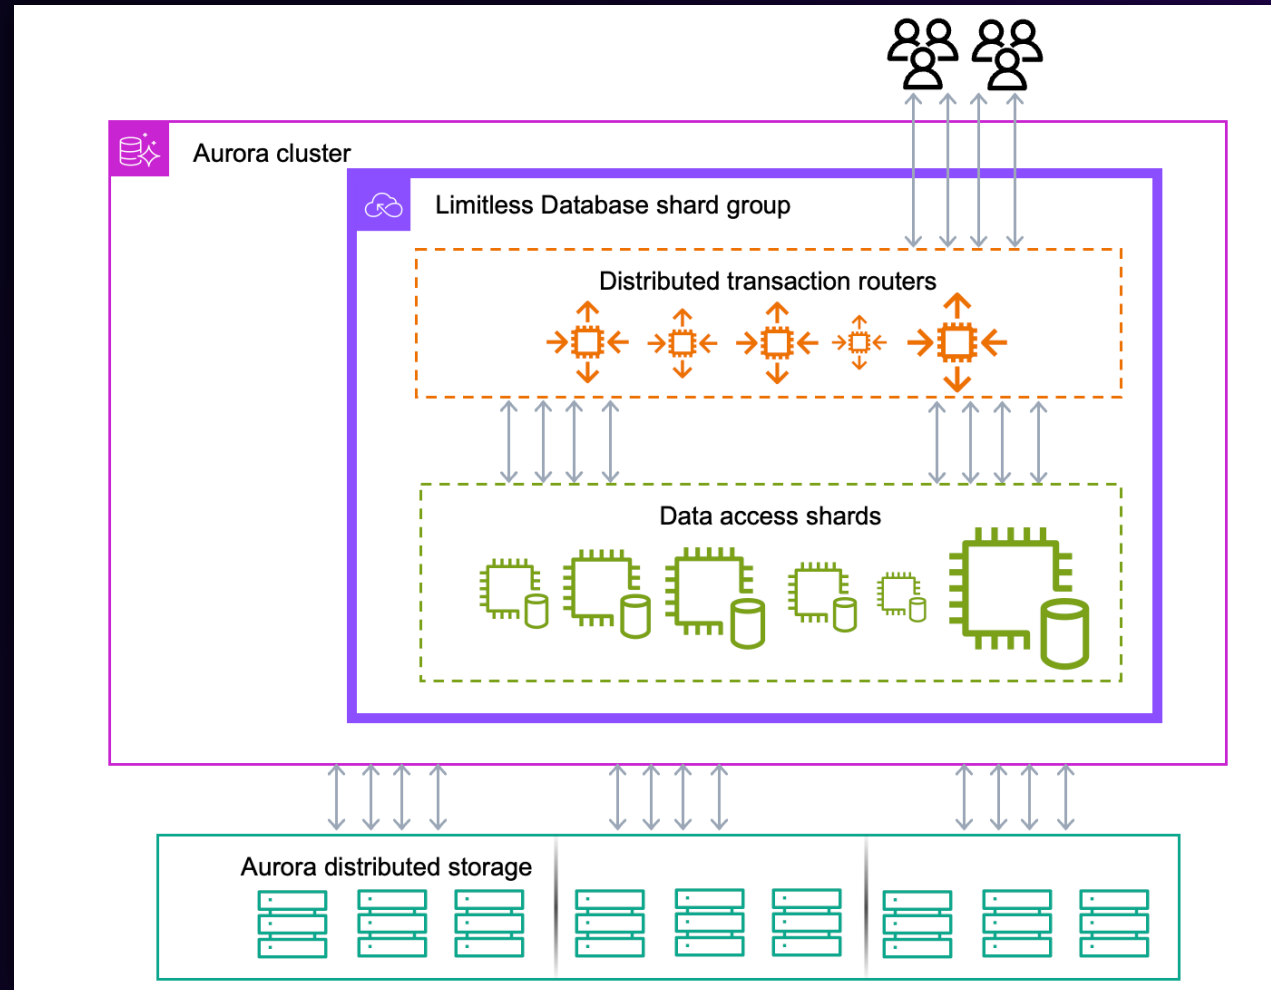


Amazon Aurora PostgreSQL Limitless

GENERAL AVAILABILITY OF LIMITLESS CLUSTERS

Scaling relational database workloads on Amazon Aurora beyond limits of a single writer instance, without a need to create custom application logic or managing multiple databases.

It does not just shard the storage layer, but also automatically scales compute up and down – depending on your workloads usage and without any need to reprovision.



GA

Oct, 31st

pre-re:Invent

10 Regions



[AWS News Blog Post](#)

Amazon DynamoDB: Warm Throughput

EASY PRE-WARMING AMAZON DYNAMODB TABLES AND INDEXES

After introducing this update, when creating table or index, a warm throughput value provides visibility into the number of read and write operations your DynamoDB tables can readily handle, while pre-warming let's you proactively increase the value to meet future traffic demands.

Warm throughput - new [Info](#)

Increasing the warm throughput value pre-warms your table to handle planned peak events without throttling or scaling delays. By default, warm throughput values are visible for all tables and global secondary indexes. These values automatically adjust as you increase your provisioned throughput or on-demand consumption without extra charges, but if you choose to change them manually, additional charges apply. [Learn more about Amazon DynamoDB pricing](#)

Read units per second

Keep default values

Increase warm throughput

Read units per second

12,000

Write operations per second

Keep default values

Increase warm throughput

Write units per second

4,000

Secondary indexes [Info](#)

[Delete](#) [Create local index](#) [Create global index](#)

Name	Type	Partition key	Sort key	Projected attributes	Warm throughput
No indexes					

Use secondary indexes to perform queries on attributes that are not part of your table's primary key.

[Create global index](#)

Encryption at rest [Info](#)

All user data stored in Amazon DynamoDB is fully encrypted at rest. By default, Amazon DynamoDB manages the encryption key and you do not incur any fees for using it.

© 2024, Amazon Web Services, Inc. or its affiliates. [Privacy](#) [Terms](#) [Co](#)



GA

Nov, 13th

pre-re:Invent

All Regions



[AWS Announcement](#)



Price reduction for Amazon DynamoDB

50% REDUCTION FOR ON-DEMAND THROUGHPUT AND UP TO 67% FOR GLOBAL TABLES

Our continued engineering investments on how efficiently we can operate Amazon DynamoDB allow us to identify and pass on cost savings to you.

From November 1st, 2024, DynamoDB has reduced prices for **on-demand throughput by 50%** and **global tables by up to 67%**, making it more cost-effective than ever to build, scale, and optimize applications.



GA

Nov, 14th

pre-re:Invent

All Regions



Scale to zero for Amazon Aurora Serverless

AMAZON AURORA SERVERLESS V2 SUPPORTS SCALING TO 0 AURORA CAPACITY UNITS (ACUS)



GA

Nov, 20th

pre-re:Invent

All Regions

Amazon Aurora Serverless v2 supports scaling to **zero** Aurora Capacity Units (ACUs).

This launch enables the database to automatically pause after a period of inactivity based on database connections.

When the first connection is requested, the database will automatically resume and scale to meet the application demand.



[AWS Database Blog Post](#)



Amazon MemoryDB Multi-Region

ACTIVE-ACTIVE, MULTI-REGION, FULLY MANAGED VALKEY AND REDIS-COMPATIBLE DATABASE

Amazon Memory DB from now supports creation of multi-region clusters in active-active configuration mode for Valkey and Redis that allows for microsecond read and single-digit write latencies.

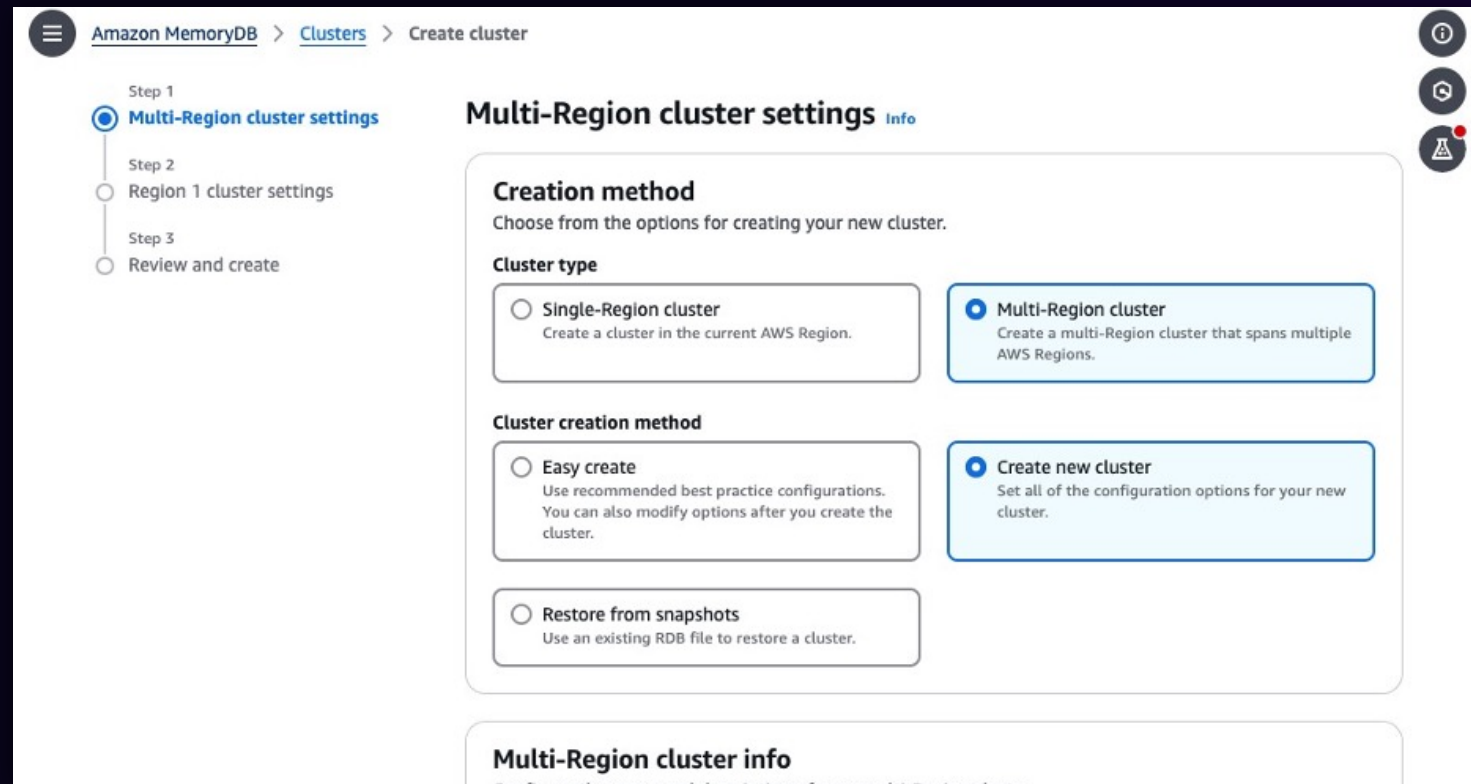


GA

Dec, 1st

re:Invent

12 Regions



[AWS News Blog Post](#)



Oracle Database@AWS

FULLY-MANAGED AND AWS-POWERED ORACLE EXADATA DATABASE SERVICE

Oracle Database@AWS is a new offering from AWS and Oracle available from now in limited preview.

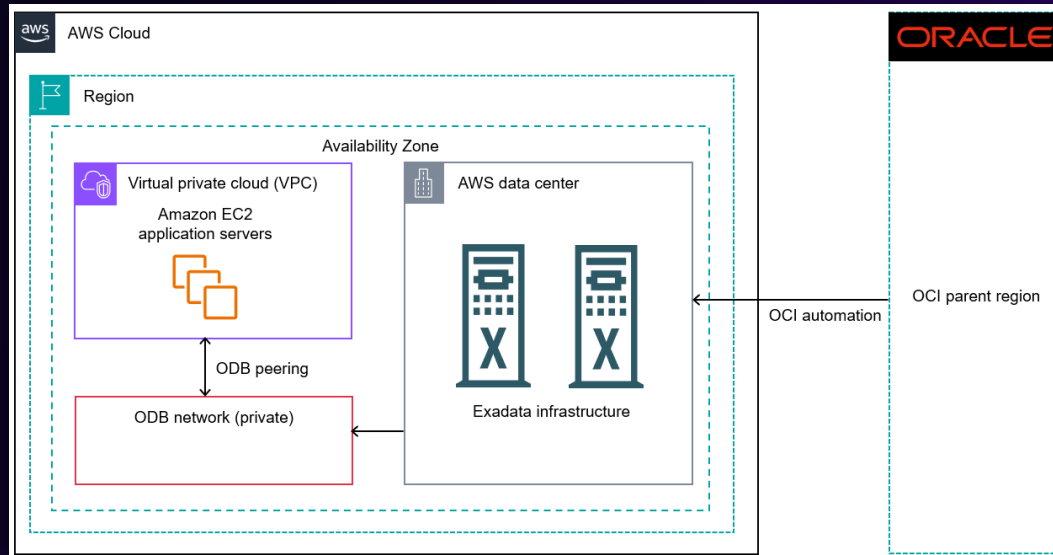
It enables customers to access Oracle Database Services on Oracle Cloud Infrastructure (OCI) managed Exadata infrastructure within AWS data centers.

Customers can easily and quickly migrate their Oracle Database workloads, including Oracle Real Application Clusters (RAC) workloads, to Oracle Exadata Database Service within AWS with minimal to no changes.

aws | ORACLE

Oracle Database@AWS

Easily and quickly migrate Oracle Databases to AWS with minimal changes



P

Dec, 1st

re:Invent

1 Region



[AWS Announcement](#)

Multi-Region Strong Consistency

AMAZON DYNAMODB GLOBAL TABLES MULTI-REGION STRONG CONSISTENCY



PP

Dec, 3rd
re:Invent
3 Regions

Amazon DynamoDB Global Tables are now supporting in preview multi-region strong consistency.

For now you can enable it only in 3 US-based regions: us-east-1 (N. Virginia), us-east-2 (Ohio), and us-west-2 (Oregon).

The screenshot shows the AWS Management Console interface for a DynamoDB Global Table named 'TestingStrongConsistency'. The 'Global tables' tab is selected, showing a summary where 'Multi-Region consistency' is set to 'Strong' and the 'Current Region' is 'US East (N. Virginia) us-east-1'. A warning message states: 'Multi-Region strong consistency is a preview feature, and should not be used for production workloads.' Below this, the 'Replicas (3)' section shows three active replicas in different regions: US West (Oregon) us-west-2, US East (Ohio) us-east-2, and US East (N. Virginia) us-east-1 (Current Region). Each replica is in 'On-demand' capacity mode with an active status. The endpoints for each replica are listed as dynamodb.us-west-2.amazonaws.com, dynamodb.us-east-2.amazonaws.com, and dynamodb.us-east-1.amazonaws.com respectively.

Replication Region	Status	Capacity mode	Endpoint
US West (Oregon) us-west-2	Active	On-demand	dynamodb.us-west-2.amazonaws.com
US East (Ohio) us-east-2	Active	On-demand	dynamodb.us-east-2.amazonaws.com
US East (N. Virginia) us-east-1 Current Region	Active	On-demand	dynamodb.us-east-1.amazonaws.com



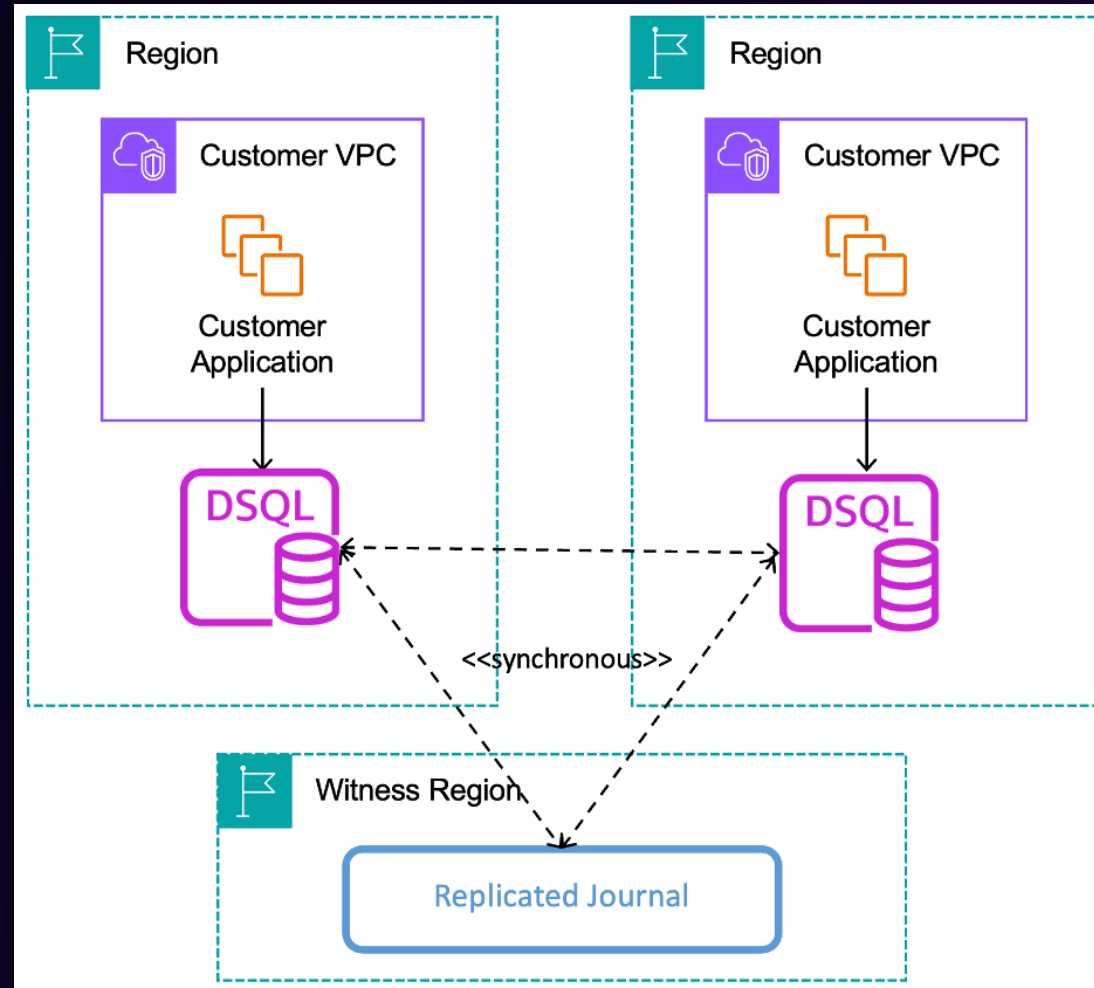
Amazon Aurora DSQL

SERVERLESS, POSTGRESQL-COMPATIBLE DATABASE WITH ACTIVE-ACTIVE MULTI-REGION SETUP

Amazon Aurora DSQL is a serverless distributed SQL database with multi-region active-active capabilities.

Aurora DSQL is PostgreSQL compatible and its active-active distributed architecture is designed for 99.99% single-region and 99.999% multi-region availability with no single point of failure, and automated failure recovery.

This ensures that all reads and writes to any regional endpoint are strongly consistent and durable.



PP

Dec, 3rd

re:Invent

3 Regions



[AWS Database Blog Post](#)

Other Announcements



1. [Provisioned Timestream Compute Units \(TCUs\) for Amazon Timestream for Live Analytics](#) (pre-re:Invent)
2. [Amazon RDS for PostgreSQL supports pgvector 0.8.0](#) (pre-re:Invent)
3. **Valkey-related AWS Services updates:**
 - a) [Support for Valkey 8.0 in Amazon ElastiCache](#) (pre-re:Invent)
 - b) [Valkey GLIDE 1.2](#) (pre-re:Invent)
4. **Amazon RDS Blue/Green Deployments enhancements:**
 - a) [Storage volume shrink](#) (pre-re:Invent)
 - b) [Minor version upgrade for PostgreSQL](#) (pre-re:Invent)

ClockBound

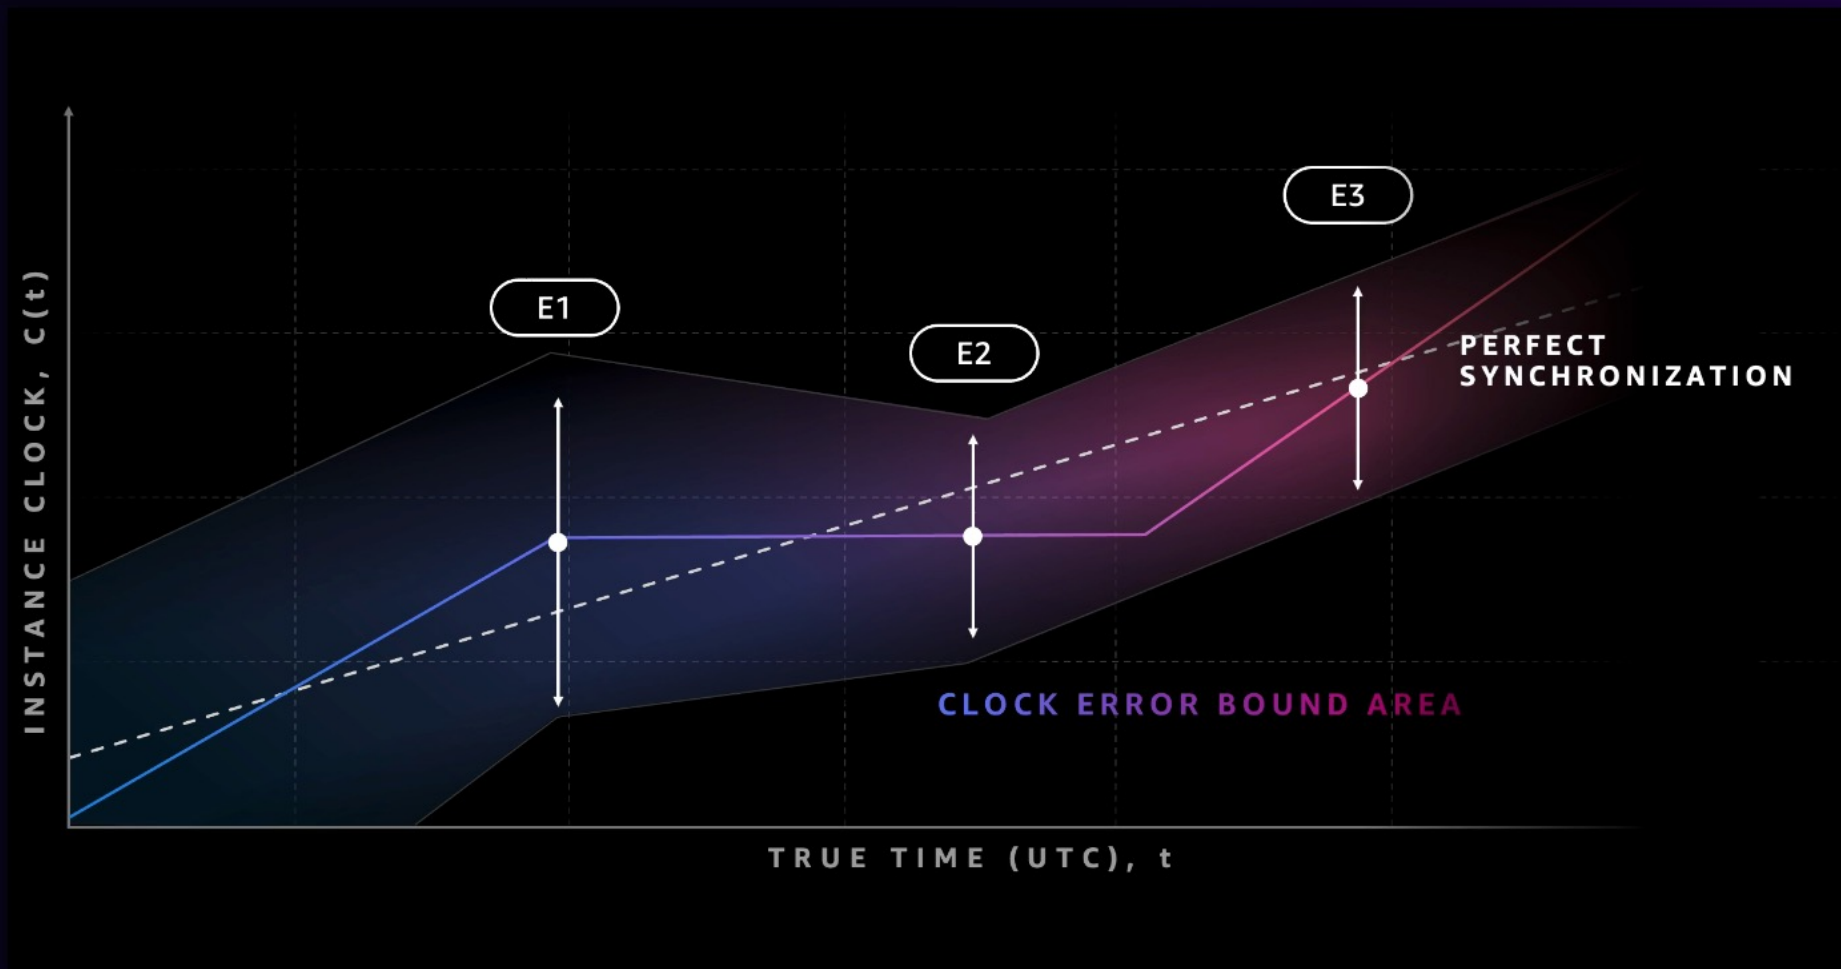
GENERATING, COMPARING, AND ERROR-CORRECTING TIMESTAMPS FROM TRUSTED TIME SERVICE

Dec 5th

re:Invent

Open Source

GA



[aws/clock-bound](#)
GitHub Repository



Networking and Content Delivery

8 announcements (6 *pre-re:Invent*)



Amazon CloudFront supports gRPC APIs

DEPLOY CLOUDFRONT DISTRIBUTIONS IN FRONT OF GRPC API ENDPOINTS

With this announcement, you can deploy Amazon CloudFront distributions in front of your gRPC API endpoints and benefit from two things:

1. Latency reduction between client and API, thanks to the nature of Content Delivery Network (CDN).
2. Additional security services deployed on edge locations – including traffic encryption, HTTP headers validation (via AWS Web Application Firewall) and AWS Shield Standard protection from distributed denial of service (DDoS) attacks.

Viewer

Viewer protocol policy

- HTTP and HTTPS
- Redirect HTTP to HTTPS
- HTTPS only

Allowed HTTP methods

- GET, HEAD
- GET, HEAD, OPTIONS
- GET, HEAD, OPTIONS, PUT, POST, PATCH, DELETE

Cache HTTP methods

GET and HEAD methods are cached by default.

- OPTIONS

Allow gRPC requests over HTTP/2 [Info](#)

- Enable



GA

Nov, 20th

pre-re:Invent

All Regions



[AWS News Blog Post](#)

Amazon CloudFront VPC Origins

EXPOSE SECURELY PRIVATE RESOURCES FROM AMAZON VPC VIA AMAZON CLOUDFRONT

Amazon CloudFront VPC Origins is a new feature that enables content delivery from applications hosted in private subnets within Amazon VPC.

This makes it easy to secure web applications, allowing you to focus on growing your businesses while improving security and maintaining high-performance and global scalability of Content Delivery Network (CDN) via Amazon CloudFront.

[AWS](#) > [Documentation](#) > [Amazon CloudFront](#) > [Developer Guide](#)

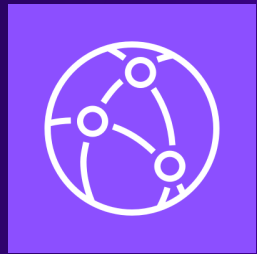
Restrict access with VPC origins

[PDF](#) [RSS](#) Focus mode

You can use CloudFront to deliver content from applications that are hosted in your virtual private cloud (VPC) private subnets. You can use [Application Load Balancers \(ALBs\)](#), [Network Load Balancers \(NLBs\)](#), and [EC2 instances in private subnets](#) as VPC origins.

Here are some reasons why you might want to use VPC origins:

- **Security** – VPC origins is designed to enhance the security posture of your application by placing your load balancers and EC2 instances in private subnets, making CloudFront the single point of entry. User requests go from CloudFront to the VPC origins over a private, secure connection, providing additional security for your applications.
- **Management** – VPC origins reduces the operational overhead required for secure connectivity between CloudFront and origins. You can move your origins to private subnets with no public access, and you don't have to implement access control lists (ACLs), secret shared headers, or other mechanisms to restrict access to your origins. This way, you don't have to invest in undifferentiated development work to secure your web applications with CloudFront.
- **Scalability and performance** – VPC origins helps you to secure your web applications, freeing up time to focus on growing your critical business applications while improving security and maintaining high-performance and global scalability with CloudFront. VPC origins streamlines security management and reduces operational complexity so that you can use CloudFront as the single point of entry for your applications.



GA

Nov, 20th

pre-re:Invent

All Regions



[AWS News Blog Post](#)

Amazon CloudFront Functions Update

CONDITIONALLY UPDATE OR CHANGE ORIGIN SERVERS ON EACH REQUEST

Amazon CloudFront now supports origin modification within CloudFront Functions, enabling you to conditionally change or update origin servers on each request.

You can now write custom logic in CloudFront Functions to overwrite origin properties, use another origin in your CloudFront distribution, or forward requests to any public HTTP endpoint.

[AWS](#) > [Documentation](#) > [Amazon CloudFront](#) > [Developer Guide](#)

Helper methods for origin modification

PDF

RSS

Focus mode

This section applies if you dynamically update or change the origin used on the request inside your CloudFront Functions code. You can update the origin on *viewer request* CloudFront Functions only. CloudFront Functions has a module that provides helper methods to dynamically update or change the origin.

To use this module, create a CloudFront function using JavaScript runtime 2.0 and include the following statement in the first line of the function code:

```
import cf from 'cloudfront';
```

For more information, see [JavaScript runtime 2.0 features for CloudFront Functions](#).



GA

Nov, 21th

pre-re:Invent

All Regions



[AWS Announcement](#)



Access VPC resources via AWS PrivateLink

AND INTEGRATE THEM DIRECTLY WITH AWS STEP FUNCTIONS OR AMAZON EVENTBRIDGE

You can now share AWS resources such as Amazon EC2 instances, Amazon ECS and Amazon EKS container services, and your own HTTPS services across Amazon VPC and AWS account boundaries and use them to build event-driven apps via Amazon EventBridge and orchestrate workflows with AWS Step Functions.

The screenshot displays the AWS Step Functions console configuration for a state named 'Call HTTPS APIs'. The state is of type 'Task: HTTP endpoint'. The configuration includes the following fields:

- State name:** Call HTTPS APIs
- State type:** Task: HTTP endpoint
- API Arguments:**
 - API endpoint:** https://service1. .com
 - Method:** GET
 - Authentication:** MyService1 (Status: Active)
- Request body - optional:** (This field is currently empty)

The left pane shows a workflow diagram starting with a 'Start' node, followed by the 'Call HTTPS APIs' state, then a 'Lambda Invoke' state, and finally an 'End' node.



GA

Dec, 1st

re:Invent

21 Regions



[AWS News Blog Post](#)

Other Announcements



1. [Built-in Amazon ECS support in Amazon VPC Lattice](#) (pre-re:Invent)
2. [Load Balancer Capacity Unit Reservation for ALBs / NLBs](#) (pre-re:Invent)
3. [UDP support for AWS PrivateLink and dual-stack NLBs](#) (pre-re:Invent)
4. [Amazon VPC Lattice adds TCP support for VPC Resources](#) (re:Invent)

Analytics

31 announcements (20 *pre-re:Invent*)



Express brokers for Amazon MSK

HIGH THROUGHPUT AND FASTER SCALING FOR AMAZON MANAGED STREAMING FOR APACHE KAFKA (MSK)

Express Broker is a new broker type for Amazon MSK that is designed to improve following capabilities in comparison to *Standard Brokers*:

1. Reduce to minimum storage operations.
2. Deliver 3x more throughput per broker.
3. Scale up to 20x faster.
4. Reduce recovery time by 90%.

In terms of billing, you pay hourly rate (with one-second resolution) varying on the size of the instance and per-GB rate for data written to the broken (with per-byte resolution).

General cluster properties

Cluster type
Choose the type of cluster you want to create. [Learn more](#)

Serverless
This option provides on-demand capacity that scales automatically as your application I/O scales.

Provisioned
This option allows you to specify the number of brokers and the amount of storage per broker.

Apache Kafka version
Select the Apache Kafka version that you want to run on the brokers. [Learn more](#)

3.6.0 (recommended) [View version compatibility](#)

Metadata mode - *new*
ZooKeeper

Broker type - *new*
Select a broker type for your MSK Provisioned cluster. You can't switch between Express and Standard broker types after you create the cluster. [Learn more](#)

Express brokers - *new* ←
Express brokers provide up to 3 times more throughput per broker, scale up to 20 times faster, and offer fully managed virtual storage capacity.

Standard brokers
Typical brokers for all versions of Apache Kafka.

Broker size
The broker size determines the broker's compute, memory, and storage capacities.

express.m7g.large
1000 recommended max partition count
vCPU : 2 Memory (GiB): 8 Network bandwidth (Gbps): Up to 12.5

Storage
Cluster storage
MSK-managed



GA

Nov, 7th
pre-re:Invent
9 Regions

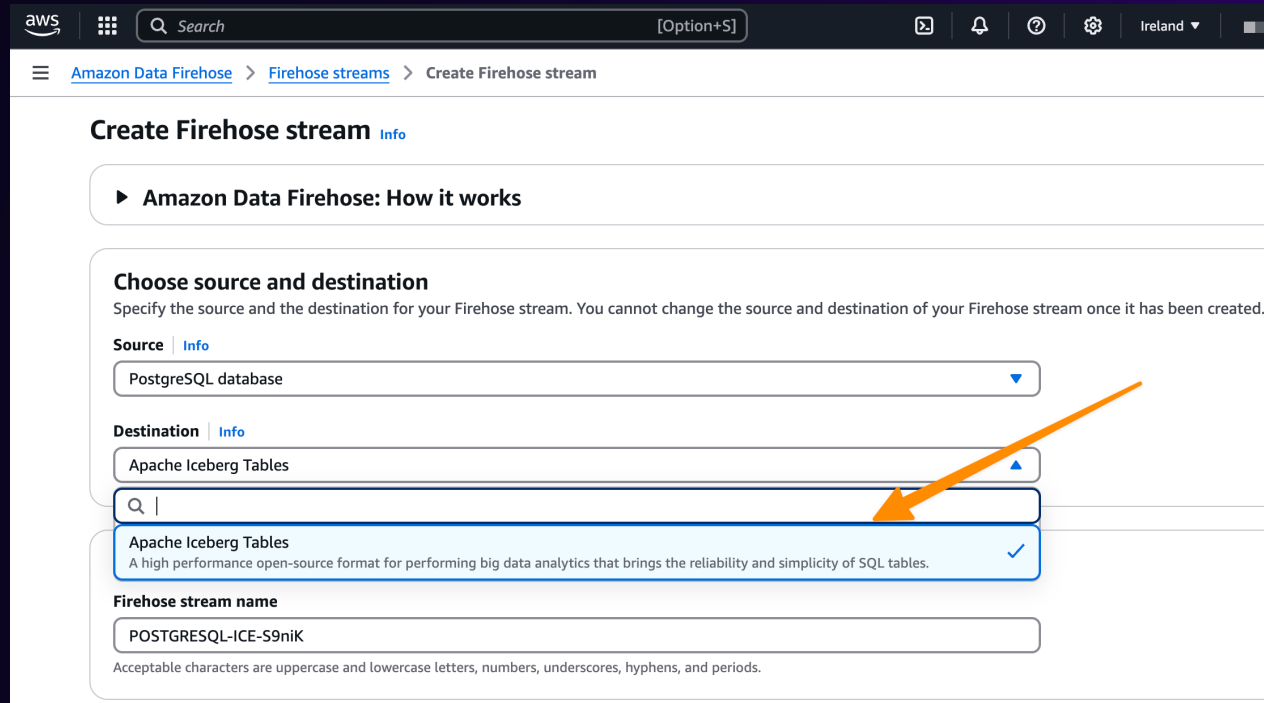


[AWS News Blog Post](#)

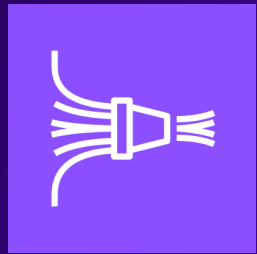
Database CDC with Amazon Data Firehose

CHANGE DATA CAPTURE (CDC) REPLICATION FROM MYSQL OR POSTGRESQL DATABASES TO APACHE ICEBERG TABLES STORED ON AMAZON S3

Amazon Data Firehose from now on supports MySQL and PostgreSQL relational databases as source, and Apache Iceberg tables (stored in Amazon S3 buckets and schema definition inside AWS Glue Data Catalog) as a destination for replicating change data capture streams with continuous schema evolution.



During *public preview*, there are no charges for your usage at the beginning of the preview. At some point in the future, it will be priced based on actual usage (e.g., bytes read and delivered).



Nov, 15th
pre-re:Invent
All Regions



[AWS News Blog Post](#)

Throughput increase for Amazon KDS

WRITING UP TO 10 GB/S FOR AMAZON KINESIS DATA STREAMS (KDS) ON-DEMAND MODE

Amazon Kinesis Data Streams On-Demand Mode now automatically scales to support streaming applications that write up to 10 GB/s per stream and consumers that read up to 20 GB/s per stream.

This is a **5x increase** from the previously supported limits.



GA

Nov, 20th

pre-re:Invent

3 Regions



[AWS Announcement](#)

AI upgrades for Apache Spark in AWS Glue

REDUCE EFFORT REQUIRED TO UPGRADE SPARK APPLICATIONS FROM WEEKS TO MINUTES

AWS Glue supports (in preview) AI upgrades for Apache Spark from an older version of AWS Glue to 4.0, which reduces the time data engineers spend on modernizing their Spark applications, allowing them to focus on building valuable analytics faster.



PP

Nov, 22nd

pre-re:Invent

5 Regions

The screenshot shows the AWS Glue Studio interface. At the top, there are three options for creating a job: 'Visual ETL', 'Notebook', and 'Script editor'. Below this is the 'Your jobs' section, which contains a table of jobs. The table has columns for 'Job name', 'Type', 'Created by', and 'Last modified'. Three jobs are listed: 'glue2-migrate-demo', 'glue2-migration', and 'glue2-count'. The 'glue2-count' job is selected. An 'Actions' dropdown menu is open over the 'glue2-count' job, showing options like 'Edit job', 'Clone job', 'Schedule job', 'Delete job(s)', 'Reset job bookmark', and 'Run upgrade analysis with AI - preview'. An orange arrow points to the 'Run upgrade analysis with AI - preview' option.

Job name	Type	Created by	Last modified
glue2-migrate-demo	Glue ETL	Script	11/6/2024, 2:14:49 AM
glue2-migration	Glue ETL	Script	11/4/2024, 9:56:12 PM
<input checked="" type="checkbox"/> glue2-count	Glue ETL	Script	11/1/2024, 12:51:39 AM



[AWS Big Data Blog Post](#)



AWS Clean Rooms Updates

SUPPORT FOR MULTIPLE CLOUDS, DATA SOURCES, AND PRIVACY ENHANCED MODEL TRAINING AND INFERENCE

Prior to re:Invent, AWS Clean Rooms ML received new capability called **custom modeling**, that enables companies and their partners on training and inferencing machine learning models on collective datasets without sharing sensitive data or proprietary models.

During AWS re:Invent, we have announced that from now you can connect data to AWS Clean Rooms also from Amazon Athena and Snowflake.



GA

Nov, 7th

Dec, 1st

pre-re:Invent

re:Invent

All Regions

AWS Clean Rooms > Tables > Configure new table

Configure new table [info](#)

Choose the data source and the data source type, then configure it for use in AWS Clean Rooms. After it is configured, you can associate it to multiple collaborations.

Data source [info](#)

AWS data sources

Amazon S3
Supported analytics engines
Spark, AWS Clean Rooms SQL

Amazon Athena
Supported analytics engines
Spark

Third-party clouds and data sources

Snowflake
Supported analytics engines
Spark

Data source credentials [info](#)

To use the selected data source, you must first create an AWS Secrets Manager secret to store your credentials.

Use existing secret ARN
Enter an existing secret ARN to link to the selected data source.

Store a new secret for this table
Enter your credentials to link to the selected data source.

Secret ARN
Store your credentials in [AWS Secrets Manager](#) and enter your secret ARN or import your secret ARN from an existing table.

Character count: 84 of 256

[Go to AWS Secrets Manager](#) [Import secret ARN from existing table](#)



[AWS News Blog Post](#)



Next generation of Amazon SageMaker

CENTER FOR ALL YOUR DATA, ANALYTICS, AI, AND GOVERNANCE - VIA UNIFIED STUDIO



PP

Dec, 3rd

re:Invent

5 Regions

A screenshot of the Amazon SageMaker Unified Studio interface. The top navigation bar includes 'Preview', 'Discover', 'Build', and 'Govern' menus, along with a 'Select a project' dropdown and utility icons. The main content area is divided into two sections: 'Discover' and 'Build'. The 'Discover' section features three cards: 'Data catalog' (Find and access data assets and ML models), 'Generative AI playground' (Experiment with chat, image, and video models), and 'Shared generative AI assets' (Explore apps and prompts shared with you). The 'Build' section has two columns of cards. The left column includes 'ML and generative AI model development' (Build, train, and deploy ML and foundation models with fully managed infrastructure, tools, and workflows) and 'Data processing and SQL analytics' (Analyze, prepare, and integrate data for analytics and AI using Amazon Athena, Amazon EMR, AWS Glue, and Amazon Redshift). The right column includes 'Generative AI app development' (Build generative AI apps and experiment with foundation models, prompts, agents, functions, and guardrails in Amazon Bedrock IDE) and 'Data and AI governance' (Publish your data products to the catalog with glossaries and metadata forms. Govern access securely in the Amazon SageMaker Catalog built on Amazon DataZone). Each card in the 'Build' section has a 'Create' button with a dropdown arrow.



[AWS News Blog Post](#)



Amazon SageMaker Lakehouse

UNIFIED, OPEN, AND SECURE DATA LAKEHOUSE ACROSS AMAZON S3 AND AMAZON REDSHIFT



GA

Dec, 3rd

re:Invent

14 Regions

The screenshot displays the Amazon SageMaker Lakehouse Query Editor interface. The main window shows a SQL editor with three queries:

```
1 create schema "dev@test-rms-catalog".salesdb;
2

1 CREATE TABLE salesdb.store_sales ( sale_id INTEGER IDENTITY(1,1) PRIMARY KE
2

1 INSERT INTO salesdb.store_sales (customer_sk, sale_date, sale_amount, product_name, last_purchase_date) VALUES (13251813, '2023-01-15', 150.00, 'Widget A', '2023-01-15'), (29033279, '2023-01-
```

A modal window is open, showing a list of connections, databases, and schemas. The connections list includes Athena (Lakehouse), Redshift (Lakehouse), and Redshift (project.redshift). The databases list includes dev, awsdatacatalog, and dev@test-rms-catalog-654654... The schemas list includes public.



[AWS News Blog Post](#)



AWS Glue 5.0

SUPPORT FOR AMAZON SAGEMAKER UNIFIED STUDIO, LAKEHOUSE, AND SPARK 3.5.2

In this version of AWS Glue you will get improved performance, enhanced security, support for Amazon Sagemaker Unified Studio, Amazon SageMaker Lakehouse, and following runtime and dependencies updates:

- Apache Spark 3.5.2
- Python 3.11
- Java 17
- Apache Hudi 0.15.0
- Apache Iceberg 1.6.1
- Delta Lake 3.2.0

The screenshot shows the AWS Glue console interface. On the left is a navigation sidebar with categories like 'AWS Glue', 'ETL jobs', 'Data Catalog', and 'Legacy pages'. The main content area displays the 'Job details' for a job named 'TestJobFor5.0'. The 'Basic properties' section includes fields for Name, Description, IAM Role, Type, Glue version, and Language. The 'Glue version' dropdown is currently set to 'Glue 5.0 - Supports spark 3.5, Scala 2, Python 3'. An orange arrow points from the 'Enable new navigation' option in the sidebar to this dropdown.



GA

Dec, 3rd

re:Invent

14 Regions



[AWS Announcement](#)

Scenario analysis in QuickSight with Q

AI-ASSISTED DATA ANALYSIS EXPERIENCE THAT HELPS YOU MAKE BETTER DECISIONS, FASTER



PP

Dec, 4th

re:Invent

2 Regions

Amazon Q in QuickSight received a new capability, that helps users perform scenario analyses to find answers to complex problems quickly.

This AI-assisted data analysis experience helps business users find answers to complex problems by guiding them step-by-step through in-depth data analysis - suggesting analytical approaches, automatically analyzing data, and summarizing findings with suggested actions - all of the above using natural language prompts.

Data to insights
Describe what you want to explore, and we will find relevant data for you

Which products should we focus our marketing efforts on in order to maximize revenue in Japan?

ADD DATA EDIT

Which marketing campaigns have been most effective for increasing sales in Japan?

Which marketing campaigns have been most effective for increasing sales in Japan?

Step 1: We'll start by focusing on the sales data from Japan and organizing it by marketing campaign to get an overview of each campaign's performance.

In Japan, campaign GC22 generated the highest sales (\$273) with 39 units sold, followed by VEG903 with \$232 in sales and 29 units, while GC33 showed the highest average profit per order at \$53.4.

Campaign_ID	Sales	Quantity	Profit	Number_of_Orders
GC22	273	39	35.33	6
VEG903	232	29	29	5
BB120	162	29	14.78	9
GC11	161	23	37.38	4
GRA27	150	30	10.91	11
BB121	124	23	34.57	3
GRA24	120	24	8	12
GC33	112	16	53.4	2



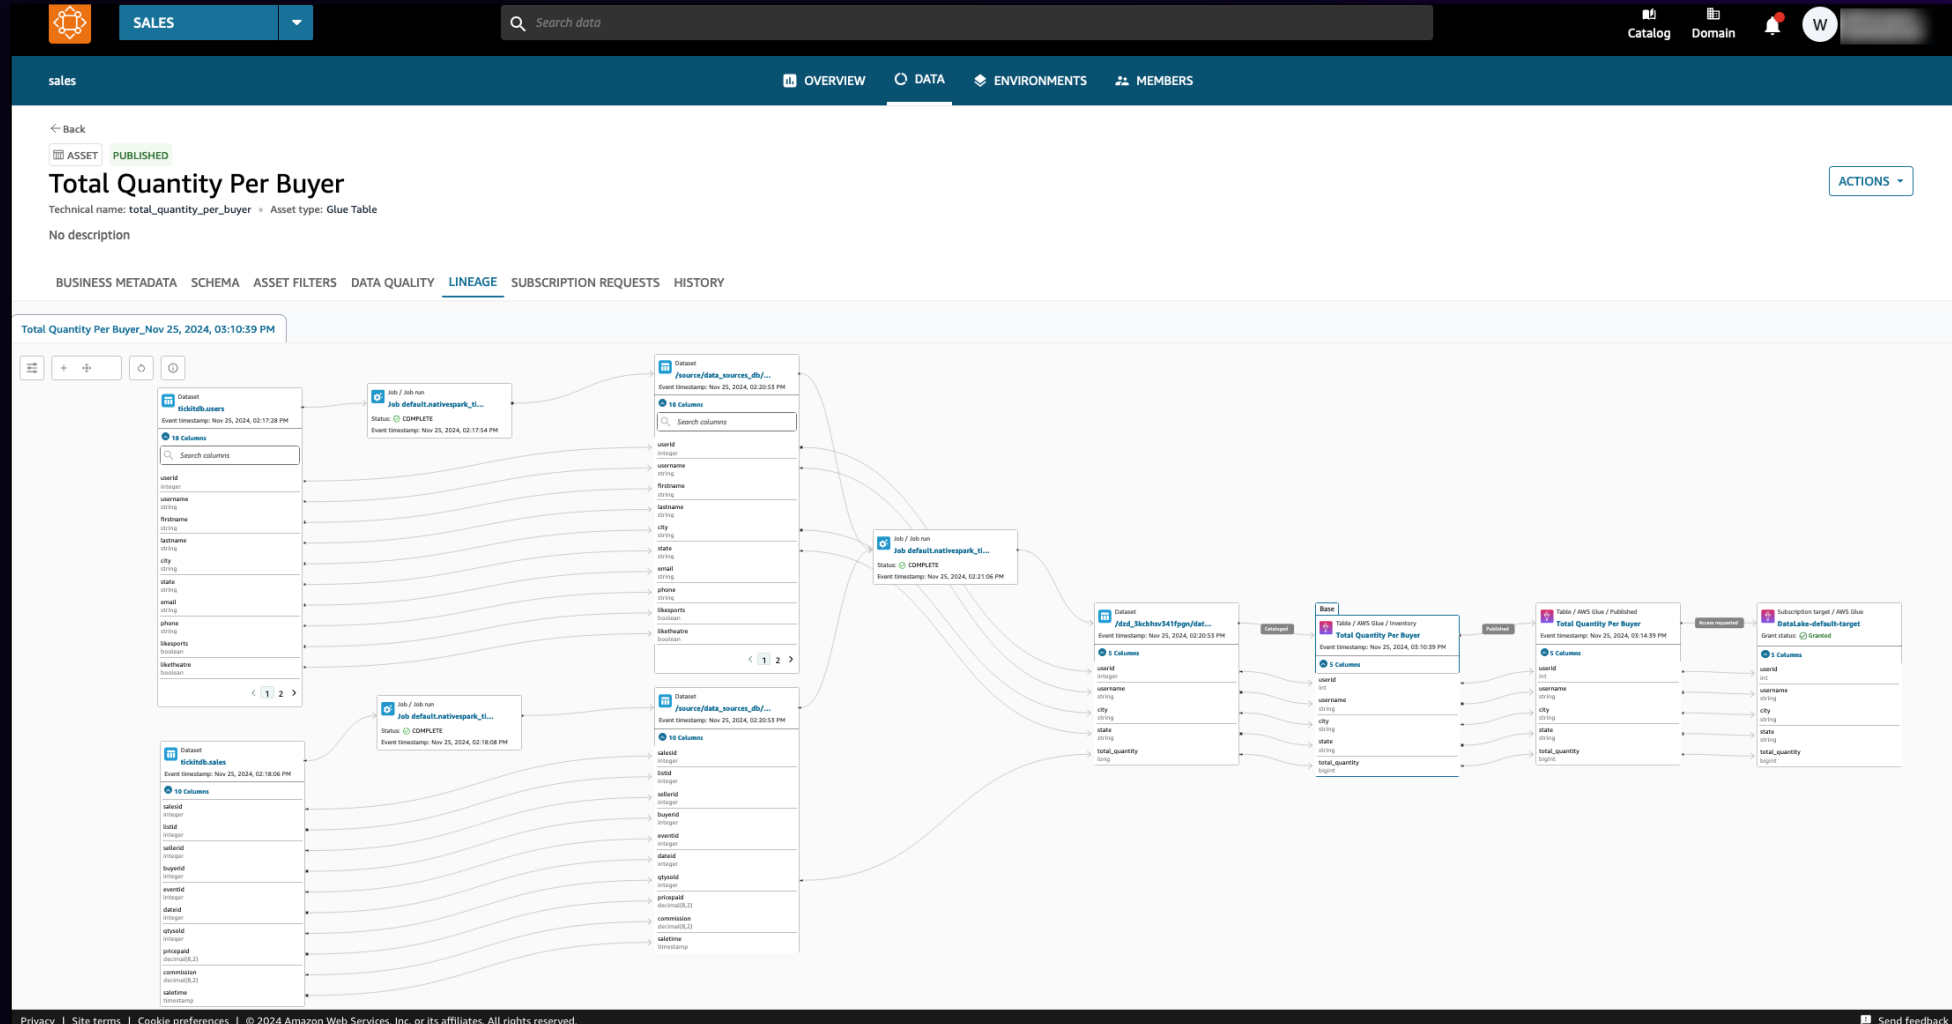
[AWS News Blog Post](#)



Data Lineage tracking across AWS tools



TRACK DATA LINEAGE WITH AMAZON DATAZONE AND AMAZON SAGEMAKER UNIFIED STUDIO



GA

Dec, 4th

re:Invent

All Regions

Open Source



[AWS News Blog Post](#)



Other Announcements (1/2)



1. **Amazon QuickSight updates:**
 - a) Image component (pre-re:Invent)
 - b) Highcharts visual (pre-re:Invent, **PP**)
 - c) Import visual capability from an existing dashboard or analysis (pre-re:Invent, **PP**)
 - d) Layer Map component (pre-re:Invent)
 - e) Amazon Q in QuickSight unifies insights from structured and unstructured data (re:Invent)
2. Amazon DataZone expands access with tools like Tableau, Power BI, and via JDBC driver (pre-re:Invent)
3. Smaller environment size for Amazon MWAA (pre-re:Invent)
4. **Other AWS Glue updates:**
 - a) Scheduled generation of column level statistics in AWS Glue Data Catalog (pre-re:Invent)
 - b) Automatic Apache Iceberg tables optimization through Amazon VPC in AWS Glue Data Catalog (pre-re:Invent)
 - c) Automation statistics generation for new tables in AWS Glue Data Catalog (re:Invent)
 - d) Amazon S3 Access Grants now integrate with AWS Glue (re:Invent)
 - e) Expanded data connectivity for Amazon SageMaker Lakehouse and AWS Glue (re:Invent)
5. **Amazon Managed Service for Apache Flink (MSF) updates:**
 - a) Generating vector embeddings from Amazon MSK with Amazon Bedrock (pre-re:Invent)
 - b) Adding support for Amazon DynamoDB Streams as a source (pre-re:Invent)
 - c) Adding support for Amazon Managed Service for Prometheus as destination (pre-re:Invent)

Other Announcements (2/2)



6. **Amazon Redshift Query Editor updates:**
 - a) **Amazon Q generative SQL is now available in additional regions** (pre-re:Invent)
 - b) **Increasing maximum result set and export size to 100 MB** (pre-re:Invent)
7. **Advanced Scaling in Amazon EMR Managed Scaling** (pre-re:Invent)
8. **AWS Lambda custom data transformation in Amazon OpenSearch Ingestion** (pre-re:Invent)
9. **Amazon DynamoDB zero-ETL integration with Amazon SageMaker Lakehouse** (re:Invent)

Machine Learning, AI, and Generative AI

46 announcements (17 *pre-re:Invent*)



Prompt Optimization in Amazon Bedrock

REWRITING PROMPTS FOR HIGHER QUALITY RESPONSES FROM FOUNDATIONAL MODELS

Amazon Bedrock now offers a tool to optimize prompts. Optimization rewrites prompts to yield inference results that are more suitable for your use case. You can choose the model that you want to optimize the prompt for and then generate a revised prompt.

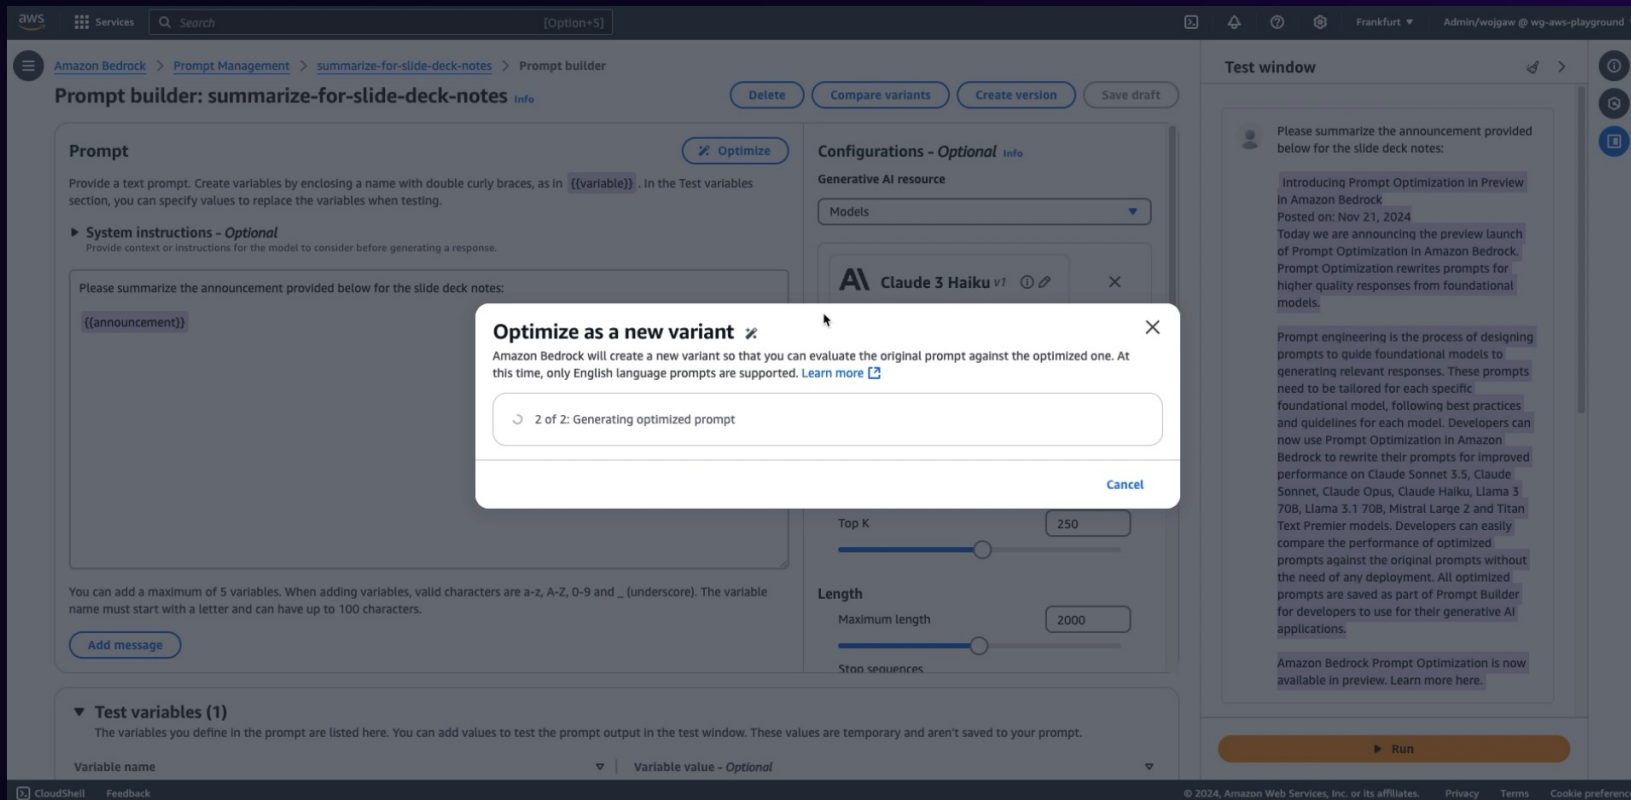


PP

Nov, 21st

pre-re:Invent

All Regions
Limited Models



[AWS Announcement](#)



Amazon Bedrock Flows

GENERAL AVAILABILITY OF PREVIOUSLY KNOWN AS PROMPT FLOWS WITH 2 NEW CAPABILITIES



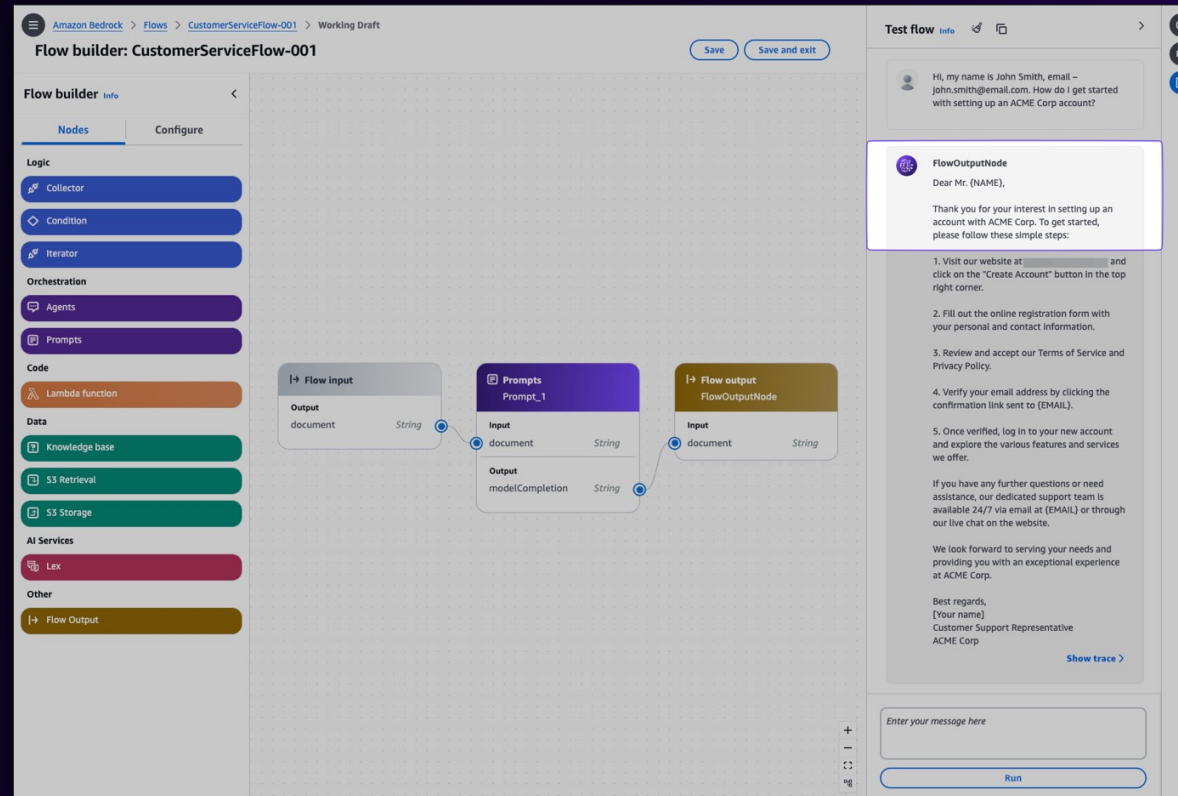
GA

Visual builder in Amazon Bedrock, previously known as Amazon Bedrock Prompt Flows, received an updated name:

Amazon Bedrock Flows

and two new capabilities:

1. **Enhanced Security** by adding Amazon Bedrock Guardrails to specific nodes.
2. **Enhanced Traceability** with validation and debugging capabilities.



Nov, 22nd

pre-re:Invent

All Regions



[AWS ML Blog Post](#)



Scale Down to Zero for Amazon SageMaker

NEW CAPABILITY OF SAGEMAKER INFERENCE THAT ALLOWS TO SCALE DOWN DURING INACTIVITY



GA

Nov, 25th

pre-re:Invent

All Regions

Amazon SageMaker Inference that allows endpoints to scale to zero instances during periods of inactivity.

This feature can significantly reduce costs for running inference using AI models, making it particularly beneficial for applications with variable traffic patterns such as chatbots, content moderation systems, and other generative AI use cases.



[AWS Announcement](#)

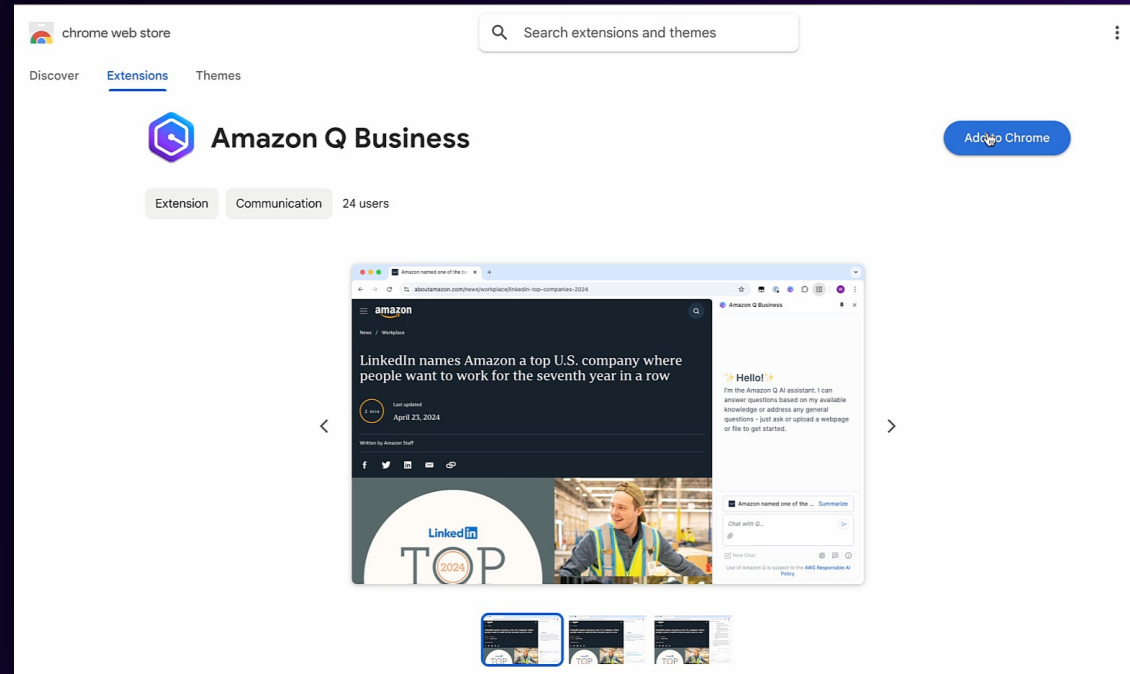


Amazon Q Business at pre-re:Invent

BROWSER EXTENSION, PRIVATE SHARING FOR Q APPS, CONNECTORS AND INTEGRATIONS

Period of pre-re:Invent was also really busy for the Amazon Q Business service team, as they announced 9 different announcements.

They were mostly around new connectors and integrations, releasing browser extension, private sharing in Q Apps, data collection in Q Apps (preview), and user experience (UX) improvements in terms of reusing files and context in the conversations, and analyzing multi-modal data like tables and visual documents.



Nov, 2024

pre-re:Invent

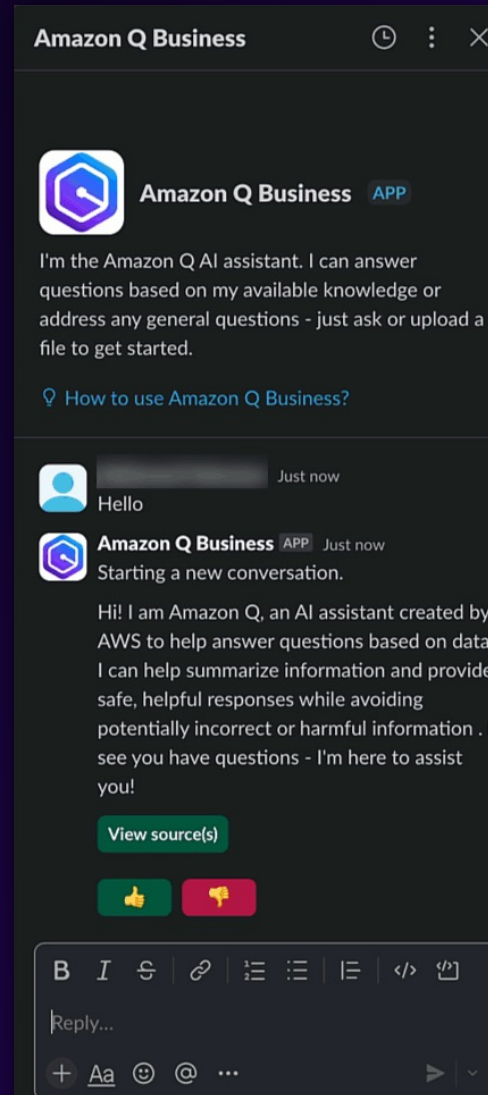
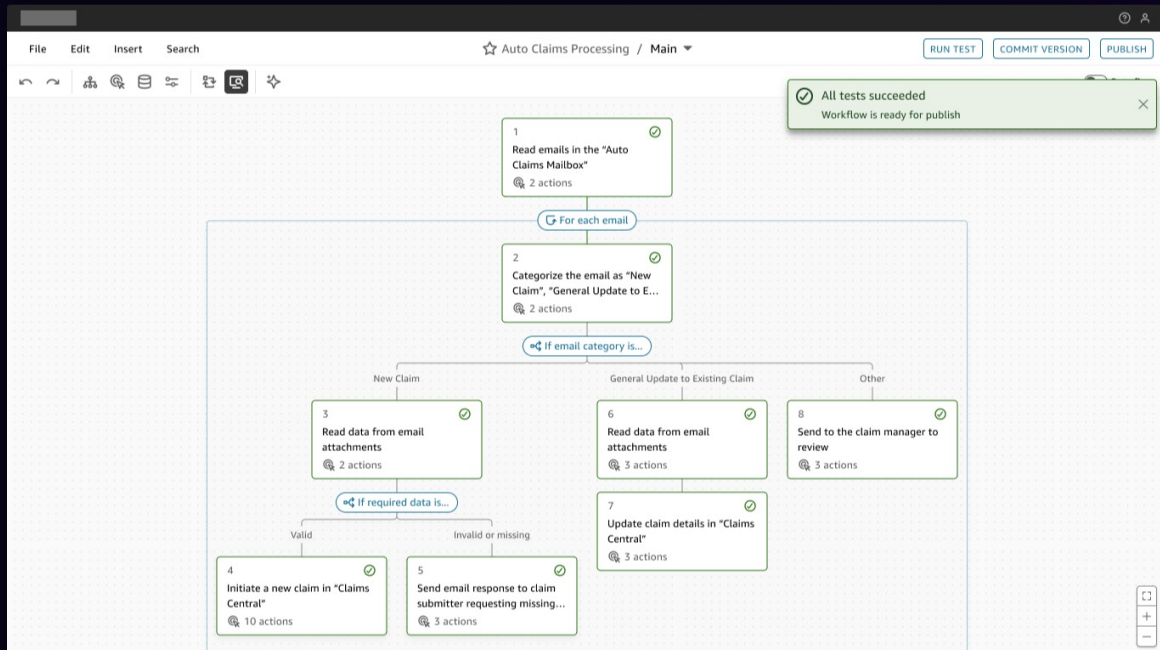
All Regions

Amazon Q Business at re:Invent

50+ ACTION INTEGRATIONS, TOOLS, WORKFLOW AUTOMATION, AMAZON Q INDEX

However, service team did not stop there! Throughout the re:Invent they've announced 4 additional capabilities

Starting from adding 50+ actions, through integrations with tools, **workflow automation capability (coming soon)**, and ability to use Q Index by Independent Software Vendor (ISV) partners for driving future integrations and new use cases.



GA

CS

Dec, 1st
Dec, 2nd
Dec, 3rd

re:Invent

All Regions



[AWS News Blog Post](#)



Other updates for Amazon Bedrock

RAG EVALUATION, LLM-AS-A-JUDGE, RERANKING API, CUSTOM CONNECTORS AND STREAMING



GA

PP

Dec, 1st

re:Invent

All Regions

With the plethora of updates for Amazon Bedrock announced in two keynotes, there are a few others announced during re:Invent that are worth your attention:

- Support for **custom connectors** and **ingestion of streaming data** in Amazon Bedrock Knowledge Bases.
- Support for **reranking models** for Bedrock Knowledge Bases.
- Bedrock Knowledge Bases now supports **RAG evaluation (preview)**.
- Bedrock Model Evaluation now includes **LLM-as-a-judge (preview)**.



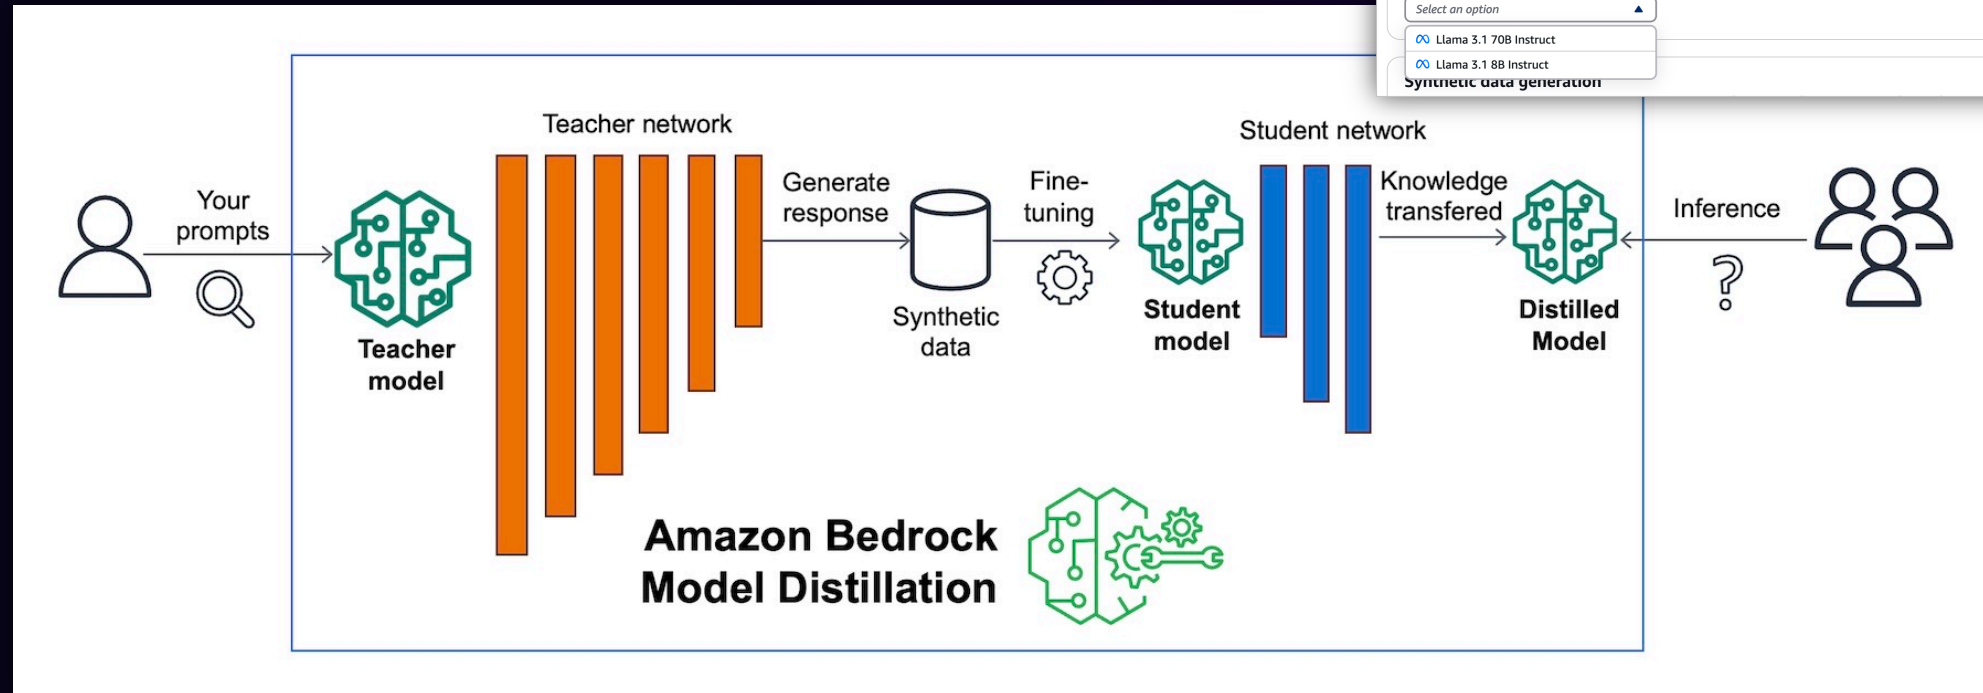
[AWS News Blog Post](#)



Amazon Bedrock Model Distillation

FINE-TUNING SMALLER MODEL FROM LARGER MODEL WITH DATA SYNTHESIS

Amazon Bedrock Model Distillation is a capability, that automates the process of creating a distilled model for your specific use case by generating responses from a large foundation model (a teacher) model and fine-tunes a smaller FM (a student model) with the generated responses via data synthesis.



Teacher model - Student model details [Info](#)
Choose the teacher model and based on your choice of the teacher model, then select a student model from the list of available student models.

Teacher model
Choose from a list of teacher models that you wish to use for model distillation. The teacher model will be used to generate synthetic data.

[Llama 3.1 405B Instruct v1](#)
[Change](#)

Student model
Select a student model. Student models are based on the same family as the teacher model.

Select a student model

Select an option

- [Llama 3.1 70B Instruct](#)
- [Llama 3.1 8B Instruct](#)

Synthetic data generation



PP

Dec, 3rd

re:Invent

2 Regions



Automated Reasoning in Bedrock Guardrails

EMPLOY SOUND MATHEMATICAL TECHNIQUES TO SAFEGUARD AMAZON BEDROCK GUARDRAILS



PP

Dec, 3rd

re:Invent

2 Regions

AWS becomes the first and only major cloud provider to integrate **automated reasoning** in our generative AI offerings.

Automated Reasoning checks in Amazon Bedrock Guardrails are the first and only generative AI safeguard that helps prevent factual errors due to hallucinations using logically accurate and verifiable reasoning that explains why generative AI responses are correct.

Amazon Bedrock > Automated Reasoning > Antje-TicketPolicy-123 > Version 1

Antje-TicketPolicy-123 Version 1 - latest Save changes

Definitions | Test playground | Details

Rules (6) Add rule

Rule	Updates	Actions
A name change is allowed if and only if changes are allowed and the submission is valid.	-	⋮
A submission is valid if and only if the submission method is email.	-	⋮
Change is allowed if and only if the airline allows it and the time limit is met.	-	⋮
Name correction is allowed if and only if changes are allowed and the submission is valid.	-	⋮
The airline is allowed if and only if the allowed airlines are either Unicorn Airlines or Cloud Nine Air.	-	⋮
The time limit is met if and only if the hours since purchase is less than or equal to 24.	-	⋮

Variables (8) Add variable

Name	Type	Description	Updates	Actions
allowed_airlines	Enumeration	The airlines that allow changes under the Soar Customer Fir...	-	⋮
allowed_changes	Enumeration	The types of changes and corrections allowed under the policy	-	⋮
change_method	Enumeration	The method by which name changes or corrections must be s...	-	⋮
change_time_limit	Integer	The time limit within which changes must be made	-	⋮
hours_since_purchase	Integer	The time limit within which name changes or corrections mu...	-	⋮
name_change_allowed	Boolean	The process of changing a passenger's name on a ticket	-	⋮
name_correction_allowed	Boolean	The process of correcting a misspelled name on a ticket	-	⋮
submission_method	Enumeration	The required method for submitting change requests	-	⋮



[AWS News Blog Post](#)



Multi-agent collaboration in Bedrock

MULTIPLE AI AGENTS WORKING TOGETHER ON A COMPLEX MULTI-STEP TASKS



PP

Dec, 3rd

re:Invent

3 Regions

When you need to handle a complex task, there is often a need to use more than just one agent. You can create multiple specialized ones – however, managing, coordination, memory and session handling are tedious.

That's why Amazon Bedrock offers multi-agent collaboration capability, which addresses all of the above challenges in a fully-managed way.

Test TestAlias: Working draft Info
Using ODT Change

Create a 2-week social campaign for EcoTech's new solar panel launch. Target: B2B (facility managers, sustainability directors) Key points: 30% more efficient, AI-optimized, 2-year ROI Need: 4 posts/week on LinkedIn/Twitter (40% educational, 30% product, 30% thought leadership).

Here's your 2-week social campaign plan for EcoTech's solar panel launch:
Week 1: Tuesday 10am EST - LinkedIn: Case study showcasing manufacturing company's 45% energy cost reduction
Wednesday 10am EST - Twitter: Educational post about AI-optimized efficiency gains
Thursday 1pm EST - LinkedIn: Thought leadership piece on AI-driven facility sustainability
Friday 9am EST - Twitter: Product feature highlighting ROI and efficiency benefits
Week 2: Tuesday 11am EST - LinkedIn: How-to guide on maximizing facility sustainability
Wednesday 10am EST - Twitter: Industry report teaser with statistics
Thursday 2pm EST - LinkedIn: Social proof post featuring sustainability

Enter your message here

Multi-agent collaboration trace timeline

For agents with Multi-agent collaboration enabled, each subsequently invoked Agent collaborator is displayed. Each bar within the trace timeline can be clicked to show the specific breakdown for that Agent collaborator.

Agent Collaborator	Duration
(1)social-media-campaign-manager	49.32s
(1.1)content-strategist	13.21s
(1.2)engagement-predictor	10.20s

Time (...) 0.00s 5.00s 10.00s 15.00s 20.00s 25.00s 30.00s 35.00s 40.00s

Trace Session summaries

Trace (10) Copy all
Request ID 77bb826a-73f4-482b-a55a-572d6a9fb518

< Trace **Orchestration and Knowledge Base** P >

Orchestration and Knowledge Base

After running a prompt in the test window, this orchestration trace allows you to explore the trace steps to understand the linear chain of thought used by the agent's orchestration prompt component. If a knowledge base was invoked, the trace also allows you to see how the results from the knowledge base were summarized to generate an observation that is used for for orchestrating the next step or generating the response.



[AWS News Blog Post](#)



Amazon Bedrock Marketplace

ACCESS OVER 100 FOUNDATION MODELS IN ONE PLACE

Amazon Bedrock Marketplace is a new capability that gives you access to over 100 popular, emerging, and specialized foundation models (FMs) through Amazon Bedrock.

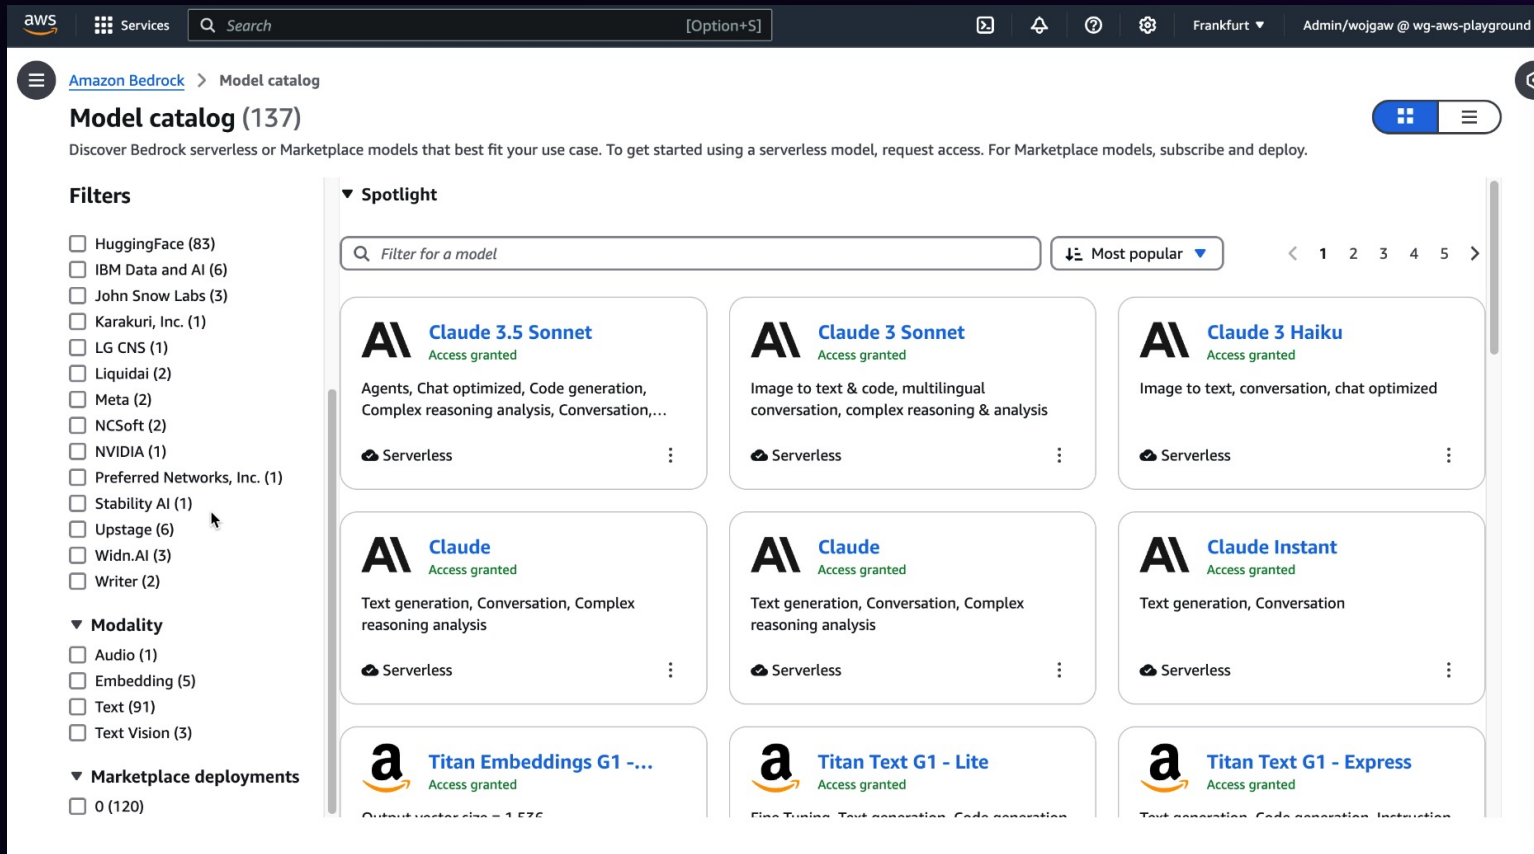


GA

Dec, 4th

re:Invent

14 Regions



Amazon Nova: New Foundation Models

FRONTIER INTELLIGENCE AND INDUSTRY LEADING PRICE PERFORMANCE

Amazon Nova is a new generation of state-of-the-art foundation models (FMs) family developed by Amazon, that deliver frontier intelligence and industry leading price performance and available exclusively in Amazon Bedrock.



GA

Dec, 3rd

re:Invent

3 Regions



Prompt routing and caching in Bedrock

PROMPT IMPROVEMENTS WITH INTELLIGENT ROUTING AND CACHING

Amazon Bedrock has introduced in preview two capabilities that help reduce costs and latency for generative AI applications: **intelligent prompt routing** within the same foundation model family and **prompt caching** that allows for reuse of frequently used context across model invocations.



PP

Dec, 4th


re:Invent

2 Regions
Limited Models

Prompt Routers Info
Efficiently routes requests between different foundational models from the same model family, optimizing for quality of responses and cost.


▼ **How it works**

Experiment with default routers



You can choose from two prompt routers that route requests between either Claude Sonnet 3.5 and Claude Haiku, or between Llama 3.1 70B and Llama 3.1 8B. Choose a prompt router to get started.

Use for production



You can save months of effort on testing different models and creating complex orchestration workflows. Use the Invoke and Converse API operations to incorporate routers into your generative AI applications.

Default routers (1/2) Open in playground

Bedrock offers pre-configured prompt routers for different model providers.

AI Anthropic Prompt Router ◯

Routing models
Claude 3 Haiku, Claude 3.5 Sonnet

Routing criteria
Response quality difference: 0%

Meta Prompt Router ◯

Routing models
Llama 3.1 8B Instruct, Llama 3.1 70B Instruct

Routing criteria
Response quality difference: 0%



Enhance processing and retrieval in Bedrock

STRUCTURED RETRIEVAL, GRAPHRAG, DATA AUTOMATION, AND MULTIMODAL KNOWLEDGE BASES



GA

PP

During Dr. Sivasubramanian's keynote, we have announced four enhancements to Amazon Bedrock that streamlines how you can analyze data within this service:

1. Amazon Bedrock Data Automation (*preview*) for fully managed insights generation from multimodal content.
2. Amazon Bedrock Knowledge Bases now processes multimodal data with either Bedrock Data Automation or another foundation model (FM) as a parser.
3. Amazon Bedrock Knowledge Bases now supports GraphRAG (*preview*) powered by Amazon Neptune Analytics.
4. Amazon Bedrock Knowledge Bases now supports structured data retrieval, that translates natural language to SQL and reaches out directly to data warehouses and data lakes.

The screenshot displays the Amazon Bedrock interface for document processing. On the left, a document preview shows a birth certificate from the City and County of San Francisco. The right-hand panel contains configuration options:

- Configuration info:** Includes buttons for "Regenerate results" and "Add to project".
- Modality detected:** Document.
- Granularity:** Options include Page level, Element level (selected), and Word level.
- Bounding boxes:** Option to Enable.
- Text Format:** Options include Plain Text, Text with Markdown (selected), Text with HTML, and CSV (tables only).
- Generative Fields:** Option to Enable.
- Output Format:** Options include JSON (selected), JSON+files (txt, md / html, csv for tables), and Return JSON and additional file formats based on text formatting controls.
- Results:** A list of extracted text blocks with confidence scores and markdown formatting. For example, "TEXT 1" has a confidence score of 98% and contains "# CERTIFICATION OF VITAL RECORD".

Dec, 4th

re:Invent

All Regions

1 Region



Multimodal toxicity detection in Guardrails

WITH IMAGE SUPPORT AND INDEPENDENT API

This new capability in Amazon Bedrock Guardrails detects and filters out undesirable image content in addition to text.

In addition to multimodality, Amazon Bedrock receives an update in the form of the independent API `ApplyGuardrail` which allows for evaluation of any inputs and outputs without invoking foundation models on Amazon Bedrock.

It means that you can apply guardrails, also for models hosted on Amazon SageMaker or regardless the underlying infrastructure.



PP

Dec, 4th

re:Invent

10 Regions

A screenshot of the Amazon Bedrock Guardrails console. The breadcrumb trail at the top reads "Amazon Bedrock > Guardrails > multimodal-toxicity-demo > Working Draft". The main area is titled "Test" and shows the configuration for "Claude 3.5 Sonnet v1 | On-demand". There is a checkbox for "Use Guardrails independent API" which is currently unchecked. The "Prompt" field contains the text "Describe this image." Below the prompt, there is a "View files (1)" link. The "Model response" field is empty. A "Guardrail action" section shows a warning icon and the text "Intervened (1 instances)", with a "View trace" button next to it. The "Final response" field contains the text "Sorry, the model cannot answer this question." To the right, the "Guardrail trace" section is active, showing a table with columns "Category", "Test result", and "Details". The table has one row: "Content filters" with a "Blocked" result. The "Details" column shows "Detected Insights", "Strength: High", and "Confidence: High". A red arrow points to the "Model response" tab, and another red arrow points to the "View trace" button.



[AWS News Blog Post](#)

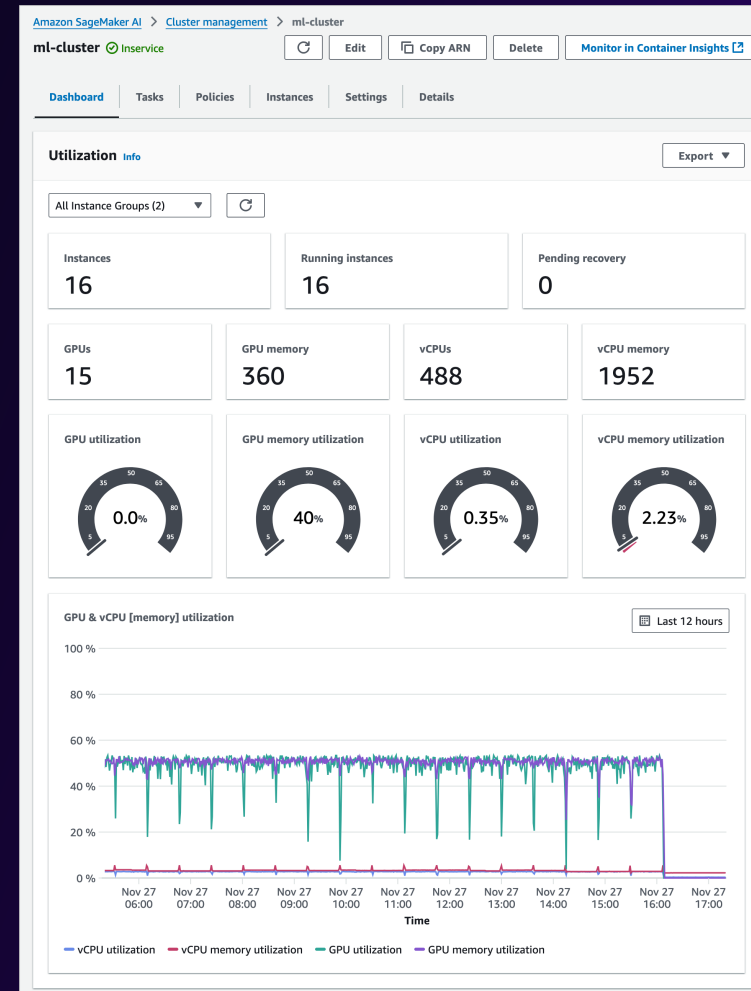


Updates for Amazon SageMaker HyperPod

FLEXIBLE TRAINING PLANS, TASK GOVERNANCE, AND RECIPES

This year, based on customer feedback, Amazon SageMaker HyperPod received 3 important updates:

1. **Flexible Training Plans**, to help data scientists train large foundation models (FMs) within their timelines and budgets and save them weeks of effort in managing the training process based on compute availability.
2. **Task Governance**, a new innovation to easily and centrally manage and maximize GPU and AWS Trainium utilization across generative AI model development tasks, such as training, fine-tuning, and inference.
3. **Recipes**, help data scientists and developers of all skill sets to get started training and fine-tuning foundation models in minutes with state-of-the-art performance.



GA

Dec, 4th

re:Invent

3 Regions
Limited Instances

12 Regions

All Regions



[AWS News Blog Post](#)



Other Announcements



1. [Fine-tuning Anthropic's Claude 3 Haiku in Amazon Bedrock](#) (pre-re:Invent)
2. Binary vector embeddings in [Amazon Bedrock Knowledge Bases](#) and [Amazon Titan Text Embeddings](#) (pre-re:Invent)
3. **Amazon Polly updates:**
 - a) [Six new synthetic generative voices \(English, French, Spanish, German\)](#) (pre-re:Invent)
 - b) [More synthetic generative voices \(English, French, Spanish, German, Italian\)](#) (pre-re:Invent)
4. [New AWS AI Service Cards for advancing responsible generative AI](#) (re:Invent)
5. [Amazon SageMaker Partner AI Apps](#) (re:Invent)
6. [Announcing GenAI Index in Amazon Kendra](#) (re:Invent)
7. [39 new or updated datasets available on the Registry of Open Data on AWS](#) (re:Invent)

App Integration and Serverless

12 announcements (12 *pre-re:Invent*)



AWS AppSync Events

SERVERLESS WEBSOCKET API TO POWER REAL-TIME WEB AND MOBILE EXPERIENCES AT SCALE



GA

Oct, 30th

pre-re:Invent

All Regions

What is AWS AppSync Events?

PDF

AWS AppSync Events lets you create secure and performant serverless WebSocket APIs that can broadcast real-time event data to millions of subscribers, without you having to manage connections or resource scaling. With AWS AppSync Events, there is no API code required to get started, so you can create production-ready real-time web and mobile experiences in minutes.

AWS AppSync Events further simplifies the management and scaling of real-time applications by shifting tasks like message transformation and broadcast, publish and subscribe authentication, and the creation of logs and metrics to AWS, while delivering reduced time to market, low latency, enhanced security, and lower total costs.

With Event APIs, you can enable the following network communication types.

- Unicast
- Multicast
- Broadcast
- Batch publishing and messaging

This allows you to build the following types of interactive and collaborative experiences.

- Live chat and messaging
- Sports and score updates
- Real-time in-app alerts and notifications
- Live commenting and activity feeds

AWS AppSync Events simplifies real-time application development by providing the following features.

- Automatic management of WebSocket connections and scaling
- Built-in support for broadcasting events to large numbers of subscribers
- Flexible event filtering and transformation capabilities
- Fine-grained authentication and authorization
- Seamless integration with other AWS services and external systems for event-driven architectures



[AWS FEWAM Blog Post](#)



Amazon EventBridge Latency Reduction

94% END-TO-END LATENCY IMPROVEMENT FOR EVENT BUSES AT NO ADDITIONAL COST

Amazon EventBridge Event Buses announces up to **94% improvement** in end-to-end latency for Event Buses, enabling you to handle highly latency sensitive applications.

You can now detect and respond to critical events more quickly, enabling rapid innovation, faster decision-making, and improved operational efficiency.



GA

Nov, 12th

pre-re:Invent

All Regions



[AWS Announcement](#)



AWS Lambda SnapStart for .NET and Python

REDUCE COLD-STARTS FOR .NET CORE AND PYTHON SERVERLESS APPLICATIONS



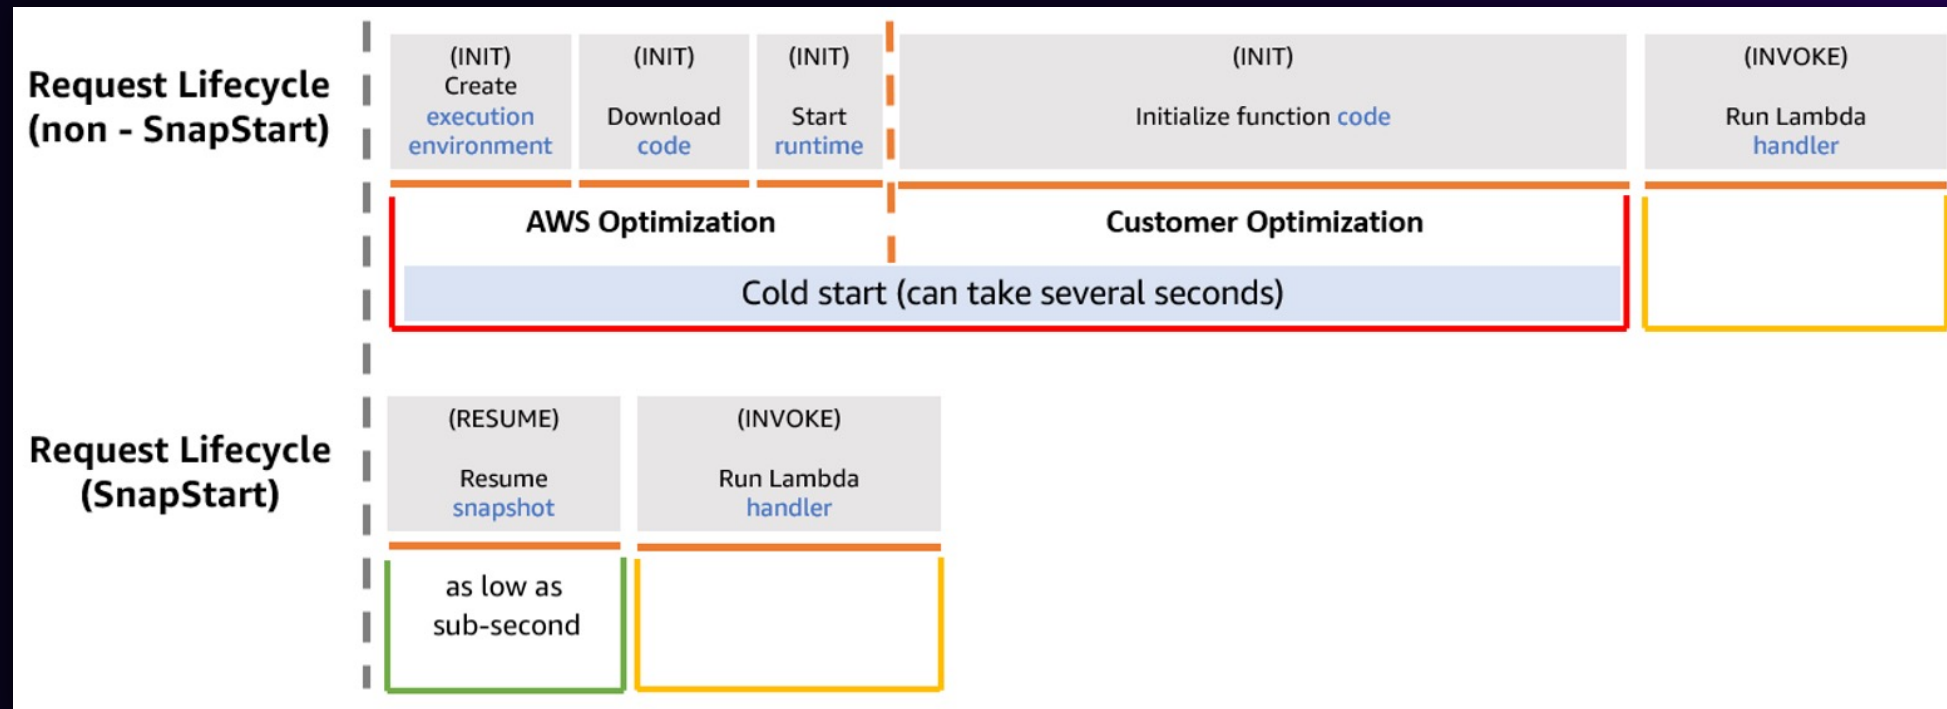
GA

Nov, 18th

pre-re:Invent

9 Regions

In addition to Java, from now you can use AWS Lambda SnapStart with your functions that use the **Python** and **.NET** managed runtimes, to deliver as low as sub-second startup performance.



[AWS News Blog Post](#)



Event Source Mappings CloudWatch Metrics



INCREASED OBSERVABILITY INTO EVENT SOURCE MAPPINGS (ESM) IN AWS LAMBDA

AWS Lambda team announced new Amazon CloudWatch metrics for Event Source Mappings (ESMs), which provide customers visibility into the processing state of events with Amazon SQS, Amazon Kinesis, and Amazon DynamoDB event source integrations.

GA

Nov, 21st

pre-re:Invent

All Regions

Event source mapping metrics

Event source mapping metrics provide insights into the processing behavior of your event source mapping. These metrics help you monitor the flow and status of events, including events that your event source mapping successfully processed, filtered, or dropped.

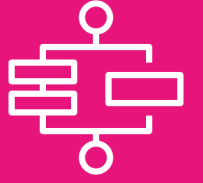
You must opt-in to receive metrics related to counts (`PolledEventCount`, `FilteredOutEventCount`, `InvokedEventCount`, `FailedInvokeEventCount`, `DroppedEventCount`, `OnFailureDestinationDeliveredEventCount`, and `DeletedEventCount`). To opt-in, you can use the console or the Lambda API.



[AWS Announcement](#)

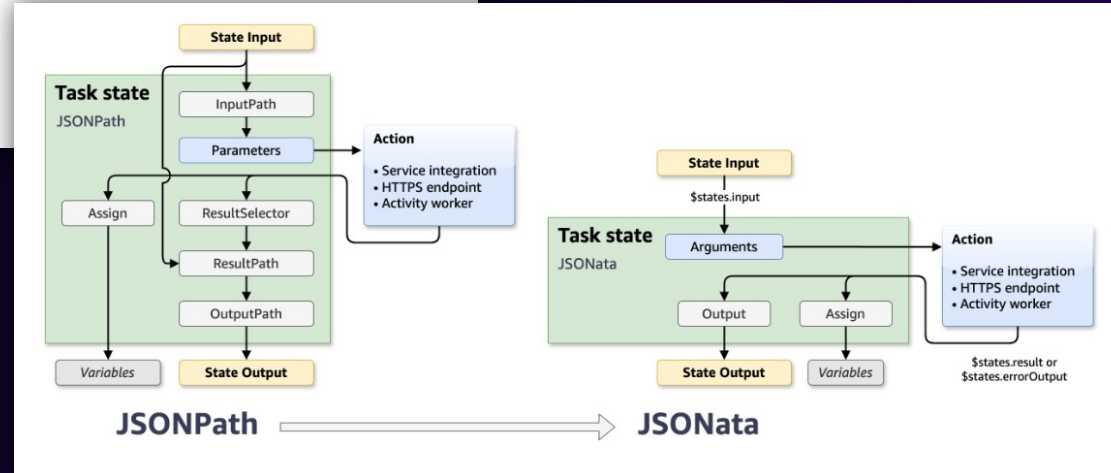
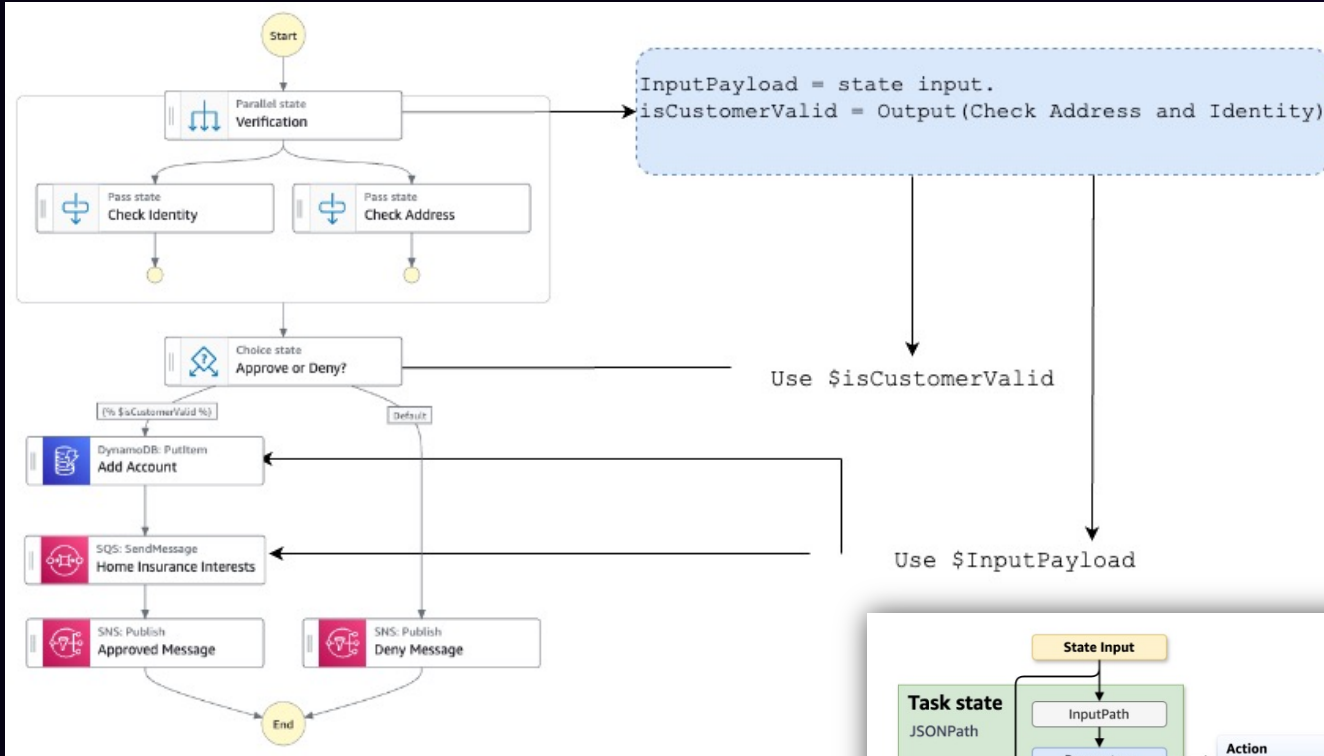
AWS Step Functions: JSONata and Variables

SIMPLIFY STATE PAYLOAD MANAGEMENT AND DATA TRANSFORMATION IN STATE MACHINES



GA

Nov, 22nd
pre-re:Invent
All Regions



[AWS Compute Blog Post](#)



Provisioned Mode for Event Source Mapping

PROVISIONING EVENT POLLING RESOURCES FOR EVENT SOURCE MAPPING (ESM) IN AWS LAMBDA



GA

Nov, 22nd

pre-re:Invent

9 Regions

AWS Lambda announces Provisioned Mode for event source mappings (ESMs) that subscribe to Apache Kafka event sources.

This feature that allows you to optimize the throughput of your Kafka ESM by provisioning event polling resources that remain ready to handle sudden spikes in traffic.

aws [Search] [Option+S] N. Virginia Admin/wojgaw @ wg-aws-playground

Lambda > Add triggers

Secrets Manager key
Choose the Secrets Manager key that contains the authentication token required to access your trigger.

Q

Encryption
Choose the secret key containing the root CA certificate used by your Kafka brokers for TLS encryption.

Required if your Kafka brokers use certificates signed by a private CA.

Configure provisioned mode - new
Select to configure provisioned mode for your event source mapping. You can configure the minimum event pollers, the maximum event pollers, or both. For more information, see the [documentation](#). For pricing estimates, see the [pricing page](#).

Minimum event pollers
If blank, Lambda sets a value of 1.

1
Specify a whole number between 1 and 200.

Maximum event pollers
If blank, Lambda sets a value of 200.

100
Specify a whole number between 1 and 2000.

► **Additional settings**
In order to read from the Apache Kafka trigger, your execution role must have proper permissions.

Cancel

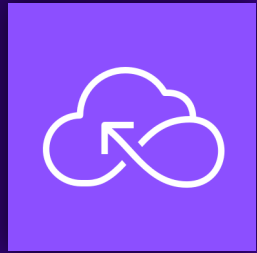
CloudShell Feedback © 2024, Amazon Web Services, Inc. or its affiliates. Privacy Terms



[AWS Announcement](#)



Other Announcements



1. **AWS Lambda runtime updates:**
 - a) **Support for Python 3.13** (pre-re:Invent)
 - b) **Support for Node.js 22** (pre-re:Invent)
2. **AWS Lambda now supports AWS Fault Injection Service (FIS) actions** (pre-re:Invent)
3. **S3 support as failed-event destination for all integrations in AWS Lambda** (pre-re:Invent)
4. **AWS AppSync now supports cross account sharing of GraphQL APIs** (pre-re:Invent)
5. **Announcing new APIs for Amazon Location Service Routes, Places, and Maps** (pre-re:Invent)

Developer Tools, Management, and Governance

32 announcements (18 *pre-re:Invent*)



Timeline view in AWS CloudFormation

PEEK INSIDE YOUR AWS CLOUDFORMATION DEPLOYMENTS WITH TIMELINE VIEW

With this release, AWS CloudFormation Timeline View complements the existing views by giving you additional visibility into what happens during deployments.

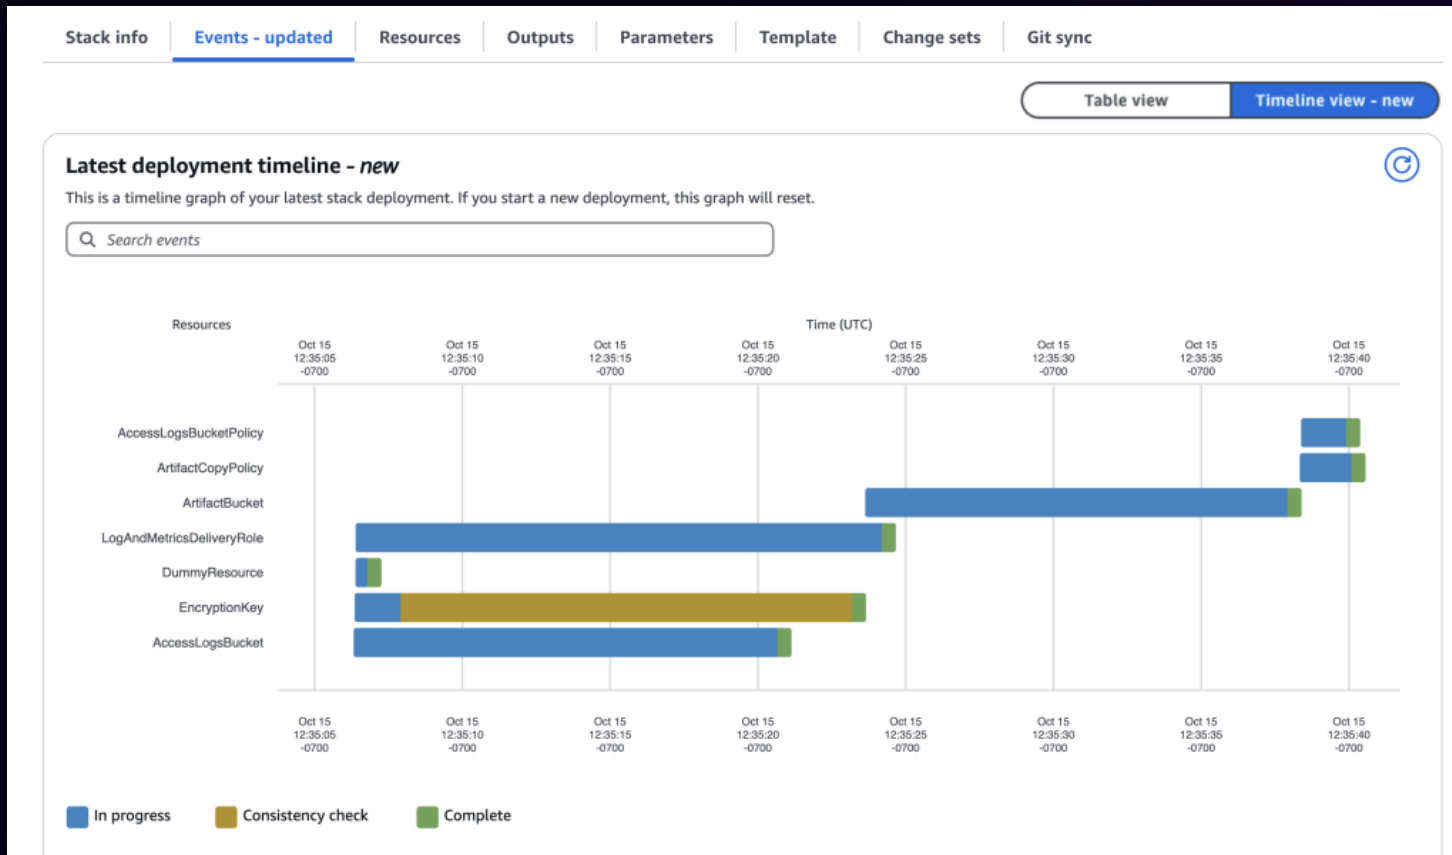


GA

Nov, 11th

pre-re:Invent

All Regions

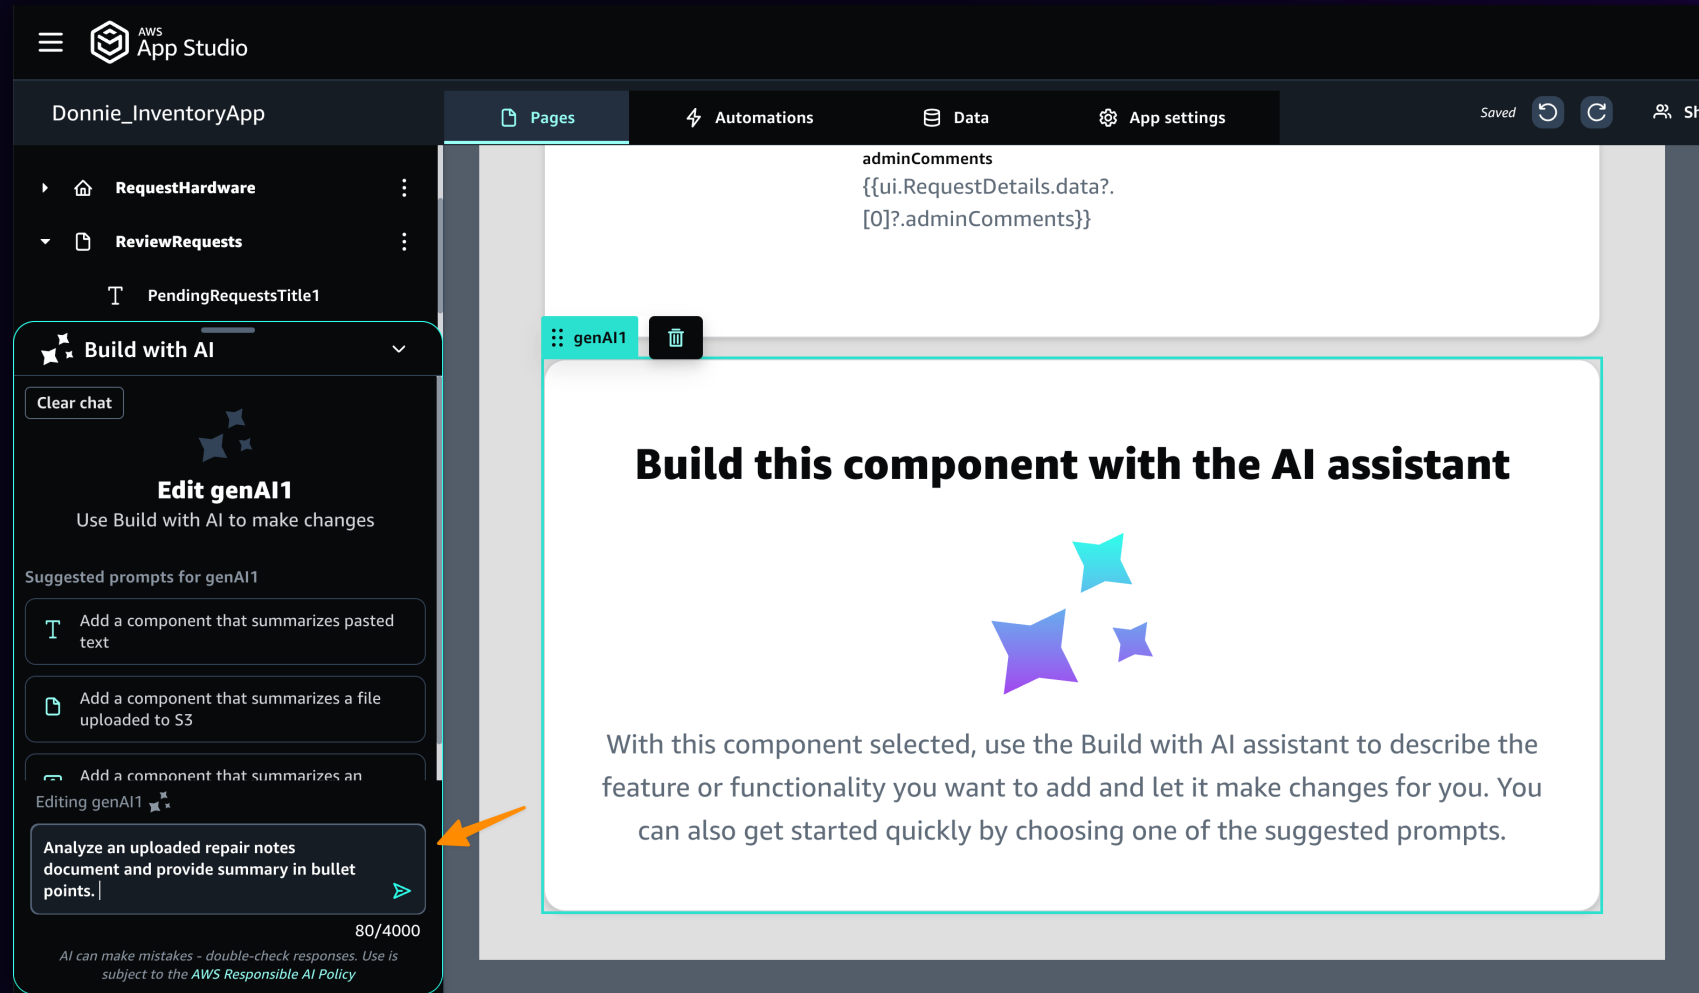


[AWS DevOps Blog Post](#)



General Availability of AWS App Studio

EMPOWERING A NEW SET OF BUILDERS TO CREATE BUSINESS APPLICATIONS



The screenshot displays the AWS App Studio interface for a project named 'Donnie_InventoryApp'. The top navigation bar includes 'Pages', 'Automations', 'Data', and 'App settings'. The main workspace shows a component named 'adminComments' with a data binding: `{{ui.RequestDetails.data?[0]?.adminComments}}`. A 'genAI1' component is highlighted with a red border. On the left, the 'Build with AI' panel is open, showing a 'Clear chat' button, the 'Edit genAI1' section, and a list of suggested prompts. An orange arrow points to the prompt: 'Analyze an uploaded repair notes document and provide summary in bullet points.' The panel also shows a character count of 80/4000 and a disclaimer: 'AI can make mistakes - double-check responses. Use is subject to the AWS Responsible AI Policy.'

GA

Nov, 18th

pre-re:Invent

2 Regions



[AWS News Blog Post](#)

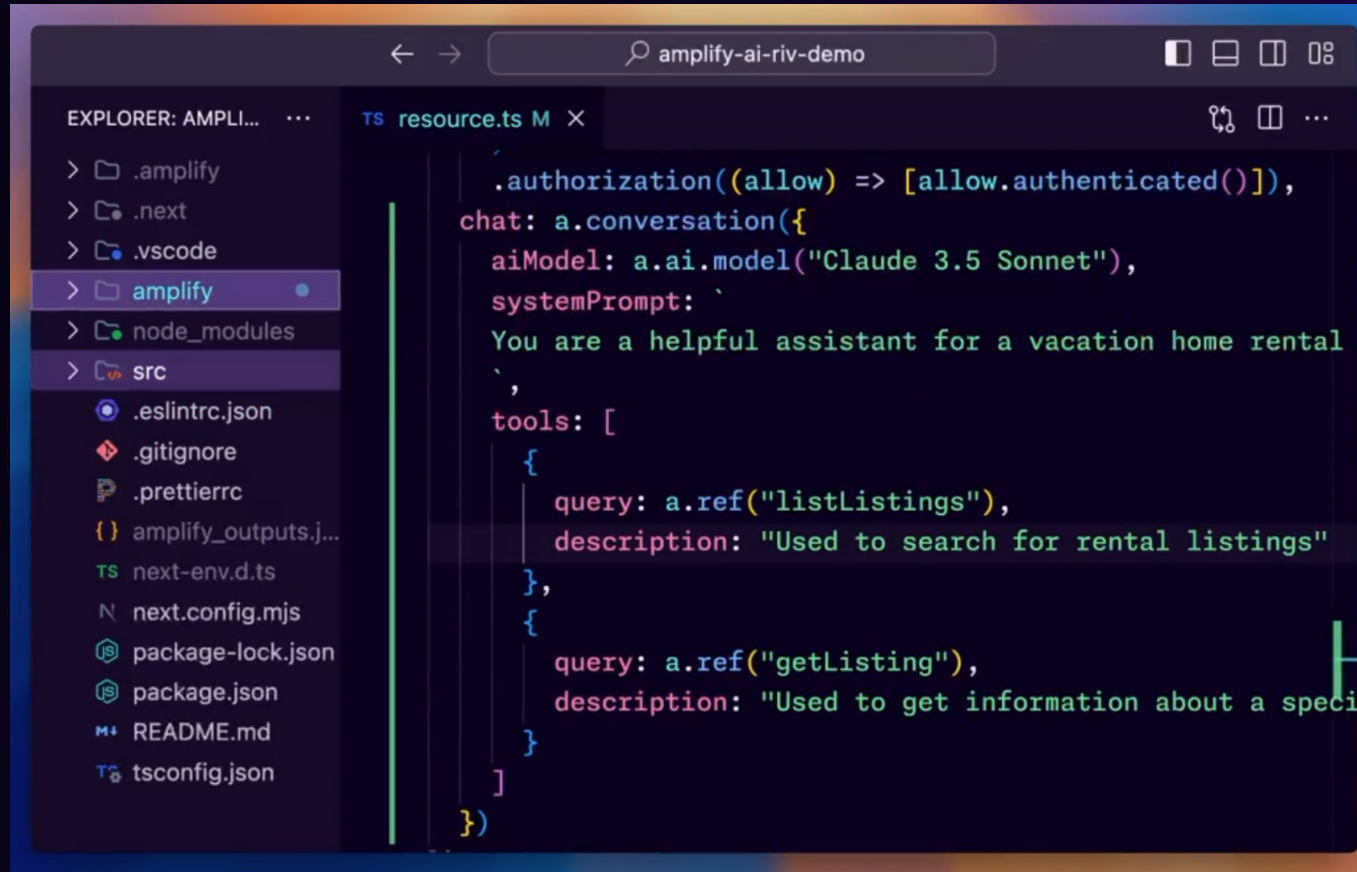


AWS Amplify AI Kit

BUILD FULLSTACK AI APPLICATIONS IN MINUTES WITH AWS AMPLIFY

AWS Amplify Kit provides following capabilities:

- `<AIConversation>` React UI component that provides a real-time, streaming chat experience with additional features.
- Type-safe client with secure server-side access to Amazon Bedrock, including secure ways to share user context.
- Fullstack TypeScript developer experience on top of Amplify Gen 2 and AWS AppSync.



```
    .authorization((allow) => [allow.authenticated()]),
  chat: a.conversation({
    aiModel: a.ai.model("Claude 3.5 Sonnet"),
    systemPrompt: `
You are a helpful assistant for a vacation home rental
`,
    tools: [
      {
        query: a.ref("listListings"),
        description: "Used to search for rental listings"
      },
      {
        query: a.ref("getListings"),
        description: "Used to get information about a speci"
      }
    ]
  })
}
```



GA

Nov, 19th

pre-re:Invent

All Regions



[AWS News Blog Post](#)



Application Signals for AWS Lambda



AMAZON CLOUDWATCH APPLICATION SIGNALS SUPPORT FOR AWS LAMBDA ENVIRONMENTS

This announcement expands initial release of Amazon CloudWatch Application Signals, an AWS built-in application performance monitoring (APM) solution, with support for AWS Lambda-based applications.

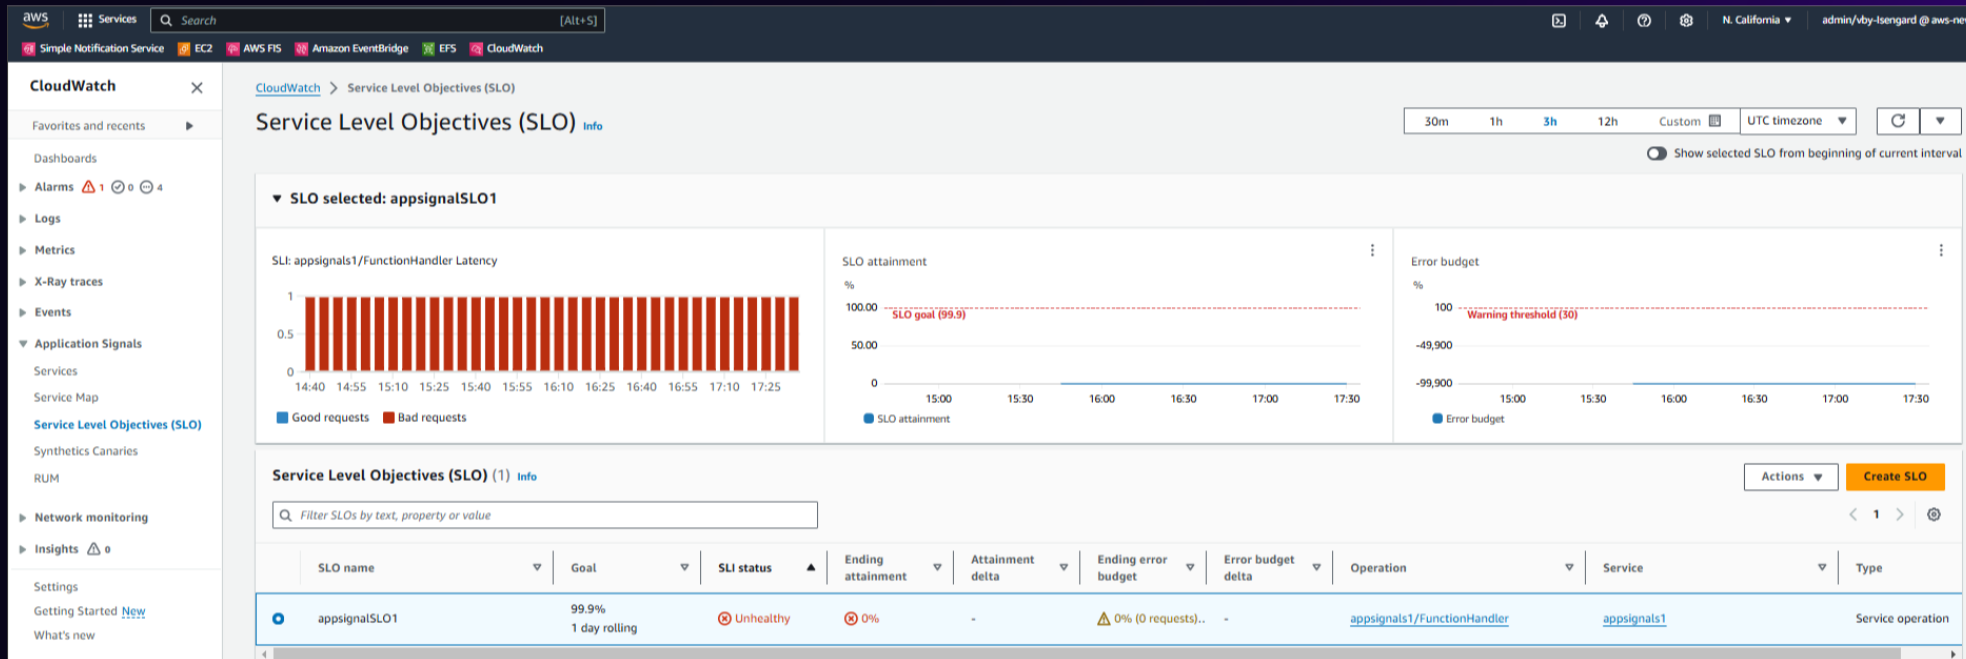
It automatically correlates telemetry across metrics, traces, and logs, to speed up troubleshooting and reduce application disruption.

GA

Nov, 21st

pre-re:Invent

All Regions



[AWS News Blog Post](#)



Amazon Q Developer at **pre-re:Invent**

ECLIPSE IDE, CLI FOR CODE TRANSFORMATION AND STEP-BY-STEP UPGRADES, COST ANALYSIS

Similarly to the Q Business, Amazon Q Developer service team provided 9 different updates in the pre-re:Invent period – and to name the most important ones:

- **Amazon Q Developer Pro Tier** received updates for administrators and regular users to track their activity.
- **Amazon Q Developer** now can help with AWS CloudFormation troubleshooting.
- Support for a new IDE: **Eclipse** (*public preview*).
- **Amazon Q Developer Code Transformation** received CLI tool (*public preview*) and ability to perform step-by-step upgrades with explanation and transforming Oracle SQL to PostgreSQL dialect in Java applications.

Commands

For step-by-step instructions for running these commands, see [Running a transformation on the command line with Amazon Q Developer](#).

To configure a transformation and authenticate to Amazon Q Developer Pro, run:

```
qct configure
```

To start a transformation, run:

```
qct transform --source_folder <path-to-folder>
```

To start a transformation with a [customization](#), run:

```
qct transform --source_folder <path-to-folder>  
--custom_transformation_file <path-to-orchestrator-file>
```

To get help with transformations, run:

```
qct -h
```



GA

PP

Nov, 2024

pre-re:Invent

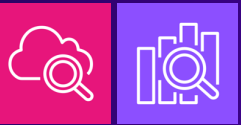
All Regions



[AWS Announcement](#)

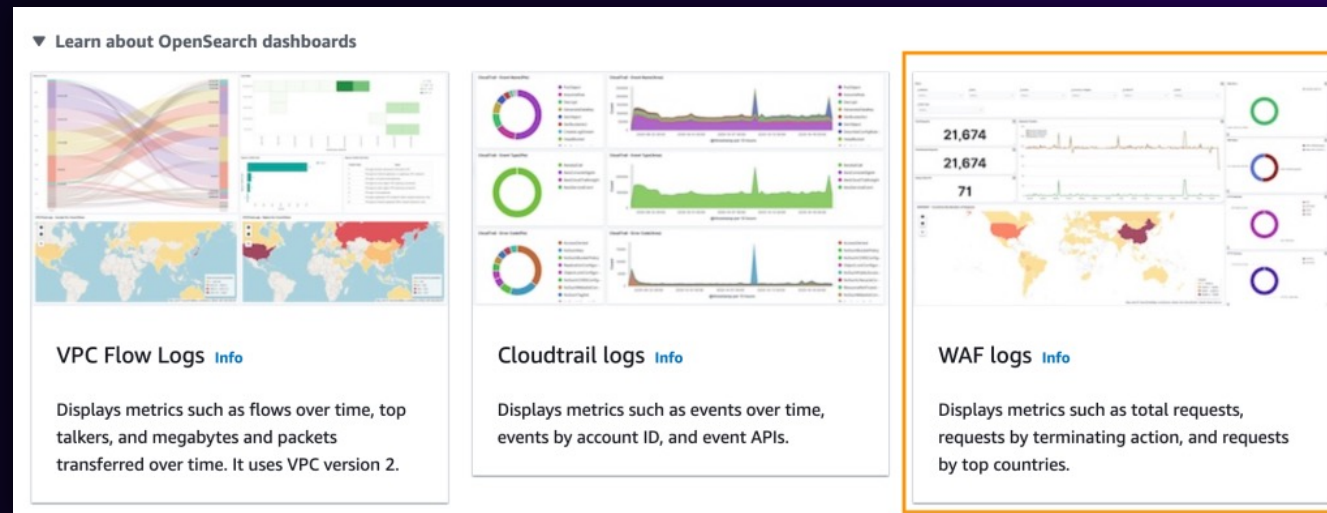
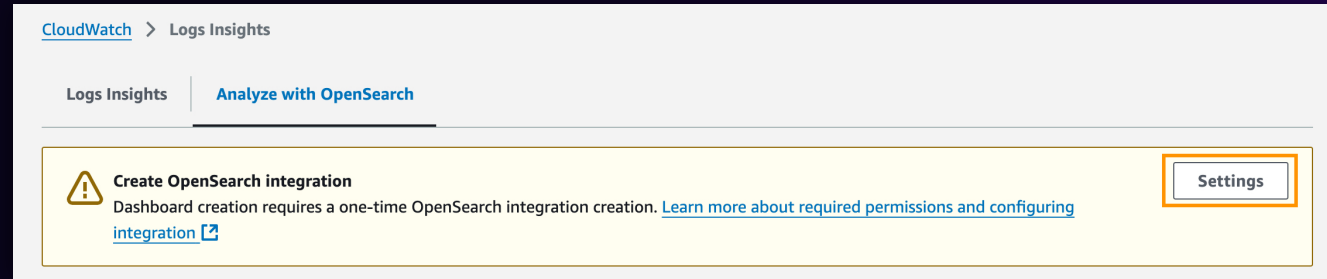
Integrated Analytics Experience

BETWEEN AMAZON CLOUDWATCH AND AMAZON OPENSEARCH SERVICE



Integrated Analytics Experience is a zero-ETL integration between Amazon CloudWatch and Amazon OpenSearch Service.

This integration simplifies log data analysis and visualization without data duplication, streamlining log management while reducing technical overhead and operational costs.



GA

Dec, 1st

re:Invent

All Regions



[AWS News Blog Post](#)



Amazon CloudWatch Database Insights

FOR AMAZON AURORA MYSQL AND POSTGRESQL



GA

Dec, 1st

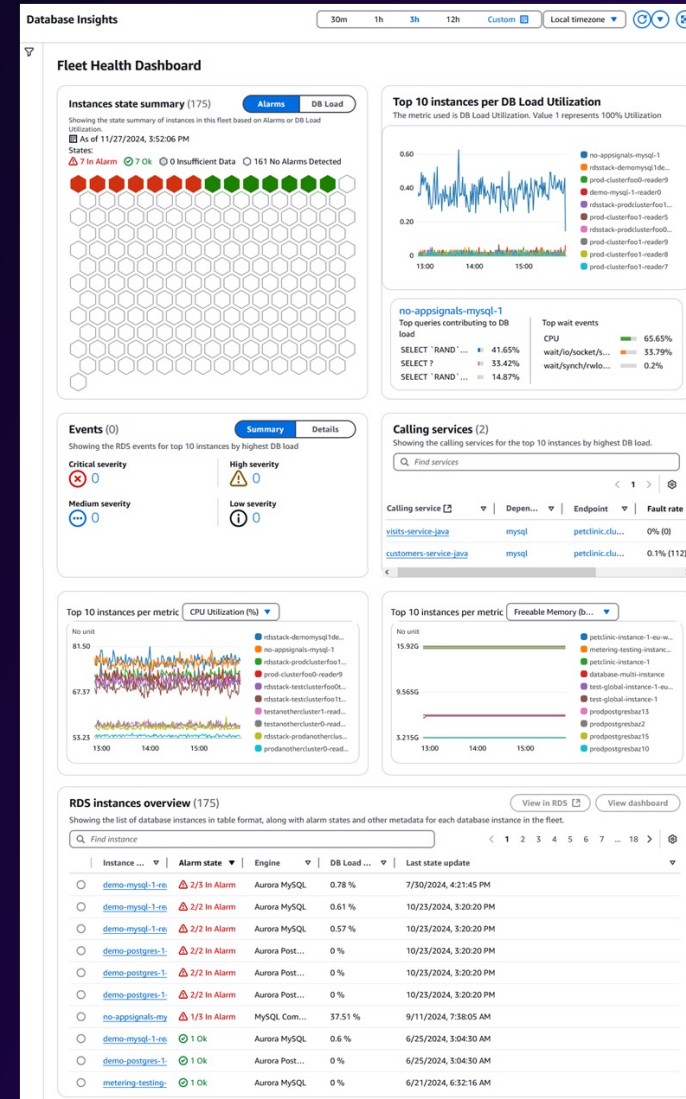
re:Invent

All Regions

Observing your Amazon Aurora databases is now a whole lot easier.

Instead of spending time setting up telemetry, building dashboards, and configuring alarms, you just open **Amazon CloudWatch Database Insights** and take a look.

With no further setup, you can monitor the health of all of your Amazon Aurora MySQL and PostgreSQL instances in the selected region with a dashboard you can see on the side.



Declarative Policies (DPs)

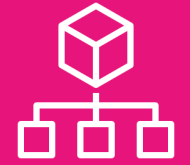
SIMPLIFY GOVERNANCE WITH DECLARATIVE POLICIES

Declarative Policies (DPs) are new capability that helps you declare and enforce desired configuration for a given AWS Service at scale across your AWS Organizations.

At the moment, they support Amazon EC2, Amazon VPC, and Amazon EBS services, with service attributes that include enforcing IMDSv2, allowing troubleshooting through serial console, allowed AMI settings, and blocking public access for Amazon EBS snapshots, Amazon EC2 AMI, and VPC.

The screenshot shows the AWS Organizations Policies page. At the top, it says 'AWS Organizations > Policies' and 'Policies'. Below that, it states 'Policies in AWS Organizations enable you to manage different features of the AWS accounts in your organization. [Learn more](#)'. The main content is a table titled 'Supported policy types' with columns for 'Policy type' and 'Status'. The table lists several policy types, all of which are currently 'Disabled'. The 'Declarative policies for EC2' row is highlighted with an orange border.

Policy type	Status
AI services opt-out policies AI services opt-out policies allow you to control data collection for AWS AI services for all the accounts in an organization. Learn more	⊖ Disabled
Backup policies Backup policies allow you to centrally manage and apply backup plans to the AWS resources across an organization's accounts. Learn more	⊖ Disabled
Chatbot policies Chatbot policies allow you to control access to an organization's accounts from chat applications such as Slack and Microsoft Teams. Learn more	⊖ Disabled
Declarative policies for EC2 Declarative policies for EC2 enable you to define and enforce desired configuration for EC2 across all your accounts. Once attached, configuration is always maintained. Learn more	⊖ Disabled
Resource control policies Resource control policies (RCPs) offer central control over the maximum available permissions for resources in an organization. Learn more	⊖ Disabled
Service control policies Service control policies (SCPs) offer central control over the maximum available permissions for IAM users and IAM roles in an organization. Learn more	⊖ Disabled
Tag policies Tag policies allow you to standardize the tags attached to the AWS resources in an organization's accounts. Learn more	⊖ Disabled



GA

Dec, 1st

re:Invent

All Regions



[AWS News Blog Post](#)



Amazon Q Developer at re:Invent

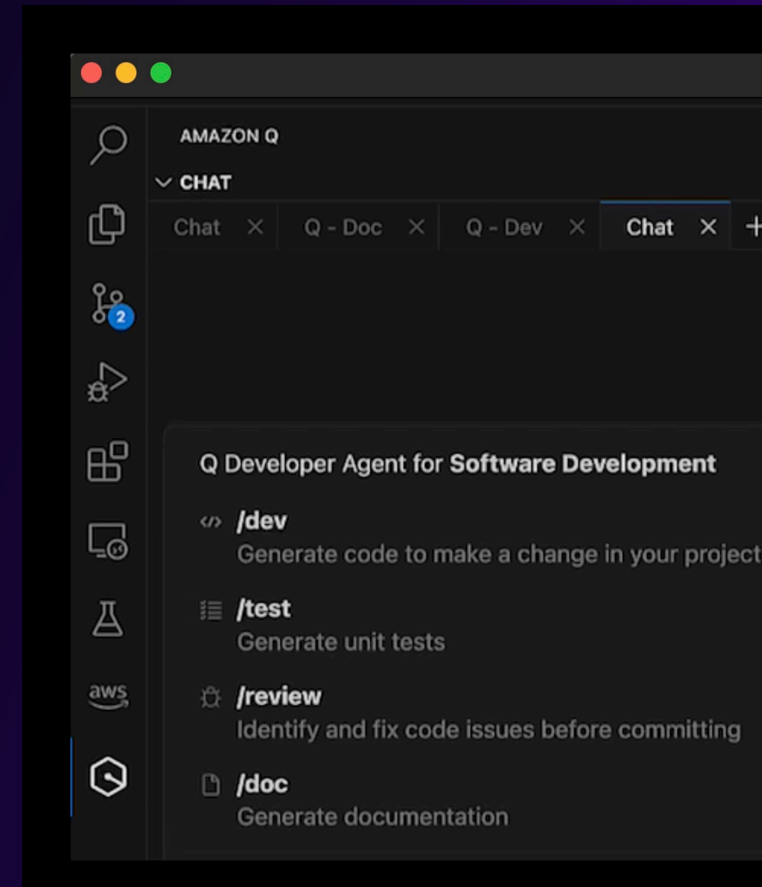
GENERATING CODE REVIEW, TESTS, DOCUMENTATION, GITLAB INTEGRATION, AND OPS HELP

With announcements presented this year, Amazon Q Developer got 3 new agent capabilities:

1. Enhancing documentation in codebases via `/doc` command.
2. Supporting code reviews to detect and resolve security and code quality issues via `/review`.
3. Generating unit tests automatically and improving test coverage `/test`.

All of the above in your preferred IDE or GitLab Duo with Amazon Q (*preview*), which is one of the most popular enterprise DevOps platforms.

Additionally, there is a new capability in Amazon Q Developer to **investigate and remediate operational issues** (*preview*).



GA

PP

Dec, 3rd

re:Invent

All Regions



[AWS News Blog Post](#)

Transformation capabilities of Q Developer

.NET APPS FROM WINDOWS TO LINUX, VMWARE WORKLOADS AND MAINFRAMES



PP

Dec, 3rd

re:Invent

All Regions

This is your workspace description. A workspace is a collection of jobs, artifacts, connectors and collaborators. You can chat with Q about any of these concepts.

Ask Amazon Q

PR Create a new job

I'd be happy to help you create a new job. To get started, could you please specify which type of transformation you'd like to work on? I support the following types of migrations:

1. Mainframe modernization
2. .NET modernization
3. VMware migration

Once you let me know which type you're interested in, I can guide you through the process of creating a new job for that specific transformation.

PR .NET modernization

Thank you for specifying that you'd like to work on .NET modernization. I'll be glad to help you create a new job for this type of transformation. For .NET modernization, I currently support one specific objective: "Modernize .NET Applications to cross-platform .NET". This objective covers the process of updating your .NET applications to make them compatible with cross-platform .NET environments.

Let's create a job for this. I'll suggest a job name, but please feel free to modify it if you prefer something different.

Suggested job name: DotNetModernization

Here are the details for the new job:

- Job Type: .NET modernization
- Objective: Modernize .NET Applications to cross-platform .NET
- Job Name: DotNetModernization

Are you happy with these details, or would you like to change the job name?

What would you like to transform today?



[AWS News Blog Post](#)



Q Developer in Amazon SageMaker Canvas

GUIDING MACHINE LEARNING MODELS DEVELOPMENT WITH NATURAL LANGUAGE



PP

Dec, 4th

re:Invent

6 Regions

Adding Amazon Q Developer to Amazon SageMaker Canvas helps to connect machine learning expertise with business needs.

Natural language conversations are making the process of building ML models much easier for domain experts.

Those roles can build accurate, production-quality models, even if they don't have machine learning expertise.

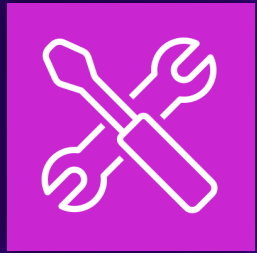
A screenshot of the Amazon SageMaker Canvas Q Developer interface. At the top, a user icon is followed by the question: "What data prep steps did you apply to transform the data?". Below this, a blue hexagonal icon with a white 'Q' is followed by a response: "I've fixed the issues below which will result in better data quality suited for model training. Additionally, you can see a preview of the transformed data. Note that these changes will not affect your stored data in Amazon S3." Underneath the response is a section titled "Data Preparation Summary" with a bullet point: "127 outliers were removed". A link "View Dataset in Canvas" is provided. At the bottom of the response area are three buttons: "Launch my training job." (highlighted with an orange border), "Show analysis report", and "Start a new conversation". Below the response area is a text input field with a downward arrow icon on the left and a rightward arrow icon on the right. At the very bottom, a small text line reads: "Use of Amazon Q is subject to the AWS Responsible AI Policy".



[AWS News Blog Post](#)



Other Announcements



1. [Introducing a new experience for AWS Systems Manager](#) (pre-re:Invent)
2. [AWS CodePipeline actions: publishing ECR image and AWS InspectorScan](#) (pre-re:Invent)
3. **AWS CloudFormation Hooks updates:**
 - a) [AWS CloudFormation Hooks now support custom AWS Lambda functions](#) (pre-re:Invent)
 - b) [AWS CloudFormation Hooks introduces stack and change set target invocation points](#) (pre-re:Invent)
 - c) [Author AWS CloudFormation Hooks using CloudFormation Guard domain specific language](#) (pre-re:Invent)
4. [Amazon CloudWatch adds network performance monitoring using flow monitors](#) (re:Invent)

Security, Identity, and Compliance

16 announcements (11 *pre-re:Invent*)



Resource Control Policies (RCPs)

CENTRALLY RESTRICT ACCESS TO AWS RESOURCES WITH NEW MANAGED AUTHORIZATION POLICY



GA

Nov, 13th

pre-re:Invent

All Regions

Resource Control Policy (RCP) is a new authorization policy managed in AWS Organizations that can be used to set the maximum available permissions on resources within your entire organization.

They are a type of preventative control that help you establish a data perimeter in your AWS environment and restrict external access to resources at scale.

Policies

Policies in AWS Organizations enable you to manage different features of the AWS accounts in your organization. [Learn more](#)

Supported policy types

Policy type	Status
AI services opt-out policies AI services opt-out policies allow you to control data collection for AWS AI services for all the accounts in an organization. Learn more	Disabled
Backup policies Backup policies allow you to centrally manage and apply backup plans to the AWS resources across an organization's accounts. Learn more	Disabled
Chatbot policies Chatbot policies allow you to control access to an organization's accounts from chat applications such as Slack and Microsoft Teams. Learn more	Disabled
Resource control policies Resource control policies (RCPs) offer central control over the maximum available permissions for resources in an organization. Learn more	Disabled
Service control policies Service control policies (SCPs) offer central control over the maximum available permissions for IAM users and IAM roles in an organization. Learn more	Disabled
Tag policies Tag policies allow you to standardize the tags attached to the AWS resources in an organization's accounts. Learn more	Disabled



[AWS News Blog Post](#)



Centralized Management of Root Access

USING AWS ORGANIZATIONS



GA

Nov, 15th

pre-re:Invent

All Regions

AWS Identity and Access Management (IAM) is launching a new capability allowing security teams to centrally manage root access for member accounts in AWS Organizations.

You can now easily manage root credentials and perform highly privileged actions with ability to establish a root session in a member AWS Account.

Centralized root access for member accounts [Info](#)

[Disable](#)

Centralized root access allows you to manage root credentials for your member accounts from the [root access management](#) page. You can launch a privileged session into your member accounts to delete root credentials and perform other privileged actions.

Available features

<p>Root credentials management</p> <p>You can delete and audit root credentials of member accounts. You can also allow password recovery for specific member accounts.</p> <p>✔ Enabled</p> <p>Disable</p>	<p>Privileged root actions in member accounts</p> <p>You can take certain root actions in member accounts like deleting misconfigured policies in Amazon SQS or Amazon S3.</p> <p>✔ Enabled</p> <p>Disable</p>	<p>Delegated administrator for centralized root access</p> <p>In addition to the management account, this member account will be authorized to have centralized root access for other member accounts.</p> <p>Assign</p>
--	--	--



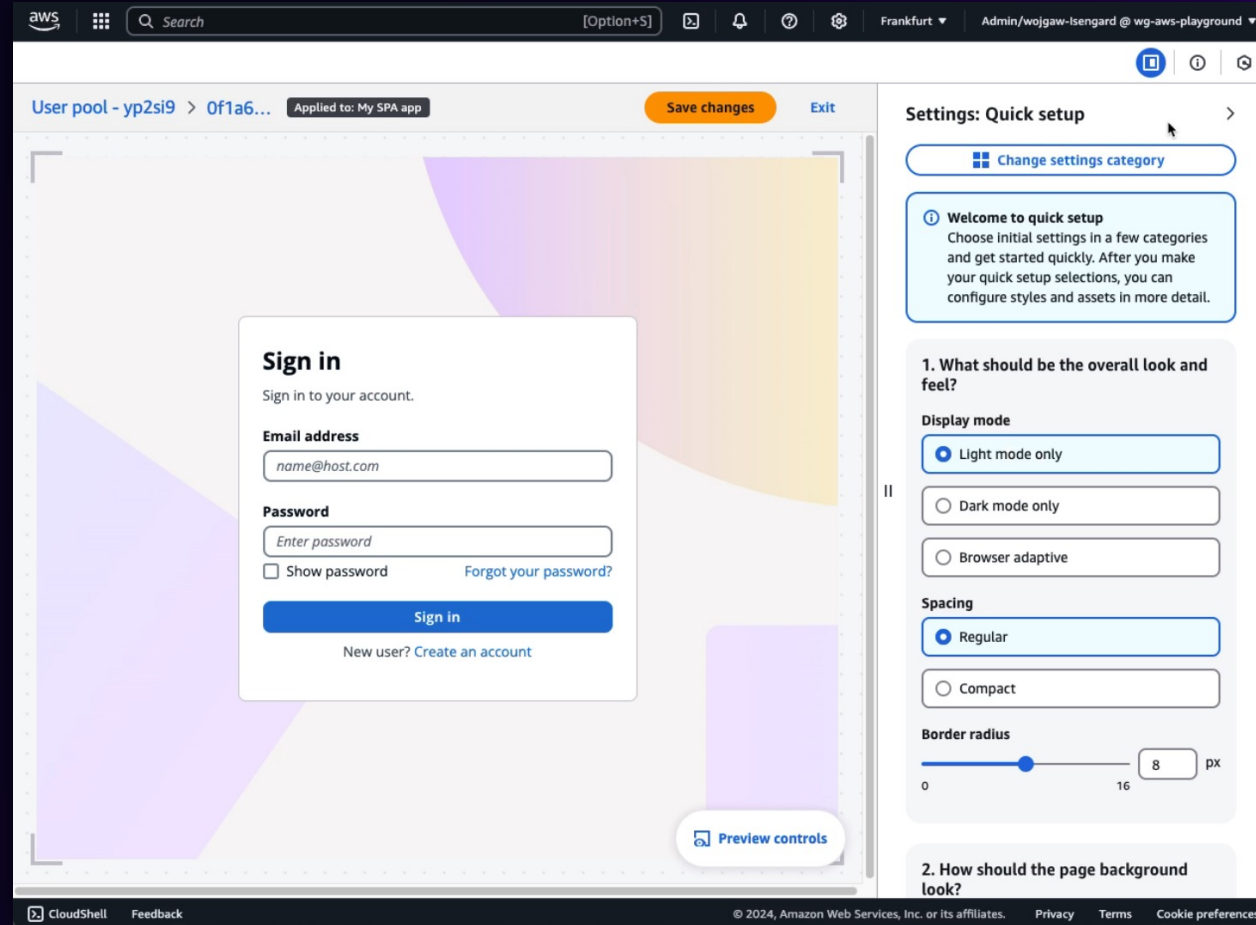
[AWS News Blog Post](#)



New Amazon Cognito features

After 10 years from the initial announcement, Amazon Cognito gets a well-expected overhaul with batch of updates:

- Improved **Developer Experience (DX)** for popular programming languages and frameworks, and improved **User Experience (UX)** with the new UI in AWS Management Console.
- **Managed Login**, as an improved version of *Hosted UI* with enhanced customization, visual builder, and templates.
- **Passwordless Login**, including Passkey authentication.
- Updated **Pricing Plans**.



GA

Nov, 22nd

pre-re:Invent

All Regions



[AWS News Blog Post](#)



Extended Threat Detection for GuardDuty

AMAZON GUARDDUTY NOW IDENTIFIES ATTACK SEQUENCE WITH AI/ML



GA

Dec, 1st

re:Invent

All Regions

GuardDuty > Findings > Finding details

Potential data compromise of one or more S3 buckets involving a sequence of actions over multiple signals associated with a user in your account.

Critical First seen a month ago, last seen a month ago

Actions ▾

A sequence of actions involving 8 signals indicating a possible data compromise of one or more S3 bucket(s) was observed for AssumedRole/Admin with principalId [redacted] in account [redacted] between 2024-10-17T17:47:24Z and 2024-10-18T17:11:44Z with the following behaviors:

- 6 MITRE ATT&CK tactics observed: Exfiltration, Impact, Persistence, Defense Evasion, Credential Access, Discovery
- 1 MITRE ATT&CK techniques observed: Data Destruction
- 7 sensitive APIs called: apigateway:GetApiKeys, cloudtrail:DeleteTrail, iam:CreateAccessKey, iam:ListGroup, s3:DeleteObject, s3:GetObject, s3:ListObjects

Finding ID	[redacted]	📄
Type	AttackSequence:S3/CompromisedData	⋮
Region	eu-west-1	📄
Account	[redacted]	⋮

Attack sequence

MITRE tactics

Discovery → Initial Access → Execution → Persistence → Privilege Escalation → Defense Evasion → Lateral Movement → Credential Access → Collection → Exfiltration → Impact

▶ Indicators (3)
Reasons why this collection of signals was identified as an attack sequence.

▶ Actors (1)
Information about the actors identified in this attack sequence.

▶ Endpoints (3)
Information about network endpoints that were involved in this attack sequence.

Signals (8) Resources (4)

Sort by: Newest first

October 18, 2024

19:11:44 (UTC+02:00)

The API GetApiKeys was invoked from a Linux computer. **Defense Evasion**

Resources involved

Access Key: [redacted]

- ▶ Indicators (2)
- ▶ Actors (1)
- ▶ Endpoints (1)

19:11:19 (UTC+02:00)

The API ListGroups was invoked from a Security Linux computer. **Discovery**

Resources involved

Access Key: [redacted]

- ▶ Indicators (1)
- ▶ Actors (1)
- ▶ Endpoints (1)



AWS Security Incident Response

MANAGE AND NEUTRALIZE SECURITY EVENTS QUICKLY AND EFFECTIVELY

AWS Security Incident Response is a new service designed to help organizations manage security events quickly and effectively.

The service is purpose-built to help customers prepare for, respond to, and recover from various security events, including account takeovers, data breaches, and ransomware attacks.

It automates the triage and investigation of security findings from Amazon GuardDuty and integrated third party threat detection tools through AWS Security Hub.

The screenshot displays the AWS Security Incident Response console for a specific case (ID: 3421132312). At the top, a green notification bar states "Case created successfully. We will acknowledge the case within 15 minutes. In the meantime, you can set up permissions or add more details to the case." Below this, the breadcrumb navigation shows "AWS Security Incident Response > Cases > 3421132312".

The main content area is titled "3421132312" and features an "Overview" section with the following details:

- Resolver:** AWS
- Name:** Active incident [2024-05-07]
- Type:** Active Incident
- Created at:** 2024-05-07T14:32:35+03:00
- State date estimate:** 2024-05-07
- Status:** Submitted (with a dropdown menu showing options: Submitted, Acknowledged, Detection & Analysis, Containment, Eradication & Recovery, Post-Incident activities, Ready to close, Closed)
- Actions:** Actions required (button)
- Last updated:** 0 seconds ago

Below the overview, there are tabs for "Details", "Communications", "Permissions", "Attachments", "Tags", and "Case activities". The "Details" tab is active, showing a "Description" section with a list of questions and answers regarding a security issue causing latency. Other sections include "Case details" (Incident start time, Main services impacted: Amazon GuardDuty, Main regions impacted: Africa (Capetown), US East(Ohio), US West (N. California), US West (Oregon), Impacted accounts: 123456789012, 2123212211211, 122422422134) and "Threat actor IPs (3)" (This section is automatically updated by the service and the AWS security professional assigned to your case. Table with columns: IP Address, User agent).

GA

Dec, 1st

re:Invent

12 Regions



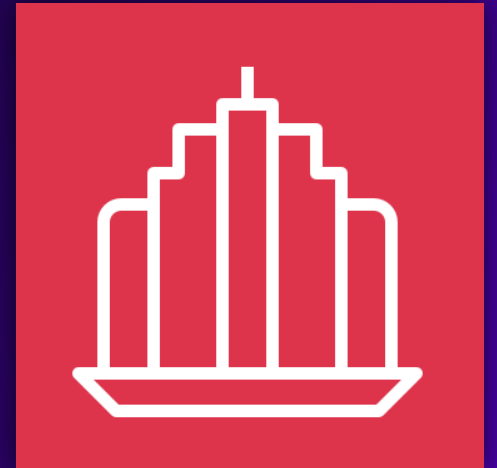
Other Announcements



1. [OCSF joins the Linux Foundation: Accelerating the Standardization of Cybersecurity Data](#) (pre-re:Invent)
2. [Verify the Safety of the Rust Standard Library](#) (pre-re:Invent)
3. [Customize scope of IAM Access Analyzer unused access analysis](#) (pre-re:Invent)
4. [Amazon Verified Permissions launches new API to get multiple policies](#) (pre-re:Invent)
5. [AWS CloudTrail Lake enhancements for cloud visibility and investigations](#) (pre-re:Invent)
6. [Amazon OpenSearch Service and Amazon Security Lake zero-ETL integration](#) (re:Invent)
7. [AWS Verified Access supports secure access for SSH and RDP protocols](#) (re:Invent, **PP**)
8. [AWS post-quantum cryptography migration plan](#) (re:Invent)

Business Applications

15 announcements (1 *pre-re:Invent*)



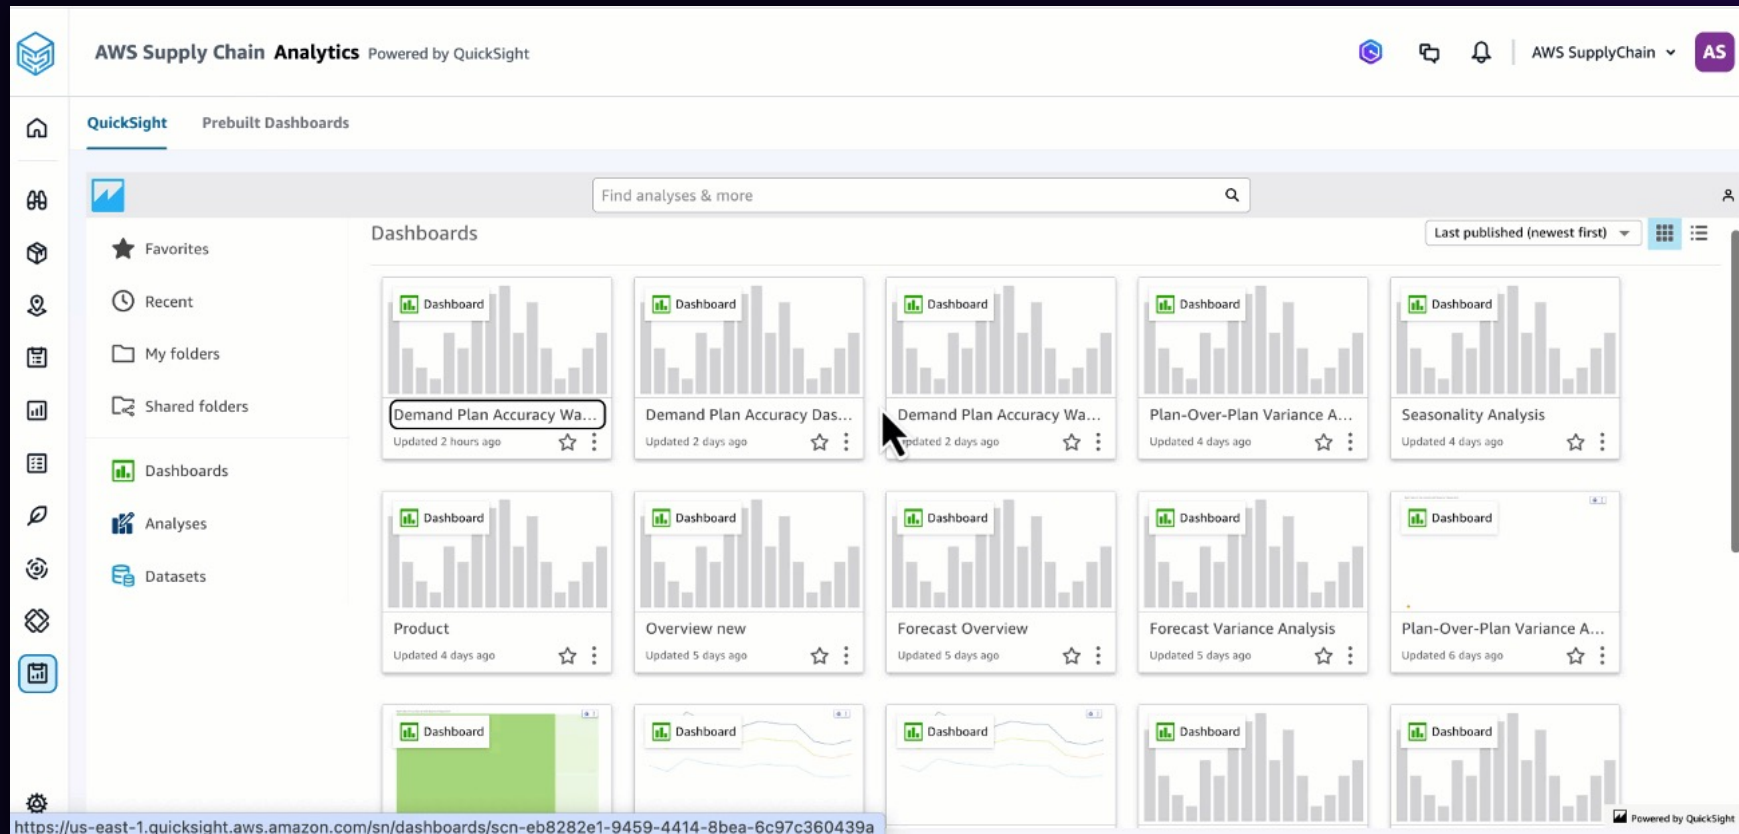
Embedded Analytics in AWS Supply Chain



EMBEDDED ANALYTICS POWERED BY AMAZON QUICKSIGHT

AWS Supply Chain Analytics are custom report dashboards powered by Amazon QuickSight, that are available for business analysts or supply chain managers inside the service.

From now they can perform custom analyses, create visualization, and gain insights about operations.



GA

Oct, 31st

pre-re:Invent

All Regions



[AWS News Blog Post](#)



New capabilities in Amazon Connect

WHATSAPP BUSINESS, SIMPLIFIED AI BOT CREATION, GUARDRAILS AND Q IMPROVEMENTS



GA

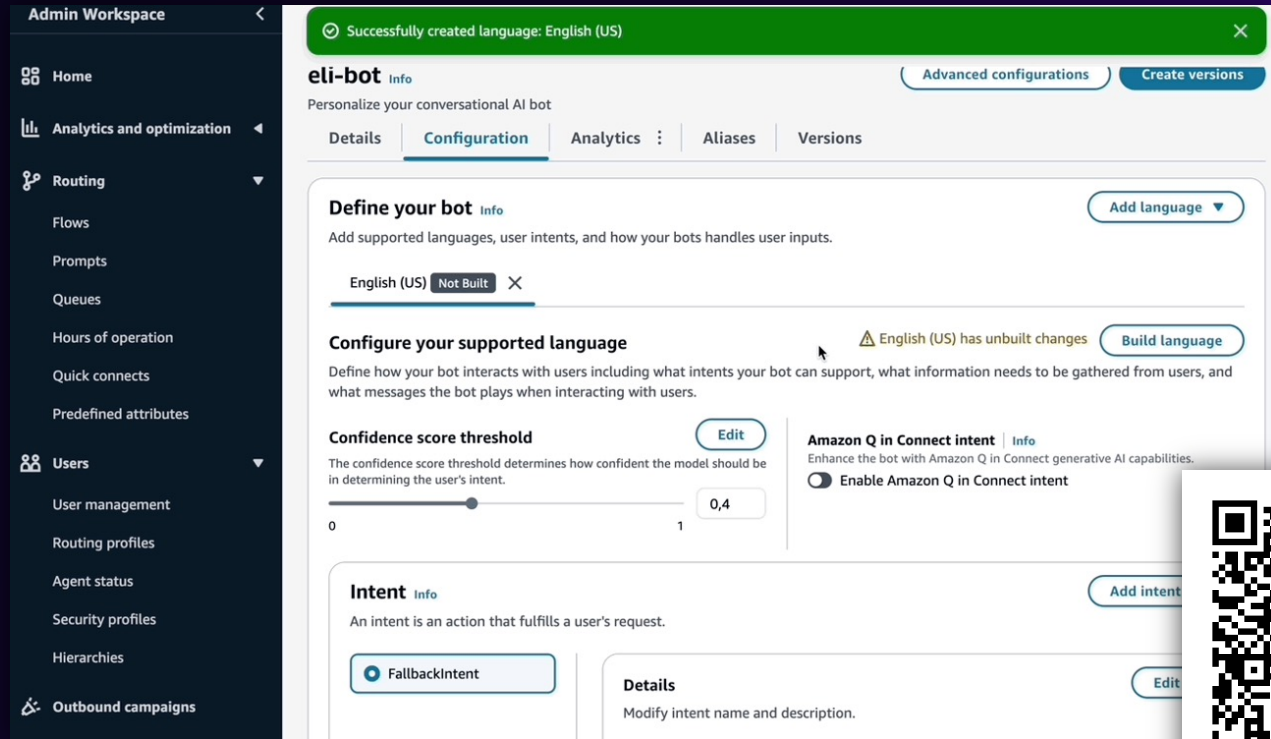
Dec, 1st

re:Invent

All Regions

A set of new features that help businesses enhance their contact center operations through generative AI, advanced security features, and streamlined bot management.

These innovations help businesses deliver better customer experiences by creating more time and space for meaningful human interactions, while maintaining security and compliance.



[AWS News Blog Post](#)



Session Recordings from AWS re:Invent 2024 (1/2)

1. [Slide 34: Store tabular data at scale with Amazon S3 Tables \(STG367\)](#)
2. [Slide 43: Amazon Elastic VMware Service: Start your modernization with AWS \(MAM119\)](#)
3. [Slide 47: Achieving scale with Amazon Aurora PostgreSQL Limitless Database \(DAT420\)](#)
4. [Slide 50: Build scalable and cost-optimized apps with Amazon Aurora Serverless \(DAT316\)](#)
5. [Slide 51: Improve resiliency using Amazon MemoryDB Multi-Region \(DAT426\)](#)
6. [Slide 52: Transform your data with Oracle Database@AWS \(DAT246\)](#)
7. [Slide 53: Multi-region strong consistency with Amazon DynamoDB global tables \(DAT425\)](#)
8. [Slide 54: Get started with Amazon Aurora DSQL \(DAT424\)](#)
9. [Slide 70: Amazon SageMaker Lakehouse: Accelerate analytics & AI \(ANT354\)](#)
10. [Slide 72: Simplify business scenario analysis with Amazon Q in QuickSight \(BSI104\)](#)
11. [Slide 73: Empower your data journey with Amazon DataZone's data lineage \(ANT207\)](#)
12. [Slide 84: Introducing automated reasoning checks in Amazon Bedrock Guardrails \(AIM393\)](#)
13. [Slide 87: Practical generative AI using Amazon Nova \(AIM398\)](#)

Session Recordings from AWS re:Invent 2024 (2/2)

14. [Slide 91: Amazon SageMaker HyperPod: Reduce costs with new governance capability \(AIM388\)](#)
15. [Slide 94: Real-time event patterns with WebSockets and AWS AppSync \(FWM302\)](#)
16. [Slide 96: Accelerate Python and .NET Lambda functions with SnapStart \(SVS218\)](#)
17. [Slide 97: What's new with Amazon CloudWatch \(COP381\)](#)
18. [Slide 98: Building advanced workflows with AWS Step Functions \(API402\)](#)
19. [Slide 102: How to choose between AWS CloudFormation, Terraform, and AWS CDK \(DEV309\)](#)
20. [Slide 103: AWS App Studio: The fastest and easiest way to build applications \(AIM218\)](#)
21. [Slide 110: What's new with Amazon Q Developer agents \(DOP223\)](#)
22. [Slide 111: Modernize mainframe applications faster using Amazon Q Developer \(DOP221\)](#)
23. [Slide 111: Accelerate .NET porting from Windows to Linux with Amazon Q Developer \(DOP208\)](#)
24. [Slide 118: Uncovering sophisticated cloud threats with Amazon GuardDuty \(SEC219\)](#)
25. [Slide 119: Respond and recover faster with AWS Security Incident Response \(SEC360\)](#)
26. [Slide 123: Reimagine end customer self-service with Amazon Connect \(BIZ223\)](#)

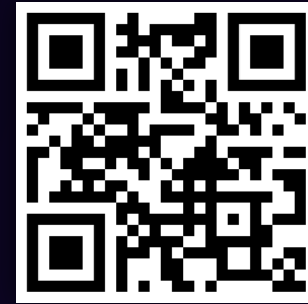
Stay Connected with AWS



[Amazon Web Services](#)



[Community.AWS](#)



[AWS Skill Builder](#)



[@AWSDevelopers](#)



[@aws-developers](#)



[@AWSDevelopers](#)



[collectives/aws](#)