



AWS re:Invent

December 1 – 5, 2025



builder.aws.com

Copyright © 2025 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



builder.aws.com

Deck Legend

Generally Available (GA)

GA

Public Preview (PP)

PP

Gated Preview (P)

P

Coming Soon (CS)

CS

Color coding:

- **pre-re:Invent**
- **re:Invent**
- **Service Name**
- **Open Source**

Each announcement contains:

- Date when it happened
- Announcement status
- pre-re:Invent or re:Invent release
- Regional availability

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 2025 re:Cap, we decided to take the date of **October 28th** as a starting date for this period.



Data and Analytics



Transform customer experiences
with Amazon Connect

Data and Analytics

Ask the Experts



About **AWS re:Invent 2025**

60.000 attendees

530 total updates

2 million viewers (*online*)

414 pre-re:Invent


3.500 sessions

116 re:Invent (Nov/30– Dec/5)

3.000 speakers




POWERING THE FUTURE



25 launches in 10 min!


STRANDS AGENTS SDK DOWNLOADS

5,299,000



Strands Agents SDK





Q&A with Matt Garman



[AWS Trainium is becoming
PyTorch native in Q1 2026](#)





Open mind for
a different view
And nothing
else matters



[The Kernel News](#)

Ready?

Fasten your seatbelts and let's re:Cap!

AWS AI League


TRANSFORM HOW BUILDERS COMPETE AND INNOVATE WITH AWS AI SERVICES



GA

Nov, 30th

re:Invent



[AWS AI League](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Updates for AWS Builder Center

REGIONAL PLANNING TOOL, WORKSHOPS, AND SPACES ARE NOW AVAILABLE FOR YOU!

aws

Search

+

⚙

☰

Sign in

Builder Center

Home

Learn

Build

Connect

Community

Heroes

Community Builders

User Groups

Cloud Clubs

Wishlist

Explore AWS

Service and features

API operations

CloudFormation resources

Services and features (39)

* All quarter dates are committed dates from the service team with potential for earlier availability as dependencies complete.

Quickly compare service and feature availability across AWS regions to eliminate architecture guesswork. Identify regional capabilities upfront, make confident deployment decisions, and accelerate your multi-region application planning with comprehensive visibility into what's available where.

Regions

N. Virginia (us-east-1)

Search capabilities

Offering name = Amazon Simple Storage Service (S3)

Show only common features

Reset all filters

Show more (+2)

Clear features selection

AWS Services

N. Virginia (us-east-1)

Seoul (ap-north-2)

Taipei (ap-east-1)

S3 Inventory Feature

Available

Available

Available

S3 Lifecycle Feature

Available

Available

Available

aws

Search

+

⚙

☰

Sign in

Builder Center

Home

Learn

Build

Connect

Community

Heroes

Community Builders

User Groups

Cloud Clubs

Wishlist

Explore AWS

AWS Builder Center

Build

Workshops

Workshops

Accelerate your AWS journey through hands-on workshops crafted by AWS experts to gain practical experience and solve real business challenges.

How it works

Find a workshop by name

Level

All levels

Language

Choose languages

Categories

Choose categories

AWS Services and Features

Choose services and features

Sort by

Newest

Get started

Aurora DSQL Immersion Day

Amazon Aurora DSQL is a completely distributed SQL database with virtually unlimited scale, the highest

aws

Search

+

⚙

☰

Sign in

Builder Center

Home

Learn

Build

Connect

Community

Heroes

Community Builders

User Groups

Cloud Clubs

Wishlist

Explore AWS

AWS Builder Center

Connect

Spaces

Spaces

Connect with builders around the world. Join spaces to share knowledge, collaborate on projects, and grow your skills together.

Create a space

How it works

Find a space by name

Country/region

All

Tag

All

Sort By

Most members

Space type

All

AWS re:Invent 2025 Chatter

aws-community

Join this space for all things re:Invent 2025! Whether you are attending virtually or are in-person in Las Vegas, this space is perfect to: • Ask and answer questions • Share advi...

561

Public

Join

AWS re:Invent 2025 Keynotes & Launches

aws-community

Join this space to chat about keynotes and launches announced during AWS re:Invent 2025. Launches are announced all week. Keynotes are streamed live December 2, 3, and ...

224


Public

Join

GA

Nov, 18th
Nov, 12th
Nov, 6th

pre-re:Invent



AWS Builder Center



Developer: Sign up to AWS Builder Center!

[AWS Builder Center](#) > [Get started](#)

Sign up. Level up.

Join AWS Builder Center to learn something new, connect with other curious minds, and expand your network in the cloud community.



Get started



01

Step 1

Activate your profile

Sign in to Builder Center with your Builder ID or create a new one.



02

Step 2

Get your alias and profile

Choose your builder alias. Then make your profile sparkle with your photo, bio, links to your social accounts, and more.



03

Step 3

Go anywhere

Explore resources, connect with builders in the AWS community, and get building!



Partners

8 announcements (3 *pre-re:Invent*)



Partner Central is now available in Console

SIMPLIFIED ACCESS AND INTEGRATION VIA APIs FOR AWS PARTNERS




GA

Nov, 30th

re:Invent

All Regions



AWS Partner Central is now available directly in the *AWS Management Console*, creating a unified experience that transforms how you engage with AWS as both customers and *AWS Partners*.



[AWS APN Blog Post](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Software Partners join Resilience programs

AWS RESILIENCE COMPETENCY EXPANDS PROGRAM FOR SOFTWARE PARTNERS



GA

Nov, 30th

re:Invent


All Regions

Why AWS Resilience Competency Partners?

AWS Resilience Competency Partners specialize in improving customer's critical workloads availability and resilience in the cloud. As each customer and their critical workloads are unique and have distinct availability requirements, AWS Resilience Competency Partners can provide guidance and solutions suitable to achieve the utmost system uptime requirements.

With AWS Resilience Competency Partners, customers can trust that the professional, consulting, or engineering services provided have been validated by AWS experts to the technical bar of its new framework in the AWS Resilience Competency Partner categories: Resilience Design, Resilience Operations, and Resilience Recovery. AWS Resilience Partners who have achieved the Core Resilience designation demonstrate the most holistic set of consulting services for customers with solutions validated by AWS in all three categories. With this standardized approach, customers can achieve their resilience in the cloud goals with the expert assistance of AWS-validated Resilience Competency Partners in a simplified, phased approach.

The AWS Resilience Competency has validated that the partners below have demonstrated they can help customers improve their resilience in the cloud.



In the **AWS Resilience Competency**, Services Partners deliver professional services for design, operation, and recovery of business critical workloads, now complemented by *Software Partners* who provide purpose-built tools to implement and maintain resilience.



[AWS APN Blog Post](#)



AI Competency adds Agentic AI categories

AI COMPETENCY EXPANDS WITH AGENTIC AI TOOLS, APPLICATIONS, AND CONSULTING

Explore AWS AI Competency Partners

From Service Partners offering professional services to Software Partners providing tools that provide customer value along all parts of the AWS AI stack, browse the Partner Solution Finder by category to quickly locate the right partner in your region for the job.

Agentic AI

Generative AI

Agentic AI

Partners in the Agentic AI category offer solutions and services that enable the safe and secure deployment of production-grade, agentic AI systems on AWS and its Agentic AI Marketplace. These systems are capable of autonomous reasoning, planning, collaboration, tool usage, execution, and self-improvement, thereby delivering measurable business value.

[Browse AWS Partners](#)

Adding three new categories for AWS Partners in the **AWS AI Competency**: *Agentic AI Applications*, *Agentic AI Tools*, and *Agentic AI Consulting Services*.

These categories recognize Software and Services Partners with deep expertise in developing and implementing enterprise-grade AI solutions that understand context, reason over goals, and act autonomously with minimal human intervention.



GA

Nov, 30th

re:Invent

All Regions



[AWS APN Blog Post](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.


Agent and AI search in AWS Marketplace

TWO NEW CAPABILITIES - AGENT MODE AND AI-POWERED SEARCH - TO ACCELERATE DISCOVERY

AWS Marketplace introduces two new AI-powered capabilities to accelerate solution discovery:

Agent mode, for a conversational discovery, comparison and proposal preparation.

AI-enhanced search, for finding right solutions fast, with AI-powered summaries to better understand your options and key considerations.



GA

Nov, 30th

re:Invent

All Regions



Multi-product solutions in AWS Marketplace

COMBINING PRODUCTS AND SERVICES FROM AWS PARTNERS FOR INDUSTRIES AND USE CASES



GA

Nov, 30th

re:Invent

All Regions

aws marketplace

Search

Hello, AWS buyer

AboutCategoriesDelivery MethodsSolutionsAWS IQResourcesYour Saved ListBecome a Channel PartnerSell in AWS MarketplaceAmazon Web Services HomeHelp

[AWS Marketplace](#) > Purchase Data Readiness Engine for Gen AI - AnyCompany

Data Readiness Engine for Gen AI - AnyCompany

Review the details on this page before purchasing. This is a multi-product offer. In the offers included table, you must review and confirm each offer individually.

Offer set details

Offer set name
Data Readiness Engine for Gen AI - AnyCompany

Offer set extended
December 1, 2025

Offer set ID
Offerset-9834029384

Offer set expires
January 1, 2026

Offer set publisher
Accenture

Number of products
4

About this offer set

The Data Readiness Engine for Gen AI solution combines Elasticsearch, the preeminent semantic and vector search engine for secure and performant AI; a full suite of enterprise connectors from Accenture and expertise for preparing multilingual content for Gen AI; and the global reach and scalability of AWS. A global integration and support team ensures that the solution is up and running quickly and efficiently. Combining best-of-breed technical solutions with deep expertise, customers can focus on building and delivering AI-powered solutions at scale.

AWS Marketplace now offers multi-product solutions - a new capability that *Partners* can use to combine multiple products and services into comprehensive solutions tailored to specific industries and customer use cases. This gives customers access to end-to-end solutions, all with the simplified procurement benefits of *AWS Marketplace*.




AWS Announcement



Billing Transfer for AWS Channel Partners

AWS CHANNEL PARTNERS CAN RESELL USING BILLING TRANSFER FOR AWS ORGANIZATIONS



With the general availability of **Billing Transfer**, a new capability to centrally manage and pay bills across multiple organizations by transferring payment responsibility to other billing administrators, it introduces a new way for centralizing billing procedures for AWS Partners.



GA

Nov, 19th

pre-re:Invent


All Regions
excluding
AWS GovCloud and
Chinese regions



[AWS News Blog Post](#)

A2A Server Support in AWS Marketplace

AWS MARKETPLACE NOW OFFERS AGENT-TO-AGENT (A2A) SUPPORT FOR 3RD PARTY AGENTS



AWS Partners can now offer A2A servers in addition to *MCP* servers and *AI agents* using *Amazon Bedrock AgentCore Runtime* containers in the **AWS Marketplace** Management Portal.



GA

Nov, 19th

pre-re:Invent

All Regions




AWS Announcement



IAM Delegation for Partners and Amazon

AWS IAM TEMPORARY DELEGATION FOR AWS PARTNER AND AMAZON PRODUCTS



AWS IAM Temporary Delegation is a new capability that streamlines customer onboarding for selected partners by automating configuration in customer AWS account through temporary, limited permissions. Customers maintain full control and visibility to manage their Partner products throughout the process.



GA

Nov, 19th

pre-re:Invent

All Regions



[AWS APN Blog Post](#)

Other Selected Announcements



1. [AWS Marketplace now supports variable payments for professional services](#) (re:Invent)
2. [AWS Marketplace introduces express private offers for fast personalized pricing](#) (re:Invent)
3. [AWS API MCP Server is now available on AWS Marketplace](#) (pre-re:Invent)
4. [AWS Network Firewall Managed Rules from AWS Partners in AWS Marketplace](#) (pre-re:Invent)
5. [AWS Marketplace now displays estimated tax and invoicing entity](#) (pre-re:Invent)
6. [AWS Marketplace now offers improvements for AI agents and tools](#) (pre-re:Invent)
7. [Amazon OCSF Ready Specialization for AWS Partners](#) (pre-re:Invent)

Storage, Migration and Transfer Services

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



Intelligent Tiering for Amazon S3 Tables

OPTIMIZE COSTS BASED ON ACCESS PATTERNS IN YOUR APACHE ICEBERG TABLES

S3 Tables Intelligent-Tiering storage class, data is automatically tiered to the most cost-effective access tier based on access patterns.

Data is stored in three low-latency tiers: **Frequent Access**, **Infrequent Access** (moved after 30 days without access, 40% lower cost than previous tier), and **Archive Instant Access** (moved after 90 days without access, 68% lower cost than previous tier).

This happens without changes to your applications or impact on performance.

Amazon S3 > Table buckets > Create table bucket

Create table bucket Info

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

General configuration

Table bucket name

my-s3-table-bucket

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 SageMaker, 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 Europe (Frankfurt) eu-central-1 Region

Disabled

☒ Enable integration

Integration details

Default storage class

S3 offers multiple storage classes designed for different use cases. The default storage class applies to all new tables created in the table bucket unless otherwise specified. [Learn more](#) or see the [Amazon pricing page](#).

Storage class

☐ Don't specify

The standard storage class is used by default by the API when storage class is not specified.

☐ Standard

Used for frequently accessed data (more than once a month) with millisecond access.

☒ Intelligent-Tiering - new

Used for data with changing or unknown access patterns. Continuously monitors access to optimize costs.



Dec, 2nd
re:Invent
All Regions



Automatic replication of Amazon S3 Tables

REPLICATE AUTOMATICALLY APACHE ICEBERG TABLES ACROSS AWS REGIONS AND ACCOUNTS

Amazon S3 Tables now support automatic replication of *Apache Iceberg* tables across AWS Regions and accounts.

This new capability replicates your complete table structure, including all snapshots and metadata to reduce query latency and improve data accessibility for global analytics workloads.

S3 Tables replication automatically creates read-only replica tables in your destination table buckets, backfills them with the latest state of the source table, and continuously monitors for new updates to keep replicas in sync.

The screenshot shows the AWS console interface for creating a table bucket replication configuration. The breadcrumb trail is: Amazon S3 > Table buckets > primary-s3-table-bucket > Create table bucket replication configuration. The page title is 'Create table bucket replication configuration'. It has two main sections: 'Destination' and 'IAM role'. The 'Destination' section explains that tables can be replicated across up to 5 table buckets in different AWS Regions (Cross-Region Replication) or in the same AWS Region (Same-Region Replication). It provides a text input for 'Table bucket ARN' with the value 'arn:aws:s3tables:eu-central-1:123456789012:bucket/amzn-s3-demo-table-bucket' and a 'Browse S3' button. Below this is an 'Add destination' button and a note: 'You can add up to 4 more table buckets.' The 'IAM role' section shows the 'IAM role selection method' with three radio buttons: 'Create new IAM role', 'Choose from existing IAM roles' (which is selected), and 'Enter IAM role ARN'. Below this is a dropdown menu for 'IAM role' with the text 'Choose IAM role' and a 'View' button. At the bottom right are 'Cancel' and 'Create replication configuration' buttons.



GA

Dec, 2nd

re:Invent

All Regions



[AWS News Blog Post](#)




General availability of Amazon S3 Vectors

GLOBAL REGIONAL COVERAGE AND 40x THE SCALE OF PREVIEW

Amazon S3 Vectors is now generally available making it ready for bigger scale - with enhanced performance and global regional coverage.

From now, S3 Vectors are available in 14 AWS regions, with read latencies for frequent queries resulting around 100ms or less, and under 1s for infrequent ones.

Additionally, you can now retrieve up to 100 results per query, and write throughput supports 1000 PUT transactions per second for a single-vector updates into your indexes.



aws [Search] [Option+S] Europe (Frankfurt)

Amazon S3 > Vector buckets > Create vector bucket

Create vector bucket [Info](#)

Vector buckets are optimized for storing and querying vectors. [Learn more](#)

General configuration

Vector bucket name

new-s3-vector-bucket

Vector bucket names must be 3 to 63 characters and unique within the account for this AWS Region. [See rules for vector bucket naming](#)

Encryption

Server-side encryption is automatically applied to new vectors stored in this vector bucket. [Learn more](#)

⚠ Encryption settings can't be changed after the vector bucket is created.

Server-side encryption

☒ Don't specify encryption type
Server-side encryption with Amazon S3 managed keys (SSE-S3) is used by default when encryption type is not specified.

☐ Specify encryption type

Tags - optional

You can organize and specify permissions for your vector buckets with tags. [Learn more](#)

No tags associated with this vector bucket.

[Add new tag](#)

You can add up to 50 tags.

[Cancel](#) [Create vector bucket](#)



GA

Dec, 2nd

re:Invent

14 Regions



[AWS News Blog Post](#)



Amazon S3 Storage Lens enhancements

PERFORMANCE METRICS, SUPPORT FOR BILLIONS OF PREFIXES, AND EXPORT TO S3 TABLES

Metrics export and publishing - optional

Metrics publishing

☐ CloudWatch publishing

Access metrics in CloudWatch without incurring separate CloudWatch metrics publishing charges. See [CloudWatch pricing](#). Prefix-level metrics aren't available in Cloudwatch.

Metrics export

Export Storage Lens dashboard data to an S3 bucket daily, for analysis with your preferred analytics tools.

Metrics export

☐ Default metrics report

This export includes prefix aggregation settings.

☒ Expanded prefixes metrics report - new

This export includes all prefixes selected in the dashboard scope.

Expanded prefixes metrics report - new

This export includes all prefixes selected in the dashboard scope.

Bucket type

☒ General purpose bucket

☐ Table bucket

General purpose bucket destination settings

Source account

☒ This AWS account

☐ A different AWS account

Destination bucket

s3://vby-test-storage-lens

[View](#) [Browse S3](#)

Your destination must be in the US East (Ohio) us-east-2 Region, where this dashboard is being created.

This announcement adds 3 new capabilities for **Amazon S3 Storage Lens** providing deeper insights into storage performance and usage patterns.

With the addition of performance metrics, support for analyzing billions of prefixes, and direct export to Amazon S3 Tables, you have the tools you need to optimize application performance, reduce costs, and make data-driven decisions about your storage strategy.



GA

Dec, 2nd

re:Invent

All Regions



[AWS News Blog Post](#)



Introducing AWS Transform custom

CRUSH TECH DEBT WITH AI-POWERED CODE MODERNIZATION

TRANSFORMATION SUMMARY

The Angular 16 to Angular 19 migration was executed through an incremental update path (16→17→18→19), adopting modern Angular 19 features including signal-based inputs/outputs, the new application builder with esbuild, functional providers, and modern control flow syntax (@if/@for). The migration involved updating all framework packages to version 19.2.15, TypeScript to 5.8.3, and modernizing code patterns while maintaining application functionality.

EXIT CRITERIA RESULTS (12/12 PASSED)

1. ✅ Application Build - Production (PASS)

****Criterion:**** Application successfully builds with `ng build --configuration production` with no errors

****Verification Method:**** Execute production build command

****Status:**** **PASS**

****Evidence:****

- Production build completed successfully in ~104 seconds (build phase: 26.6s)
- Bundle size: 308.05 kB (83.12 kB gzipped)
- Build output: "Application bundle generation complete" with no errors
- Output files:
 - main-S37LC5BQ.js: 273.47 kB (71.80 kB gzipped)

AWS Transform Custom is a first agentic AI service to accelerate transformation of Windows, mainframe, and VMware workloads.

Automates transformations reducing execution time by over 80% and includes pre-built transformations for Python, Node.js, and Java upgrades. It uses custom transformations defined using natural language, reference documents, and code samples.



GA

Dec, 1st

re:Invent

US East
(N. Virginia)



[AWS News Blog Post](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Transform expands .NET capabilities

EXPANDED .NET TRANSFORMATION CAPABILITIES AND AN ENHANCED DEVELOPER EXPERIENCE



GA

Dec, 1st

re:Invent

8 Regions

The screenshot displays the AWS Transform console interface. On the left, the 'Job Plan' sidebar lists five steps: 'Complete pre-requisites' (Completed), 'Connect resources' (Awaiting user input, highlighted with a yellow arrow), 'Run validation', 'Discover resources', and 'Assess and plan waves'. The main area shows a guidance card for 'Connect resources' with a text box explaining the next steps: 'Now that you've completed the prerequisites, let's setup a connector to your SQL Server database and source code repositories. Once the connectors are setup, I can begin assessing your databases and source code. Let's start with a connector to your SQL Server databases.' Below this are three blue buttons: 'Which AWS account should I use?', 'Why do I need to provide AWS account ID?', and 'Invite collaborators'. At the bottom of the guidance card is a text input field 'Ask a question or chat with the agent' with a send icon. To the right, the 'Database connector' configuration panel is shown. It has a 'Required' label and a description: 'Database connector connects to database instances and performs necessary steps for modernization'. There are two radio buttons: 'Choose an existing connector' (disabled) and 'Create a new connector' (selected, highlighted with a yellow arrow). Below is a text input field for 'Connector name' containing 'BobsBookstore-SQLServer' (highlighted with a yellow arrow). Another text input field for 'AWS Account ID' is empty (highlighted with a yellow arrow). A blue 'Submit' button is at the bottom (highlighted with a yellow arrow).

AWS Transform now can modernize *.NET Framework* and *.NET* code to *.NET 10* or *.NET Standard*, allows for UI porting of *ASP.NET Web Forms* to *Blazor* on *ASP.NET Core*.

New capabilities are available through an enhanced IDE experience with *AWS Toolkit for Visual Studio 2026* or *2022*, and allow for customisation of transformation plans with real-time updates and next steps guidance.



[AWS News Blog Post](#)



Attribute-Based Access Control (ABAC) for S3

GENERAL PURPOSE AMAZON S3 BUCKETS CAN LEVERAGE TAGS FOR PERMISSIONS MANAGEMENT



GA

Nov, 20th

pre-re:Invent

All Regions

Amazon S3 > Buckets > wg-abac-s3-demo-bucket > Edit bucket ABAC

Edit bucket ABAC

Bucket ABAC

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. With ABAC, you can attach tags to your general purpose buckets and AWS Identity and Access Management (IAM) entities (users or roles), then scale access to objects in your S3 general purpose buckets using tag-based policies. [Learn more](#)

ABAC status

☐ Disable

Disabling ABAC blocks the ability to manage access using tags.

☒ Enable

Enabling ABAC allows you to manage access to buckets with tags. To manage tags in this bucket, you must have s3:ListTagsForResource, s3:TagResource, and s3:UntagResource permissions.

Enabling ABAC will require permissions to s3:ListTagsForResource, s3:TagResource, and s3:UntagResource to manage your tags in the bucket. [Learn more](#)

☒ I acknowledge that enabling ABAC will allow tag-based access control to the objects in this bucket.

Cancel

Save changes

[AWS News Blog Post](#)

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

34

Post-quantum key exchange in Amazon S3

POST-QUANTUM CRYPTOGRAPHY OPTIONS FOR ENCRYPTION DATA IN-TRANSIT FOR AMAZON S3

Amazon S3 now supports post-quantum TLS key exchange on regional endpoints providing customers with post-quantum cryptography options for encryption of their data in-transit.

All regional S3, S3 Tables, and S3 Express One Zone endpoints now support *Module Lattice-Based Key Encapsulation Mechanisms (ML-KEM)*.

What is post-quantum cryptography at AWS?

At AWS, the confidentiality, integrity, and authenticity of our customers' data is a top priority. Today's widely-used public-key cryptographic schemes rely on mathematical problems - like integer factoring and discrete logarithms - that could be efficiently solved by future quantum computers. To address this challenge, AWS is deploying new NIST-standardized post-quantum cryptographic algorithms that are designed to resist both classical and quantum computing attacks. These algorithms, including the Module-Lattice-based Key-Encapsulation Mechanism (ML-KEM) and Module-Lattice-based Digital Signature Algorithm (ML-DSA), are based on different mathematical foundations that are believed to be resistant to quantum computing attacks.

AWS has already deployed post-quantum cryptography across several key services. [AWS Key Management Service \(AWS KMS\)](#), [AWS Secrets Manager](#), and [AWS Certificate Manager](#) have implemented post-quantum hybrid key establishment combining Elliptic Curve Diffie-Hellman (ECDH) with ML-KEM to protect against "harvest now, decrypt later" attacks. At the foundation of these implementations is [AWS-LC, our FIPS-140-3-validated cryptographic library](#), which was the first open-source cryptographic module to include ML-KEM in its FIPS validation.



GA

Nov, 19th
pre-re:Invent
All Regions



AWS Announcement



Air-gapped **AWS Backup** vault as a primary

NOW YOU CAN USE AIR-GAPPED VAULTS AS A PRIMARY TARGET IN AWS BACKUP



GA

Nov, 17th

pre-re:Invent

All Regions

[Documentation](#) > [AWS Backup](#) > [Developer Guide](#)

Primary backups to logically air-gapped vaults

PDF

RSS



Focus mode

Overview

The logically air-gapped vault primary backup feature offers you the option to specify a logically air-gapped vault as a primary backup destination within the same account, for both scheduled and on-demand backup jobs. This eliminates the need to maintain separate copies in both a standard backup vault and a logically air-gapped vault, reducing costs and simplifying workflows while preserving the security benefits of logical air-gapping.

You can assign a logically air-gapped vault as the primary target in backup plans, organization-wide policies, or on-demand backups. Previously, to back up to a logically air-gapped vault, you first had to create a backup in a backup vault and then copy it to a logically air-gapped vault. With this feature, depending on the resource type, AWS Backup can create backups directly in your logically air-gapped vault or automatically manage temporary backups that are copied to your logically air-gapped vault and then deleted.



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Announcement](#)

AWS Backup now supports Amazon EKS

FULLY MANAGED EXPERIENCE FOR BACKING UP EKS CLUSTER STATE AND PERSISTENT APP DATA



GA

Nov, 10th

pre-re:Invent

All Regions

composite:eks/eks-backup-test-pvs-20251030203644-5bebbe2d - EKS [Info](#) Copy Delete Restore

Details

ARN
[arn:aws:backup:us-east-1:047481191117:recovery-point:composite:eks/eks-backup-test-pvs-20251030203644-5bebbe2d](#)

Creation time
October 30, 2025, 13:36:44 (UTC-07:00)

Resource type
EKS

Resource name
EKS-Backup-Test-PVs

Status
✔ Completed

Resource ID
cluster/EKS-Backup-Test-PVs

Backup type
Composite

Storage tier
Warm

Backup summary

Backup type
Manual

Expiration date
December 4, 2025, 12:36:44 (UTC-08:00)
35 days after creation

Backup vault
[Default](#)

Move to cold date
-

Backup plan
-

IAM role
[ekstesting](#)

Edit

Nested recovery points (3) [Info](#)

Scan Copy Delete Disassociate Restore

Recovery point ID

Status

Resource name

Resource ID

Resource type

Backup type

[getbackexportstore-20251030203707-4a42ec49](#)

✔ Completed

getbackexportstore

getbackexportstore

S3

Snapshot

[snapshot/snap-017ca0b7ba1de3379](#)

✔ Completed

EKS-Backup-Test-PVs-dynamic-pvc-68749721-5048-4f8a-ae14-7aeae9fc94e9

volume/vol-0be4b3822f00c6812

EBS

Snapshot

[eks/eks-backup-test-pvs-20251030203706-df6a59f1](#)

✔ Completed

EKS-Backup-Test-PVs

cluster/EKS-Backup-Test-PVs

EKS

Backup

Backup encryption [Info](#)

Encryption status
✔ Encrypted

[AWS News Blog Post](#)

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

37


KMS CMK for air-gapped AWS Backup vault

AWS BACKUP NOW SUPPORTS AWS KMS CUSTOMER MANAGED KEYS FOR AIR-GAPPED VAULTS



GA

Nov, 6th
pre-re:Invent
All Regions



Single-action snapshot copy in AWS Backup

COPYING DATABASE SNAPSHOTS ACROSS REGIONS AND AWS ACCOUNTS WITH SINGLE ACTION



GA

Oct, 30th

pre-re:Invent

All Regions

AWS Backup now supports copying database snapshots across AWS Regions and accounts using a single copy action for *Amazon RDS, Amazon Aurora, Amazon Neptune, and Amazon DocumentDB* snapshots.

For information on opt-in Regions and what resources and features are supported within, see [Feature availability by AWS Region](#).

AWS Backup supports	Cross-Region backup	Cross-account backup	AWS Backup Audit Manager	Incremental backup	Continuous backup and point-in-time restore	Full management
Amazon S3	✓	✓	✓	✓	✓	✓
Amazon EBS	✓	✓	✓	✓		
Amazon RDS single instance	✓ ³	✓ ³	✓ ⁴	✓	✓	
Amazon RDS cluster	✓ ³	✓ ³	✓ ⁴	✓		
Amazon Aurora	✓ ³	✓ ³	✓	✓ ⁶	✓	
Amazon DocumentDB	✓ ³	✓ ³				
Amazon Neptune	✓ ³	✓ ³				
Amazon DynamoDB			✓			
Amazon EKS ¹⁰	✓	✓				✓

³ Amazon RDS, Aurora, DocumentDB, and Neptune now support cross-Region and cross-account snapshot copying in a single action. RDS multi availability zone (Multi-AZ) database instances can be copied, but Multi-AZ clusters do not currently support cross-Region or cross-account copy. See [Cross-Region copy considerations with specific resources](#) for further information.



AWS Announcement



Conditional write in copy for Amazon S3

EXPANDING CONDITIONAL WRITE FUNCTIONALITY TO COPY OPERATIONS



GA

Oct, 29th

pre-re:Invent

All Regions

[Documentation](#) > [Amazon Simple Storage Service \(S3\)](#) > [User Guide](#)

Add preconditions to S3 operations with conditional requests

[PDF](#) [RSS](#) ☒ Focus mode

You can use conditional requests to add preconditions to your S3 operations. To use conditional requests, you add an additional header to your Amazon S3 API operation. This header specifies a condition that, if not met, will result in the S3 operation failing.

Conditional reads are supported for `GET`, `HEAD`, and `COPY` requests. You can add preconditions to return or copy an object based on its Entity tag (ETag) or last modified date. This can limit an S3 operation to objects updated since a specified date. You can also limit an S3 operation to a specific ETag. This could ensure you only return or copy a specific object version. For more information about the object metadata, see [Working with object metadata](#).

Conditional writes can ensure there is no existing object with the same key name in your bucket during `PUT` operations. This prevents overwriting of existing objects with identical key names. Similarly, you can use conditional writes to check if an object's ETag is unchanged before updating the object. This prevents unintentional overwrites on an object without knowing the state of its content. You can use conditional writes for [PutObject](#), [CompleteMultipartUpload](#), or [CopyObject](#) requests. For more information about key names, see [Naming Amazon S3 objects](#).

Amazon S3 expands conditional write functionality to *copy operations* for general purpose and directory buckets.

With conditional copy, you can now verify if the object exists or has been modified in your destination S3 bucket before copying it.



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Announcement](#)

Other Selected Announcements (1/2)



1. Amazon FSx for NetApp ONTAP now supports Amazon S3 access (re:Invent)
2. Amazon S3 increases the maximum object size to 50 TB (re:Invent)
3. Amazon S3 Batch Operations introduces performance improvements (re:Invent)
4. **AWS Transform:**
 - a) Launches an AI agent for full-stack Windows modernization (re:Invent)
 - b) Adds new agentic AI capabilities for enterprise VMware migrations (re:Invent)
 - c) Adds new testing automation capabilities for mainframe modernization (re:Invent)
 - d) Mainframe now supports application reimaging (re:Invent)
5. Amazon S3 Block Public Access now supports organization-level enforcement (pre-re:Invent)
6. Amazon S3 Metadata expands to 22 additional AWS Regions (pre-re:Invent)
7. AWS Backup now supports Amazon FSx Intelligent-Tiering (pre-re:Invent)
8. Recycle Bin adds support for Amazon EBS Volumes (pre-re:Invent)
9. Amazon S3 adds new bucket-level setting to standardize encryption types (pre-re:Invent)
10. AWS Backup launches a low-cost warm storage tier for Amazon S3 backups (pre-re:Invent)


Other Selected Announcements (2/2)



11. [AWS Transfer Family announces Terraform module to automate scanning](#) (pre-re:Invent)
12. [Amazon FSx for Lustre improves directory listing performance by up to 5x](#) (pre-re:Invent)
13. [Amazon S3 now supports IPv6 for gateway and interface VPC endpoints](#) (pre-re:Invent)
14. [Amazon S3 Express One Zone now supports IPv6](#) (pre-re:Invent)
15. [Amazon FSx now integrates with AWS Secrets Manager for Active Directory credentials](#) (pre-re:Invent)


Compute and Containers

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



Introducing Amazon EC2 m9g with Graviton5

NEW GENERAL PURPOSE AMAZON EC2 INSTANCES POWERED BY AWS GRAVITON5 PROCESSORS



P

Dec, 4th

re:Invent



Preview

AWS Graviton5/Amazon EC2 M9g instances

AWS Graviton5/Amazon EC2 M9g instances is now available in preview. You will need an AWS Account Number in order to request access. Please submit the information below to request an invitation to the preview. We will contact you with instructions if you are accepted.

The newest general purpose **Amazon EC2 m9g instances** powered by AWS *Graviton5* offer up to 25% better compute performance, and higher networking and Amazon EBS bandwidth than *Graviton4*-based *m8g* instances.

They are up to 30% faster for databases, up to 35% faster web applications, and up to 35% faster for machine learning workloads compared to *m8g*.




[AWS Announcement](#)



Amazon EC2 trn3 UltraServers

FASTER, LOWER-COST GENERATIVE AI TRAINING POWERED BY TRAINIUM3 CHIPS



P

Dec, 2nd

re:Invent



Amazon EC2 Trn3 UltraServer

Up to 144 Trainium3 Chips | 362 FP8 PFLOPS | 706 TB/s aggregate bandwidth | 1M+ Token Context Windows

4.4x higher performance | 3.9x higher memory bandwidth | 3.5x higher tokens/Megawatt




© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Announcement](#)

Amazon EC2 trn4

LATEST GENERATION AI SILICON DESIGNED TO POWER NEXT-GEN AI MODELS



CS


Dec, 2nd

re:Invent



Amazon EC2 x8aedz instances

NEW GENERATION MEMORY OPTIMIZED INSTANCES POWERED BY 5TH GEN AMD EPYC (TURIN)



GA

Dec, 2nd
re:Invent
2 Regions

Instance name	vCPUs	Memory (GiB)	NVMe SSD storage (GB)	Network bandwidth (Gbps)	EBS bandwidth (Gbps)
x8aedz.large	2	64	158	Up to 18.75	Up to 15
x8aedz.xlarge	4	128	316	Up to 18.75	Up to 15
x8aedz.3xlarge	12	384	950	Up to 18.75	Up to 15
x8aedz.6xlarge	24	768	1,900	18.75	15
x8aedz.12xlarge	48	1,536	3,800	37.5	30
x8aedz.24xlarge	96	3,072	7,600	75	60
x8aedz.metal-12xl	48	1,536	3,800	37.5	30
x8aedz.metal-24xl	96	3,072	7,600	75	60

Performance Benefits

Highest 5GHz CPU frequency
Up to 2x compute performance

Memory Optimization

32:1 memory to vCPU ratio
Up to 3,072 GiB memory

Enterprise Features

Sixth-generation AWS Nitro Cards
Up to 75 Gbps network bandwidth




[AWS News Blog Post](#)



Amazon EC2 p6e-GB300 UltraServers

NEW ULTRASERVERS POWERED BY NVIDIA GB300 NVL72 RACKS



P

Dec, 2nd
re:Invent

AWS is the best place to run NVIDIA GPUs

HIGHEST RELIABILITY & AVAILABILITY FOR GPU-BASED SYSTEMS

P6e

GB200 [NVL72]

GB300 [NVL72] NEW

P6

B200

B300





© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Announcement](#)

48

Announcing Amazon EKS Capabilities

FULLY MANAGED AND EXTENSIBLE SET OF K8S-NATIVE PLATFORM FEATURES (ARGO, ACK, KRO)




GA

Nov, 30th

re:Invent

All Regions



Platform Features

Continuous deployment with Argo CD
AWS resource management through ACK

Developer Experience

Ready-to-use platform capabilities
Faster workload deployment

Operational Benefits

AWS-owned infrastructure management
Auto scaling and patching



[AWS News Blog Post](#)



Interruptible Capacity Reservations

BETTER UTILIZATION OF YOUR RESERVED CAPACITY AND REDUCED COSTS

aws

Search

[Option+S]

United States (Oregon)

EC2

Capacity Reservations

Capacity Reservations (1)

Info

Actions

Create Capacity Reservations


Type

All reservation types

Filter Capacity Reservations by text, property or value

ID	Status	Reservation type	Total instan...	Available instanc...	Instance type	Interruptible	Availability Zone		
☐	☐	cr-0300b3a3cc7566f90	Active	ODCR	5 instances	5 instances	m5a.4xlarge	No	us-west-2a

Capacity Reservations overview



GA

Nov, 24th

pre-re:Invent

5 Regions



AWS Announcement



Provisioned Control Plane in Amazon EKS

ENSURE PREDICTABLE AND HIGH PERFORMANCE FOR THE MOST DEMANDING WORKLOADS



GA

Nov, 21st

pre-re:Invent

All Regions

Manage control plane scaling tier - new

Control plane scaling tier [Info](#)

Choose the capacity required for your cluster's Kubernetes control plane. By default, Amazon EKS automatically scales your cluster's control plane as needed. You can use a scaling tier for more control over its capacity. [View tier pricing](#)

☒ Use a scaling tier
For predictable and high performance from your cluster's control plane, choose a scaling tier.

XL

API request concurrency
1700 seats

Pod scheduling rate
100 pods per second

Cluster database size
16 GB

2XL

API request concurrency
3400 seats

Pod scheduling rate
100 pods per second

Cluster database size
16 GB

4XL

API request concurrency
6800 seats

Pod scheduling rate
100 pods per second

Cluster database size
16 GB

Cancel

Save changes

Amazon EKS Provisioned Control Plane gives you the ability to pre-allocate control plane capacity from a set of new scaling tiers, ensuring predictable and high performance for your most demanding workloads. By provisioning capacity ahead of time, your cluster handles instantaneous traffic bursts without the need for control plane scaling.

A square QR code that, when scanned, likely leads to the AWS Containers Blog Post mentioned in the footer.

The AWS logo, consisting of the letters 'aws' in a stylized font with a curved line underneath.

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Containers Blog Post](#)

51

Amazon ECS Express Mode

STREAMLINE DEPLOYMENT AND MANAGEMENT OF CONTAINERS WITH AWS BEST PRACTICES

Successfully created Express service hello-world-f798 in cluster default
A deployment will begin automatically. Your service will be accessible after deployment completes and resources are provisioned. Check the Resources tab to monitor provisioning progress.

View service

hello-world-f798 Express

Last updated November 8, 2025, 11:43 (UTC+8:00)

Delete service Update service

Express service overview

Service ARN
arn:aws:ecs:south-1-:service/default/hello-world-f798

Application URL
-

Service status
Active

Active deployment
-

Created at
November 8, 2025, 11:43 (UTC+8:00)

Updated at
November 8, 2025, 11:43 (UTC+8:00)

Observability Logs Resources Deployments Configuration Tags

Table view Timeline view

Resources

Your AWS resources for this service. Updates automatically.

Cluster: default

Service: hello-world-f798 Active

Target service revision: 4964277530704639669

Task definition: hello-world-f798:1

Service security group: sg-0666a448970cfc159 Active

Public ingress path: hello-world-f798-8e2f73f1.ecs.af-south-1.on.aws

Certificate: hello-world-f798-8e2f73f1.ecs.af-south-1.on.aws Active

Load balancer security group: sg-0af0367a5123bf3e1 Active

Load balancer: app/ecs-route-ee5e83fb-e6f7-477e-996/a1d92898fc0955cd Active

Listener: app/ecs-route-ee5e83fb-e6f7-477e-996/a1d92898fc0955cd/764169341c81a363 Active

Target group: ecs-tg-b0cd30c82152070d/ea18a248a08a16d7 Active

Listener rule: app/ecs-route-ee5e83fb-e6f7-477e-996/a1d92898fc0955cd/764169341c81a363/c3f4ddddd3893c996 Active

Scaling target: service-default-hello-world-f798 Active

Scaling policy: default-hello-world-f798-AVERAGE_CPU-70 Active

Deployment mn02o5KJxDFww6JozZfw4 is in progress...

Deploy containerized applications on AWS aligned with Well-Architected best practices via a single command with **Amazon ECS Express Mode**.

And when your applications evolve and require advanced features, you can seamlessly configure and access the full capabilities of AWS services and the resources - including *Amazon ECS*.



GA

Nov, 21st

pre-re:Invent

All Regions



[AWS News Blog Post](#)



Fully managed MCP servers on ECS and EKS

ENTERPRISE-GRADE CAPABILITIES FOR MCP SERVERS HOSTED ON AMAZON ECS OR AMAZON EKS




Nov, 21st

pre-re:Invent

All Regions

```
mcp-config.json x
{} mcp-config.json > {} mcpServers >
1 {
2   "mcpServers": {
3     "eks-mcp": {
4       "disabled": false,
5       "type": "stdio",
6       "command": "uvx",
7       "args": [
8         "mcp-proxy-for-aws@latest",
9         "https://eks-mcp.us-west-2.api.aws/mcp",|
10      "--service",
11      "eks-mcp",
12      "--profile",
13      "default",
14      "--region",
15      "us-west-2"
16      ]
17     }
18   }
19 }
```



[AWS Containers Blog Post](#)



AI-powered experience for ECS and EKS

AMAZON Q DEVELOPER IN AWS MANAGEMENT CONSOLE FOR AMAZON ECS AND AMAZON EKS



GA

Nov, 21st

pre-re:Invent

All Regions

Tasks (21)

Last updated November 15, 2025, 20:34 (UTC-8:00)

Manage tags

Stop

Run new task

Filter tasks by property or value

Filter desired status Any desired status

Filter launch type Any launch type

< 1 >

☐	Task	Last status	Desired st...	Task definition	Health sta...	Created at	Started by
☐	 6bd86a67...	 Pending	 Running	sample-ui:1	 Unknown	3 minutes ago	ecs-svc/0831507
☐	 9f7a3029...	 Pending	 Running	sample-ui:1	 Unknown	3 minutes ago	ecs-svc/0831507
☐	 ad6b9be6...	 Pending	 Running	sample-appserver:1	 Unknown	5 minutes ago	ecs-svc/4935780
☐	 af018b04...	 Pending	 Running	sample-appserver:1	 Unknown	5 minutes ago	ecs-svc/4935780
☐	 b1b07bf2...	 Pending	 Running	sample-db:1	 Unknown	2 minutes ago	ecs-svc/4182217
☐	 b5b084b2...	 Pending	 Running	sample-ui:1	 Unknown	3 minutes ago	ecs-svc/0831507
☐	 09411f35...	 Stopped ...	 Stopped	sample-appserver:1	 Unknown	31 minutes ago	ecs-svc/4935780
☐	 0b972fbb...	 Stopped ...	 Stopped	sample-ui:1	 Unknown	30 minutes ago	ecs-svc/0831507
☐	 19a28c0d...	 Stopped ...	 Stopped	sample-ui:1	 Unknown	57 minutes ago	ecs-svc/0831507
☐	 26be4a8f...	 Stopped ...	 Stopped	redis-server:1	 Unknown	28 minutes ago	ecs-svc/6110215
☐	 2b92527c...	 Stopped ...	 Stopped	sample-db:1	 Unknown	29 minutes ago	ecs-svc/4182217



Managed image signing for Amazon ECR

ENHANCE YOUR SECURITY POSTURE WITH FULLY MANAGED CONTAINER IMAGE SIGNING

[Documentation](#) > [Amazon ECR](#) > [User Guide](#)

Managed signing

PDF

RSS

Focus mode

Amazon ECR managed signing automatically signs your container images by generating cryptographic signatures using [AWS Signer](#) when images are pushed to Amazon ECR. This eliminates the need to install and configure client-side tools and allows you to centrally govern signing as a registry configuration.

Prerequisites

To configure managed signing, you create a signing configuration with Amazon ECR that references one or more Signer signing profiles and, optionally, repository filters that restrict which repositories should have their images signed. Once configured, Amazon ECR managed signing automatically signs images as they are pushed using the identity of the entity pushing the image.

Before you can configure managed signing, you must have the following:

- **A Signer signing profile** — Create at least one Signer [signing profile](#). A signing profile is a unique AWS Signer resource that you can use to perform signing operations in Amazon ECR. Signing profiles enable you to sign and verify code artifacts, such as container images and AWS Lambda deployment bundles. Each signing profile designates the signing platform to sign for, a platform ID, and other platform-specific information. For example, a signing profile ARN looks like this:
`arn:partition:signer:region:account-id:/signing-profiles/profile-name`.
- **IAM permissions** — The IAM principal that pushes the image must have the necessary IAM permissions to access the relevant Signer signing profile and the relevant ECR repository. You need to modify the identity-based policy for the IAM principal to include permissions for both ECR repository operations and Signer signing operations.



GA

Nov, 21st

pre-re:Invent

All Regions



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Auto-versioning for EC2 Image Builder

AUTO-VERSIONING AND ENHANCED INFRASTRUCTURE AS CODE DEVEX FOR EC2 IMAGE BUILDER



GA

Nov, 21st

pre-re:Invent

All Regions

```
TS L2-aws-cdk-ec2-image-builder.ts 3 X
> TS L2-aws-cdk-ec2-image-builder.ts > ...
4
5 const exampleContainerRecipe = new imagebuilder.ContainerRecipe(this, 'MyContainerRecipe', {
6   baseImage: imagebuilder.BaseContainerImage.fromDockerHub('amazonlinux', 'latest'),
7   targetRepository: imagebuilder.Repository.fromEcr(
8     ecr.Repository.fromRepositoryName(this, 'Repository', 'my-container-repo')
9   )
10 });
11
12 const containerPipeline = new imagebuilder.ImagePipeline(this, 'MyContainerPipeline', {
13   recipe: exampleContainerRecipe,
14   schedule: {
15     expression: events.Schedule.cron({
16       minute: '0',
17       hour: '6',
18       weekday: 'MON'
19     })
20   }
21 });
```

[Documentation](#) > [EC2 Image Builder](#) > [User Guide](#)

Semantic versioning in Image Builder

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

Image Builder uses semantic versioning to organize resources and ensure that they have unique IDs. The semantic version has four nodes:

`<major>.<minor>.<patch>/<build>`

You can assign values for the first three, and can filter on all of them.

Semantic versioning is included in each object's Amazon Resource Name (ARN), at the level that applies to that object as follows:

1. Versionless ARNs and Name ARNs do not include specific values in any of the nodes. The nodes are either left off entirely, or they are specified as wildcards, for example: x.x.x.
2. Version ARNs have only the first three nodes: <major>.<minor>.<patch>
3. Build version ARNs have all four nodes, and point to a specific build for a specific version of an object.

Assignment: For the first three nodes you can assign any positive integer value, or zero, with an upper limit of $2^{30}-1$, or 1073741823 for each node. Image Builder automatically assigns the build number to the fourth node.

Patterns: You can use any numeric pattern that adheres to the assignment requirements for the nodes that you can assign. For example, you might choose a software version pattern, such as 1.0.0, or a date, such as 2021.01.01.

Selection: With semantic versioning, you have the flexibility to use wildcards (x) to specify the most recent versions or nodes when selecting the base image or components for your recipe. When you use a wildcard in any node, all nodes to the right of the first wildcard must also be wildcards.




[AWS Announcement](#)



Instance Lifecycle Policy for EC2 Auto Scaling

DEFINE HOW TO RETAIN INSTANCES DURING TERMINATION FAILS FOR GRACEFUL SHUTDOWNS



GA

Nov, 20th

pre-re:Invent

4 Regions

EC2 Auto Scaling announces a new feature called instance lifecycle policy.

Customers can now configure a way to retain their instances when their termination lifecycle hooks fail or timeout, providing greater confidence in managing instances for graceful shutdown.

[Documentation](#) > [Auto Scaling](#) > [User Guide](#)

Control instance retention with instance lifecycle policies

[PDF](#)

[RSS](#)

☐ Focus mode

Prevent instance termination when lifecycle hooks fail or timeout. Instance lifecycle policies give you control over what happens when termination lifecycle hooks don't complete successfully.

When to use instance lifecycle policies

Use lifecycle policies when you need guaranteed graceful shutdown for:

- Stateful applications that must complete data persistence before termination
- Applications requiring extended draining periods (beyond 48 hours)
- Workloads handling sensitive data that needs proper cleanup
- Mission-critical services where abrupt shutdown causes availability impact

How instance lifecycle policies work

Instance lifecycle policies control what happens when termination lifecycle hooks fail, timeout, or are abandoned by calling `CompleteLifecycleAction`. Without a policy, Amazon EC2 Auto Scaling terminates instances immediately when hooks fail. With a retention policy, instances move to a Retained state where they remain until you manually authorize termination.

Retained instances:

- Continue running as active Amazon EC2 instances
- Don't count toward your Amazon EC2 Auto Scaling group's desired capacity
- Incur standard Amazon EC2 charges until terminated
- Remain in Retained state indefinitely until you take action



Enhanced Network Observability in EKS

DEEPER INSIGHTS INTO YOUR CONTAINER NETWORKING ENVIRONMENT ON AMAZON EKS




GA

Nov, 19th

pre-re:Invent

All Regions



[AWS News Blog Post](#)



Archive storage class for Amazon ECR

REDUCE STORAGE COSTS FOR LARGE VOLUMES OF RARELY USED IMAGES



GA

Nov, 19th

pre-re:Invent

All Regions

Create lifecycle rule

Destructive action warning

Amazon ECR recommends using a Lifecycle Policy Preview to determine the immediate impact of adding a new rule will have on your existing images. See [Lifecycle Policy Preview](#) to run a preview with a new rule.

Rule priority

Specify a rule priority. Rule priority sets the order in which rules are applied.

2

Rule priority must be a positive integer which is unique in lifecycle policy.

Rule description - optional

Specify a rule description.

Move images to archive if image is not pulled for 30 days (based on the last recorded pull time).

Image tag status

Specify image tag status for which images will be evaluated.

☒ Tagged (wildcard matching)

☐ Tagged (prefix matching)

☐ Untagged

☐ Any

Image tag filters

List of image tag filters matching on image tag substring with wildcards. Wildcard (*) matches zero or more characters.

Add filter

Filters must only contain letters, numbers, and special characters (.,_*,-). Each filter is limited 4 wildcards (*).

Match criteria

Specify the type of matching.

☐ Days since image created

☒ Days since last recorded pull time

☐ Days since image archived

☐ Image count

Days before action

Specify number of days since image last recorded pull time before action taken.

30

Rule action

Specify the action taken when rule criteria is met


☐ Expire (not supported for 'days since image last recorded pull time')
Deletes the image

☒ Archive
Moves image into archive storage class



Supplementary Packages for Amazon Linux

SUPPLEMENTARY PACKAGES FOR AMAZON LINUX (SPAL) EPEL9 REPO WITH 1000+ PACKAGES



GA

Nov, 18th

pre-re:Invent

Open Source

aws

Search

[Option+S]

Europe (Ireland)

CloudShell

eu-west-1

~ \$ dnf repolist --all

repo id

amazonlinux

amazonlinux-debuginfo

amazonlinux-source

repo name

Amazon Linux 2023 repository

Amazon Linux 2023 repository - Debug

Amazon Linux 2023 repository - Source packages

status

enabled

disabled

disabled

~ \$ sudo dnf install ffmpeg

Last metadata expiration check: 0:04:27 ago on Tue 02 Dec 2025 09:57:50 AM UTC.

No match for argument: ffmpeg

Error: Unable to find a match: ffmpeg

~ \$ dnf search ffmpeg

Last metadata expiration check: 0:02:03 ago on Tue 02 Dec 2025 10:00:17 AM UTC.

===== Summary Matched: ffmpeg =====

nv-codec-headers.noarch : Ffmpeg version of Nvidia Codec SDK headers

~ \$

Feedback


© 2025, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences



AWS Announcement




Other Selected Announcements (1/2)



1. [Announcing Amazon EC2 Compute-optimized C8a instances](#) (re:Invent)
2. [Announcing Amazon EC2 Compute-optimized C8ine instances in preview](#) (re:Invent, **P**)
3. [Announcing Amazon EC2 M4 Max Mac instances in preview](#) (re:Invent, **P**)
4. [Announcing Amazon EC2 Memory-optimized X8i instances in preview](#) (re:Invent, **P**)
5. [Announcing Amazon EC2 General purpose M8azn instances in preview](#) (re:Invent, **P**)
6. [Amazon EKS add-ons now supports the AWS Secrets Store CSI Driver provider](#) (pre-re:Invent)
7. [Amazon CloudWatch Introduces In-Console Agent Management on EC2](#) (pre-re:Invent)
8. [Announcing flexible AMI distribution capabilities for EC2 Image Builder](#) (pre-re:Invent)
9. [Amazon EC2 Fleet adds new encryption attribute for instance type selection](#) (pre-re:Invent)
10. [Amazon EC2 Auto Scaling supports root volume replacement through instance refresh](#) (pre-re:Invent)
11. [Amazon EC2 introduces AMI ancestry for complete AMI lineage visibility](#) (pre-re:Invent)
12. [Amazon EC2 Mac instances now support Apple macOS Tahoe](#) (pre-re:Invent)
13. [Amazon EC2 P6-B300 instances with NVIDIA Blackwell Ultra GPUs are now available](#) (pre-re:Invent)
14. [EC2 Auto Scaling now offers a synchronous API to launch instances inside an ASG](#) (pre-re:Invent)

Other Selected Announcements (2/2)



15. Amazon ECS:

- a) Managed Instances adds configurable scale-in delay (pre-re:Invent)
 - b) Improves Service Availability during Rolling deployments (pre-re:Invent)
 - c) Amazon ECS announces non-root container support for managed EBS volumes (pre-re:Invent)
 - d) Service Connect enhances observability with Envoy Access Logs (pre-re:Invent)
 - e) Supports built-in Linear and Canary deployments (pre-re:Invent)
- 16. Amazon Braket introduces spending limits feature for quantum processing units (pre-re:Invent)
 - 17. Amazon Braket adds new quantum processor from Alpine Quantum Technologies (pre-re:Invent)
 - 18. Amazon EC2 reduces costs for Microsoft SQL Server High-Availability deployments (pre-re:Invent)
 - 19. Announcing EC2 Image Builder support for Lambda and Step functions (pre-re:Invent)
 - 20. Announcing New EC2 R8a Memory-Optimized Instances (pre-re:Invent)
 - 21. EC2 Auto Scaling announces warm pool support for ASGs with mixed policies (pre-re:Invent)
 - 22. Announcing larger instances for Amazon Lightsail (pre-re:Invent)
 - 23. AWS Elastic Beanstalk adds support for Amazon Corretto 25 and Tomcat 11 (pre-re:Invent)

Databases

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



Database Savings Plans

FLEXIBLE PRICING MODEL THAT SAVES UP TO 35% FOR A CONSISTENT USAGE COMMITMENT

Recommendations

Info

Refresh recommendations

Recommendation parameters

Info

Savings Plans type

☐ Compute Savings Plans

Applies to EC2 instance usage, AWS Fargate, and AWS Lambda service usage, regardless of Region, instance family, size, tenancy, and operating system.

☒ Database Savings Plans

Applies to Aurora, RDS, DynamoDB, ElastiCache, Amazon DocumentDB, Neptune, and other services regardless of supported engine, instance family, size, deployment option, or Region.

☐ EC2 Instance Savings Plans

Applies to instance usage within the committed EC2 family and Region, regardless of size, tenancy, and operating system.

☐ SageMaker Savings Plans

Applies to SageMaker service usage, regardless of Region, instance family, and component.

Recommendation options

Savings Plans term

☒ 1-year

☐ 3-year

Payment option

☐ All upfront

☐ Partial upfront

☒ No upfront


Based on the past

☐ 7 days

☒ 30 days

☐ 60 days

Database Savings Plans, extend *Savings Plans* flexible pricing model to a wide portfolio of AWS managed database services to help customers reduce database costs by up to 35%, when they commit to a consistent amount of usage (\$/hour) over a 1-year term.



Dec, 2nd
re:Invent
All Regions



[AWS News Blog Post](#)



Vectors performance boosts in OpenSearch

GPU-ACCELERATED AND AUTO-OPTIMIZED VECTOR INDEXES IN AMAZON OPENSEARCH SERVICE




GA

Dec, 2nd

re:Invent

5 Regions

9 Regions



Amazon OpenSearch Service received a serverless GPU acceleration and auto-optimization for vector indexes. *GPU acceleration* helps you build large-scale vector databases with lower net costs and time savings through improved performance and more efficient infrastructure utilization (a quarter of the indexing cost comparing to non-GPU acceleration for large scale). *Auto-optimization* for vector indexes finds optimal trade-offs between search quality, speed, and cost. This optimization happens without needing any vector expertise from your end.



[AWS News Blog Post](#)



Agentic Search in OpenSearch Service

AN AGENT-DRIVEN SYSTEM THAT BASED ON USER'S INTENT GENERATE QUERIES AND SUMMARIES



GA

Nov, 25th

pre-re:Invent

All Regions

With the availability of *OpenSearch* version 3.3, **Amazon OpenSearch Service** launches *Agentic Search*, transforming how users interact with their data through intelligent, agent-driven search.

Test flow

Index View details

sports_and_outdoors

Query Form JSON Search

highest rated golf clubs between \$200 and \$300

[Add query fields](#)

Generated query View details

```
{
  "query": {
    "bool": {
      "filter": [
        {
          "range": {
            "price": {
              "include_lower": true,
              "include_upper": true,
              "from": 200,
              "boost": 1,
              "to": 300
            }
          }
        }
      ],
      "adjust_pure_negative": true,
      "must": [
        {
          "match_phrase": {
            "title": {
              "query": "golf club",
              "zero_terms_query": "NONE",
              "boost": 1,
              "slop": 0
            }
          }
        }
      ]
    },
    "boost": 1
  },
  "sort": [
    {
      "average_rating": {
        "order": "desc"
      }
    }
  ]
}
```



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Dynamic data masking in Aurora PostgreSQL

LEVERAGE THE POWER OF PG_COLUMNMASK INSIDE AMAZON AURORA CLUSTERS




GA

Nov, 24th

pre-re:Invent

All Regions



[AWS Database Blog Post](#)



New Connectors for Amazon Aurora DSQL

PYTHON, NODE.JS, AND JDBC CONNECTORS THAT SIMPLIFY IAM AUTHENTICATION FOR DSQL



GA

Nov, 21st

pre-re:Invent

All Regions

[Documentation](#) > [Amazon Aurora DSQL](#) > [User Guide](#)

Connectors for Aurora DSQL

↓ PDF

↓ RSS

🔵 Focus mode

Aurora DSQL provides specialized connectors that extend existing database drivers to enable seamless IAM authentication and integration with AWS services. These connectors are designed to work with popular programming languages and frameworks while maintaining compatibility with existing PostgreSQL workflows.

Amazon Aurora DSQL provides a set of new connectors that simplify AWS IAM authorization for customers using standard *PostgreSQL* drivers - including *Python*, *Node.js* and *JDBC* driver. These connectors act as transparent authentication layers that automatically handle IAM token generation.



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Query Editor for Amazon Aurora DSQL

AWS MANAGEMENT CONSOLE NOW PROVIDES INTEGRATED QUERY EDITOR FOR AURORA DSQL



GA

Nov, 21st

pre-re:Invent

All Regions

Aurora DSQL

Clusters

Query Editor

Clusters in US East (N. Virginia) (1) [Info](#)

Find a cluster

Tag: Name	Cluster ID	Status	Peer cluster L	Peer witness
☐ shopping-list-dsql-cluster	xuabtzm7eogeabow46flgpmwuu	Active	US East (Ohio) - us-east-2	US West (Oregon) - us-west-2

Connect with Query editor

Actions

Create cluster

CloudShell

Feedback

© 2025, Amazon Web Services, Inc. or its affiliates.

Privacy

Terms

Cookie preferences



[AWS Announcement](#)



GA

pre-re:Invent



Blue/Green deployments in Aurora Global

AMAZON RDS BLUE/GREEN DEPLOYMENTS ARE NOW SUPPORTED IN AMAZON AURORA GLOBAL



GA

Nov, 14th

pre-re:Invent

All Regions

Aurora and RDS > Databases

Databases (11)

Filter by databases

Group resources

Modify

Actions

Create database

DB identifier	Status	Role	Engine	Version	Region	Size
gdb-prod Available Global database Aurora PostgreSQL 12.22 eu-west-1 2 clusters						
apg-prod-cluster-dublin Available Primary cluster Aurora PostgreSQL 12.22 eu-west-1 3 instances						
apg-prod-cluster-dublin-instance-1 Available Writer instance Aurora PostgreSQL 12.22 eu-west-1a db.r6g.2xlarge						
apg-prod-cluster-dublin-instance-2 Available Reader instance Aurora PostgreSQL 12.22 eu-west-1b db.r6g.2xlarge						
apg-prod-cluster-dublin-instance-3 Available Reader instance Aurora PostgreSQL 12.22 eu-west-1c db.r5.xlarge						
apg-prod-cluster-london Available Secondary cluster Aurora PostgreSQL 12.22 eu-west-2 3 instances						
apg-prod-cluster-london-instance-1 Available Reader instance Aurora PostgreSQL 12.22 eu-west-2a db.r6g.2xlarge						
apg-prod-cluster-london-instance-2 Available Reader instance Aurora PostgreSQL 12.22 eu-west-2b db.r6g.2xlarge						
apg-prod-cluster-london-instance-3 Available Reader instance Aurora PostgreSQL 12.22 eu-west-2c db.r6g.2xlarge						

With this announcement, AWS is introducing *Blue/Green Deployment* support for **Amazon Aurora Global Database**, enabling database upgrades and updates with minimal downtime.



[AWS Database Blog Post](#)



Other Selected Announcements (1/2)



1. Amazon OpenSearch Service now supports automatic semantic enrichment (re:Invent)
2. Amazon RDS for Oracle and SQL Server now support up to 256 TiB storage volumes (re:Invent)
3. Amazon RDS for SQL Server launches new instances for up to 55% lower price (re:Invent)
4. Amazon RDS for SQL Server now supports Developer Edition (re:Invent)
5. **Amazon Aurora DSQL:**
 - a) Clusters support up to 256 TiB of storage volume (pre-re:Invent)
 - b) Provides statement-level cost estimates in query plans (pre-re:Invent)
6. Amazon RDS for PostgreSQL now supports major version 18 (pre-re:Invent)
7. Amazon Keyspaces (for Apache Cassandra) now supports Logged Batches (pre-re:Invent)
8. Amazon RDS extends IPv6 support for publicly accessible databases (pre-re:Invent)
9. Amazon ElastiCache now supports dual-stack (IPv4 and IPv6) service endpoints (pre-re:Invent)

Other Selected Announcements (2/2)



10. Amazon OpenSearch Service:

- a) Enhances Log Analytics with New PPL Experience (pre-re:Invent)
- b) Supports OpenSearch version 3.3 (pre-re:Invent)
- c) Launches Cluster Insights for improved operational visibility (pre-re:Invent)
- d) OpenSearch Serverless now adds audit logs for data plane APIs (pre-re:Invent)

11. Amazon CloudWatch Database Insights:

- a) Adds cross-account cross-region monitoring (pre-re:Invent)
- b) Expands anomaly detection in on-demand analysis (pre-re:Invent)

12. Amazon DocumentDB (with MongoDB compatibility):

- a) Introduces version 8.0 (pre-re:Invent)
- b) Announces upgraded query planner that can run queries up to 10x faster (pre-re:Invent)

Edge, Networking and Content Delivery

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



AWS AI Factories

RAPIDLY DEPLOYABLE, HIGH-PERFORMANCE AI INFRASTRUCTURE IN YOUR DATA CENTERS



GA

Dec, 2nd

re:Invent

NEW

AWS AI Factories

Dedicated customer-specific AI infrastructure,
built fast and managed by AWS

Integrate with existing hardware, data center,
and networking investments

Make use of the latest hardware
and AWS AI services



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Announcement](#)

AWS Interconnect Multicloud

SIMPLE, RESILIENT, HIGH-SPEED PRIVATE CONNECTIONS TO OTHER CLOUD SERVICE PROVIDERS




PP

Nov, 30th

re:Invent

5 Regions



AWS Interconnect is a new managed private connectivity service. It enables you to create high-speed network connections between your Amazon Virtual Private Clouds (VPCs) and your VPCs on other public clouds.



AWS Interconnect Last Mile

GATED PREVIEW OF A FULLY MANAGED CONNECTIVITY OFFERING FOR HYBRID CLOUD

Additionally, **AWS Interconnect** provides a *last mile*, a fully managed connectivity offering that allows customers to connect their branch offices, data centers, and remote locations to AWS with just a few clicks.

This *gated preview* is a collaboration between AWS and Lumen, combining AWS cloud innovation with Lumen's extensive network footprint to redefine how businesses connect to the cloud.

Customers can now instantly establish private, high-speed connections to AWS by simply entering their location, selecting bandwidth, and choosing their AWS Region. The launch simplifies the connectivity experience by automating complex network configuration including *BGP peering, VLAN configuration, and ASN assignment*.



Gated Preview

AWS Interconnect - last mile

AWS Interconnect - last mile (Gated Preview) is a fully managed connectivity offering that allows customers to connect their remote locations to AWS with just a few clicks. It makes it easier for customers to onboard to AWS Direct Connect from a broader set of locations and removes the operational burden by automating network configuration. Customers experience:

- 1. Turnkey Solution:** Enterprise customers can now setup a private AWS Direct Connect connection in minutes using their current last-mile network provider. For locations their provider cannot serve, this auto-discovers eligible partners based on the address and provides a guided experience to quickly sign-up and request a connection.
- 2. Simplified Network Setup:** Interconnect - Last mile abstracts away the complexity of network plumbing such as ordering cross-connects, configuring VLANs, ASNs, and BGP peering by automating these steps and enforcing AWS best practices. This significantly lowers the risk of misconfiguration errors.
- 3. Flexible Scaling:** Customers increasingly need the ability to adjust bandwidth speeds as workloads fluctuate. AWS Interconnect - last mile supports scaling natively, from 1 Gbps to 100 Gbps with increments at 1, 2, 5, 10, 25, 50, and 100 Gbps. Customers can scale up or down via the AWS console without needing to reprovision connections or open a support case.

AWS Interconnect - last mile is now available in gated preview. You will need an AWS Account Number in order to request access. Please submit the information below to request an invitation to the preview. We will contact you with instructions if you are accepted.



P

Nov, 30th

re:Invent



AWS Announcement



Route 53 Global Resolver for anycast DNS

PUBLIC PREVIEW OF THE NEW INTERNET-REACHABLE DNS RESOLVER WITH ANYCAST DNS



PP

Nov, 30th

re:Invent

11 Regions


Networking and Security

Amazon Route 53 Global Resolver

Secure and reliable global DNS resolution for public and private domains

Amazon Route 53 Global Resolver is a fully managed DNS resolver service that provides secure DNS resolution with advanced filtering capabilities. Simplify DNS connectivity, apply security policies across all clients, and ensure high availability with anycast IPs and multi-Region deployment.

How it works



Amazon Route 53 Global Resolver is internet-reachable DNS resolver providing split DNS resolution for public and private domains.

Supports *DNS over UDP*, *DNS-over-HTTPS*, and *DNS-over-TLS* protocols. It has integrated DNS Firewall with threat filtering and centralized logging capabilities.

It provides global anycast architecture with automatic failover to closest available region.



[AWS News Blog Post](#)



Accelerated recovery in Amazon Route 53

60-MINUTE RECOVERY TIME OBJECTIVE (RTO) FOR MANAGING PUBLIC DNS RECORDS




GA

Nov, 26th

pre-re:Invent

All Regions




[AWS News Blog Post](#)



AWS Network Firewall Proxy

EFFICIENTLY MANAGE AND SECURE WEB AND INTER-NETWORK TRAFFIC



PP

Nov, 25th

pre-re:Invent

US East (Ohio)



[AWS N&CD Blog Post](#)

mTLS support in Amazon CloudFront

VALIDATE CLIENT IDENTITIES AT AMAZON CLOUDFRONT EDGE LOCATIONS USING X.509 CERTS

With this announcement, **Amazon CloudFront** supports viewer *mTLS* authentication from end users to CloudFront, enhancing security for highly distributed and sensitive applications.

mTLS adds an extra layer of security by requiring both parties to present and verify digital certificates, creating two-way authentication.

Viewer mutual authentication (mTLS)
Require and validate viewer (client) certificates against your certificate authorities.

☐ Off

☒ On

Mode
Specify if client certificates are required to access your distribution.

☒ Required
Clients must provide a valid client certificate.

☐ Optional
Allow access to clients without certificates, and deny access to clients with invalid certificates. Optionally, you can use connection functions to validate clients without certificates, based on other criteria.

Trust store
The trust store validates client certificates in order to grant access to the distribution.

Choose a trust store to validate client certificates

☒ Ignore certificate expiration date
Allow expired (but otherwise valid) client certificates to access the distribution.

☐ Advertise trust store CA names
Present the CA names to the client during the TLS handshake.

Connection function - optional | [Info](#)
A function that triggers after the TLS handshake completes to selectively gate access based on client certificate validity and expiration (mTLS).

None

Description - optional

GA

Nov, 24th

pre-re:Invent

All Regions

[AWS N&CD Blog Post](#)

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

81

BYOIP for Anycast Static IPs in CloudFront

BRING YOUR OWN IP (BYOIP) VIA VPC IP ADDRESS MANAGER (IPAM) FOR ANYCAST STATIC IPs



GA

Nov, 24th

pre-re:Invent

All Regions

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

Bring your own IP to CloudFront using IPAM



Focus mode


This tutorial shows how to use IPAM to manage your BYOIP CIDRs for CloudFront Anycast Static IP lists.

Topics

- [What is BYOIP for Anycast Static IPs?](#)
- [Why use this feature?](#)
- [Prerequisites](#)
- [Step 1: Request an Anycast static IP list](#)
- [Step 2: Create an Anycast static IP list](#)
- [Step 3: Create a CloudFront distribution](#)
- [Step 4: Associate with CloudFront resources](#)
- [Step 5: Prepare for migration](#)
- [Step 6: Advertise CIDR globally](#)

What is BYOIP for Anycast Static IPs?

CloudFront supports bringing your own IPv4 addresses through IPAM's BYOIP for global services. Through IPAM's unified interface, customers can create dedicated IP address pools using their own IP addresses (BYOIP) and assign them to CloudFront distributions while leveraging the AWS worldwide content delivery network to deliver their applications and content. Your IP addresses are advertised from multiple CloudFront edge locations simultaneously using anycast routing.



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

VPC Encryption Controls for Amazon VPC

AUDIT AND ENFORCE ENCRYPTION IN TRANSIT WITHIN AND ACROSS VIRTUAL PRIVATE CLOUDS



GA

Nov, 21st

pre-re:Invent

26 Regions

VPC > Encryption controls > Create VPC encryption control

Create encryption control

Encrypt all network traffic within and between VPCs, while enforcing encryption-capable resources and providing audit visibility through flow logs.

Encryption control settings

Name - optional

Creates a tag with a key of 'Name' and a value that you specify.

AWS News Blog Demo

Character count: 18/255

VPC

Select the VPC for which you want to create an encryption control.

vpc-979f65f0

Create a VPC

Encryption mode

Monitor mode provides visibility into encryption status without blocking traffic. Enforce mode prevents unencrypted traffic.

Monitor mode

See which resources in your VPC are unencrypted but allow the creation of unencrypted resources.

Enforce mode

Requires all resources, except exclusions, in your VPC to be encryption-capable and blocks creation of unencrypted resources.

You may switch to enforce mode once all VPC resources are either encrypted or excluded from encryption requirements. Note that these exclusions can only be marked at the time of switching to enforce mode.

Tags

Add tags to help organize and identify your encryption control.

Key

Q Name

X

Value - optional

Q AWS News Blog Demo

X

Remove

Add new tag

You can add 49 more tags.

Cancel

Create encryption control

[AWS News Blog Post](#)

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

83

Post-quantum key exchange in ALB and NLB

SUPPORTING HYBRID POST-QUANTUM KEY AGREEMENT FOR TRANSPORT LAYER SECURITY (TLS)



GA

Nov, 21st

pre-re:Invent

All Regions

AWS Application Load Balancer (ALB) and **Network Load Balancer (NLB)** are now supporting post-quantum key exchange options for TLS protocol.

Post-quantum TLS (PQ-TLS) security policies protect your data in transit against potential *Harvest Now, Decrypt Later (HNDL)* attacks, where adversaries collect encrypted data today with the intention to decrypt it once quantum computing capabilities mature.

What is post-quantum cryptography at AWS?

At AWS, the confidentiality, integrity, and authenticity of our customers' data is a top priority. Today's widely-used public-key cryptographic schemes rely on mathematical problems - like integer factoring and discrete logarithms - that could be efficiently solved by future quantum computers. To address this challenge, AWS is deploying new NIST-standardized post-quantum cryptographic algorithms that are designed to resist both classical and quantum computing attacks. These algorithms, including the Module-Lattice-based Key-Encapsulation Mechanism (ML-KEM) and Module-Lattice-based Digital Signature Algorithm (ML-DSA), are based on different mathematical foundations that are believed to be resistant to quantum computing attacks.

AWS has already deployed post-quantum cryptography across several key services. [AWS Key Management Service \(AWS KMS\)](#), [AWS Secrets Manager](#), and [AWS Certificate Manager](#) have implemented post-quantum hybrid key establishment combining Elliptic Curve Diffie-Hellman (ECDH) with ML-KEM to protect against "harvest now, decrypt later" attacks. At the foundation of these implementations is [AWS-LC](#), our [FIPS-140-3-validated cryptographic library](#), which was the first open-source cryptographic module to include ML-KEM in its FIPS validation.



AWS Announcement



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Target Optimizer for AWS ALB

MAXIMUM CONCURRENT REQUESTS PER TARGET IN AWS APPLICATION LOAD BALANCER (ALB)




GA

Nov, 20th

pre-re:Invent

All Regions




AWS Application Load Balancer (ALB) gains a new powerful feature that allows for delivering optimal concurrency per each target attached to the load balancer. **Target Optimizer** works with the help of an AWS-provided agent by involving targets in load distribution.



[AWS N&CD Blog Post](#)

AWS Cloud WAN Routing Policy

FINE-GRAINED CONTROLS FOR ROUTES AND PATTERNS ACROSS GLOBAL WIDE-AREA NETWORKS



GA

Nov, 20th

pre-re:Invent

All Regions




[AWS N&CD Blog Post](#)



Weighted Target Groups in AWS NLB

SIMPLIFYING DEPLOYMENT FOR AWS NETWORK LOAD BALANCER (NLB)

AWS Network Load Balancer (NLB) now supports weighted target groups, allowing you to distribute traffic across multiple target groups with configurable weights for advanced deployment strategies.





GA

Nov, 19th

pre-re:Invent

All Regions




[AWS N&CD Blog Post](#)



Regional availability for AWS NAT Gateway

AUTOMATED EXPANSION AND CONTRACTION ACROSS AVAILABILITY ZONES (AZ)

This announcement introduces **AWS Regional NAT Gateway (RNAT)** that allows for automatic enhancement and contraction across availability zones in the designated VPC with improved security posture, built-in protection for port exhaustion, integration with VPC IPAM, and many more.



GA

Nov, 19th

pre-re:Invent

All Regions



AWS Announcement

GA

```
pre-re:Invent
```

All Regions



AWS N&CD Blog Post

Flat-rate pricing plans in Amazon CloudFront

NO OVERAGES FOR WEBSITE DELIVERY AND SECURITY WITH FLAT-RATE PRICING PLANS




GA

Nov, 18th

pre-re:Invent

All Regions



Amazon CloudFront introduces new pricing plans, that combine global content delivery (CDN) with multiple AWS services and features into a monthly price with no overage charges, regardless of whether your website or application goes viral or faces a DDoS attack.

Flat-rate pricing plans are now available in *Free* (\$0/month), *Pro* (\$15/month), *Business* (\$200/month), and *Premium* (\$1,000/month) tiers to match your application's needs.



[AWS News Blog Post](#)



QUIC passthrough in AWS NLB

QUIC PROTOCOL SUPPORT IN PASSTHROUGH MODE FOR AWS NETWORK LOAD BALANCER (NLB)



GA

Nov, 13th

pre-re:Invent

All Regions

Security groups [Info](#)

A security group is a set of firewall rules that control the traffic to your load balancer. Select an existing security group, or you can [create a new security group](#).

Security groups - recommended

Security groups support on Network Load Balancers can only be enabled at creation by including at least one security group. You can change security groups after creation. The security group for your load balancer must allow it to communicate with registered targets on both the listener port and the health check port. For PrivateLink Network Load Balancers, security group rules are enforced on PrivateLink traffic; however, you can turn off inbound rule evaluation after creation within the load balancer's Security tab or using the API.

Select up to 5 security groups

⚠ Network Load Balancers that are created without security groups can't support them after creation.

Listeners and routing [Info](#)

A listener is a process that checks for connection requests using the port and protocol you configure. The rules that you define for a listener determine how the load balancer routes requests to its registered targets.

▼ Listener TCP:80

Protocol TCP **Port** 80

Default action [Info](#)

Forward to Select a target group

[Create target group](#)

TCP
Suitable for web browsing, email, text messaging, and file transfers.

TLS
Offloads encryption and decryption to your load balancer. Requires an SSL server certificate on the listener.

UDP
High transfer speeds for video/audio streaming, games, VoIP.

TCP_UDP
Supports both TCP and UDP on the same port.

QUIC
Fast transport optimized for mobile applications.

TCP_QUIC
Supports both TCP and QUIC on the same port.



AWS N&CD Blog Post

JWT Verification in AWS ALB

SECURE MACHINE-TO-MACHINE (M2M) AND SERVICE-TO-SERVICE (S2S) COMMUNICATION

[Documentation](#) > [Elastic Load Balancing](#) > Application Load Balancers

Verify JWTs using an Application Load Balancer

↓ PDF

↓ RSS

🔒 Focus mode

You can configure an Application Load Balancer (ALB) to verify JSON Web Tokens (JWT) provided by clients for secure service-to-service (S2S) or machine-to-machine (M2M) communications. The load balancer can verify a JWT no matter how it was issued and without human interaction.

ALB will validate the token signature and requires two mandatory claims: 'iss' (issuer) and 'exp' (expiration). Additionally, if present in the token, ALB will also validate 'nbf' (not before) and 'iat' (issued at time) claims. You can configure up to 10 additional claims for validation. These claims support three formats:

- Single-string: A single text value
- Space-separated values: Multiple values separated by spaces (maximum 10 values)
- String-array: An array of text values (maximum 10 values)

If the token is valid, the load balancer forwards the request with token as is to the target. Otherwise, it rejects the request.



GA

Nov, 12th

pre-re:Invent

All Regions



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Announcement](#)

Other Selected Announcements (1/2)



1. [Amazon Kinesis Video Streams now supports a new warm storage tier](#) (pre-re:Invent)
2. [Amazon Route 53 DNS service adds support for IPv6 API service endpoint](#) (pre-re:Invent)
3. [AWS WAF announces Web Bot Auth support](#) (pre-re:Invent)
4. [AWS announces Flexible Cost Allocation on AWS Transit Gateway](#) (pre-re:Invent)
5. [AWS Application Load Balancer now supports Health Check Logs](#) (pre-re:Invent)
6. [AWS Network Firewall now supports flexible cost allocation via Transit Gateway](#) (pre-re:Invent)
7. [Amazon CloudFront now supports TLS 1.3 for origin connections](#) (pre-re:Invent)
8. [Amazon CloudFront announces 3 new CloudFront Functions capabilities](#) (pre-re:Invent)
9. [Amazon CloudFront now supports CBOR Web Tokens and Common Access Tokens](#) (pre-re:Invent)
10. [Amazon Route 53 DNS service now supports AWS PrivateLink](#) (pre-re:Invent)
11. [Active threat defense now enabled by default in AWS Network Firewall](#) (pre-re:Invent)
12. [Amazon Route 53 DNS Firewall adds protection against Dictionary-based DGA attacks](#) (pre-re:Invent)
13. [Amazon Route 53 Profiles now supports Resolver query logging configurations](#) (pre-re:Invent)
14. [AWS IoT Services expand support of VPC endpoints and IPv6 connectivity](#) (pre-re:Invent)

Other Selected Announcements (2/2)



- 15. [Amazon Kinesis Video Streams WebRTC Multi-Viewer](#) (pre-re:Invent)
- 16. [AWS Site-to-Site VPN announces 5 Gbps bandwidth tunnels](#) (pre-re:Invent)
- 17. [Amazon VPC Lattice now supports custom domain names for resource configurations](#) (pre-re:Invent)
- 18. [Amazon CloudFront announces cross-account support for VPC origins](#) (pre-re:Invent)
- 19. [Amazon Cloudfront adds IPv6 support for Anycast Static IPs](#) (pre-re:Invent)
- 20. [Amazon Route 53 Resolver now supports AWS PrivateLink](#) (pre-re:Invent)
- 21. [Amazon VPC IPAM automates prefix list updates](#) (pre-re:Invent)
- 22. [Announcing an AI agent context pack for AWS IoT Greengrass developers](#) (pre-re:Invent)

Analytics

13 announcement (11 *pre-re:Invent*)



AI-powered Data Classification in SageMaker

AUTOMATED AND AI-GENERATED SUGGESTIONS FOR BUSINESS TERMS IN THE GLOSSARIES

Amazon SageMaker Catalog now provides automated data classification that suggests business glossary terms during data publishing, reducing manual tagging effort and improving metadata consistency across organizations.

It analyzes table metadata and schema information using Amazon Bedrock's language models to recommend relevant terms from organizational business glossaries.

[Documentation](#) > [Amazon SageMaker Unified Studio](#) > [User Guide](#)

Using machine learning and generative AI in Amazon SageMaker Unified Studio

[PDF](#)

[RSS](#)

☐ Focus mode

Note

Powered by Amazon Bedrock: AWS implements automated abuse detection. Because the AI recommendations for assets in Amazon SageMaker Unified Studio is built on Amazon Bedrock, users inherit the controls implemented in Amazon Bedrock to enforce safety, security, and the responsible use of AI.

In the current release of Amazon SageMaker Unified Studio, you can use the AI recommendations for names, descriptions, and glossary terms functionality to automate data discovery and cataloging.

Powered by Amazon Bedrock's large language models, the AI recommendations for data asset names, descriptions, and glossary terms in Amazon SageMaker Unified Studio help you to ensure that your data is comprehensible and easily discoverable. The AI recommendations also suggest the most pertinent analytical applications for datasets. By reducing manual documentation tasks and advising on appropriate data usage, auto-generated names and descriptions can help you to enhance the trustworthiness of your data and minimize overlooking valuable data to accelerate informed decision making.

AI recommendations for glossary terms is a feature that automatically analyzes asset metadata and context to determine the most relevant business glossary terms for each asset and its columns. Instead of relying on manual tagging or static rules, it reasons about the data and performs iterative searches across what already exists in the customer's environment to identify the best-fit glossary term concepts. Because the system suggests terms only from glossaries and definitions already present in the system, customers are encouraged to maintain high-quality, well-described glossary entries so the AI can return accurate and meaningful suggestions. This improves metadata quality, strengthens governance, accelerates data onboarding, and reduces manual stewardship effort at scale.



GA

Nov, 30th

re:Invent

12 Regions



AWS Announcement



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Synthetic dataset generation in Clean Rooms

PRIVACY-ENHANCING SYNTHETIC DATASETS FROM COLLECTIVE DATA FOR TRAINING ML MODELS



GA

Nov, 30th

re:Invent

All Regions

Synthetic data - new

☒ Require analysis template output to be synthetic
Generates synthetic data for model training. Analysis templates with synthetic outputs can only be used to create ML input channels.

Column classification [Info](#)

Classify all columns that appear in your dataset output schema as numerical or categorical. ML input channel creation fails if columns aren't classified.

Column	Classification	New column name
label	Categorical	
pixel0	Numerical	
pixel1	Numerical	
pixel2	Numerical	
Add new column name	Categorical	composite_column

Character count: 16/128.

Add another column

At least 5 columns are required.

Predictive value

Choose the column for model training data to predict.

composite_column

Advanced settings [Info](#)

Adjust the default settings to align with your organizations' privacy requirements.

Privacy level in epsilon [Info](#)

Indicates the largest epsilon the synthetic data generator can use when generating your synthetic dataset. Generally, a lower value provides more privacy, while a higher value provides more model accuracy.

4.855

Must be between 0.0001 and 10.

Privacy threshold in % [Info](#)

The highest allowed probability that a membership inference attack could identify members of the original dataset. Scores of 50% indicate that a membership inference attack cannot successfully distinguish members from non-members better than a random guess. For no privacy limit, enter 100.

51.5

Must be between 50.0 and 100.

AWS Clean Rooms are now providing generation of privacy-enhancing synthetic datasets from collective data for ML training.

Throughout the process it de-identifies subjects while maintaining similar statistical properties to original data.

Unlocks ML use cases previously restricted by privacy or competition concerns, as it enables collaborative training without exposing raw data between organizations.



[AWS News Blog Post](#)



Deletion and row lineage in Iceberg for EMR

AWS ANALYTICS SUPPORT DELETION VECTORS AND ROW LINEAGE FROM APACHE ICEBERG V3



GA

Nov, 26th

pre-re:Invent

All Regions

With that announcement **AWS Analytics** services (Apache Spark on **Amazon EMR** 7.12, **AWS Glue**, **Amazon SageMaker** notebooks, **Amazon S3 Tables**, and **AWS Glue Data Catalog**) are supporting *deletion vectors* and *row lineage*, as it is defined in the **Apache Iceberg** Version 3 (V3) specification.

```
deletion-vectors-and-row-lineage.spark.sql ×
deletion-vectors-and-row-lineage.spark.sql
1 CREATE TABLE customer_data (
2     customer_id bigint,
3     name string,
4     email string,
5     last_purchase timestamp,
6     total_spent decimal(10,2)
7 )
8 USING iceberg
9 TBLPROPERTIES (
10     'format-version' = '3',
11     'write.delete.mode' = 'merge-on-read',
12     'write.update.mode' = 'merge-on-read',
13     'write.merge.mode' = 'merge-on-read'
14 )
```



[AWS Big Data Blog Post](#)



Glue Zero ETL for self-managed databases

SETUP ZERO ETL INTEGRATION FROM ORACLE, SQL SERVER, MYSQL, OR POSTGRESQL

Zero-ETL capabilities in AWS Glue

Zero-ETL integrations in AWS Glue simplify data ingestion and replication from AWS data service, third-party applications, and self-managed databases to AWS destinations.

AWS services supported as zero-ETL sources in AWS Glue include:

- Amazon DynamoDB

Third-party applications supported as zero-ETL sources in AWS Glue include:

- Facebook Ads
- Instagram Ads
- Salesforce
- Salesforce Marketing Cloud Account Engagement
- SAP OData
- ServiceNow
- Zendesk
- Zoho CRM

Self-managed databases supported as zero-ETL sources in AWS Glue include:

- Oracle
- SQL Server
- MySQL
- PostgreSQL

For more information about zero-ETL integrations from self-managed sources, see [AWS zero-ETL integration for self-managed database sources](#).



GA

Nov, 25th

pre-re:Invent

7 Regions



[AWS Announcement](#)



Multi-warehouse federation for Redshift

PERMISSION FEDERATION ACROSS MULTIPLE WAREHOUSES IN AMAZON REDSHIFT

[Documentation](#) > [Amazon Redshift](#) > [Database Developer Guide](#)

Amazon Redshift federated permissions

PDF

RSS

Focus mode

Key concepts

- **Redshift warehouse with federated permissions:** The producer warehouse that is registered with Data Catalog and where data and Redshift permissions are stored.
- **Consuming Redshift warehouse:** Any warehouse that queries data from remote warehouse. Consuming warehouse can optionally be enabled for Redshift federated permissions.
- **Global identity:** IAM and IAM Identity Center provides global identity across all warehouses enabled for Redshift federated permissions. Users authenticate once through their existing identity provider and receive consistent access based on their global identity, regardless of which warehouse they connect to.
- **Auto-mounting:** All warehouses enabled for Redshift federated permissions are automatically visible in all warehouses within your account. This auto-mounting capability enables catalog and database discovery for cross-warehouse analytics.
- **Identity propagation:** When you execute a cross-warehouse query, Redshift propagates your global identity (IAM role or IAM Identity Center user) to the remote warehouse.
- **Authorization across warehouses:** The remote warehouse enabled Redshift federated permissions validates your permissions for cross-warehouse queries and are enforced in the consuming warehouses.
- **Fine grained access control:** Policies for Row Level Security (RLS), Column Level Policies (CLP), and Dynamic Data Masking (DDM) that can be enforced across warehouses.



GA

Nov, 24th

pre-re:Invent

4 Regions



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Announcement](#)

Iceberg catalog federation for AWS Glue

DIRECT AND SECURE ACCESS TO ICEBERG TABLES CATALOGED REMOTELY INSIDE AWS ANALYTICS



GA

Nov, 24th

pre-re:Invent

All Regions

☰

AWS Lake Formation > Catalogs

⛔

Catalogs Updated

Databases

Tables

Views

Data filters

Data sharing

Crawlers ↗

▼ Permissions

Data permissions

LF-Tags and permissions

Hybrid access mode

Data locations

▼ Administration

Administrative roles and tasks

Data lake locations

Application integration settings

IAM Identity Center integration

Data Catalog settings ⚠

▼ Ingestion

Blueprints

Catalogs

▼ How it works



Create a catalog

Register Redshift databases as catalogs in the Data Catalog. [Learn more ↗](#)



Manage catalog permissions

Manage permissions for specific catalogs, databases, tables and fine-grained data access. [Learn more ↗](#)



Access from query editors

Access catalog objects from [Redshift Query Editor v2 ↗](#) and [Athena Console ↗](#). 

ⓘ Automatically create a catalog from an S3 table bucket. Allows external engines to access data in Amazon S3 locations with full table access.

Enable S3 Table integration ✕

Catalogs (2) Info

🔄 Actions ▾ View ▾ Create catalog

A catalog is the top level in the Data Catalog's three-level data hierarchy and contains Data Catalog objects.

🔍 Find catalogs by name

< 1 > ⚙

Name ▲ Type ▾ Source ▾ Owner a... ▾ Connection... Shared r... ▾ Shared r... ▾ Shared ... ▾

CloudShell Feedback

© 2025, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences


[AWS Big Data Blog Post](#)

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

101

One-click dataset onboarding in SageMaker

IMMEDIATE PRODUCTIVITY WITH EXISTING DATASETS IN AMAZON SAGEMAKER UNIFIED STUDIO




GA

Nov, 21st

pre-re:Invent

15 Regions




[AWS News Blog Post](#)



Introducing Amazon SageMaker Data Agent

BUILT-IN AGENT THAT ACCELERATES DEVELOPMENT OF DATA ANALYTICS AND ML APPLICATIONS

Amazon SageMaker Data Agent is available inside a new serverless notebooks in **Amazon SageMaker Unified Studio**, and supports *SQL*, *Python*, *Spark*, or natural language and gives data engineers, analysts, and data scientists one place to develop and run both SQL queries and code – that includes authoring code for data processing and machine learning (ML) pipelines.



```
1from pyspark.sql.connect import functions as F
2import matplotlib.pyplot as plt
3import seaborn as sns
4import pandas as pd
5
6# Read the churn dataset
7df = spark.read.table("sagemaker_sample_db.churn")
8
9# Convert to pandas for visualization
10df_pd = df.toPandas()
11
12# Convert churn boolean to string for better visualization
13df_pd['churn_label'] = df_pd['churn'].map({True: 'Churned', False: 'Retained'})
14
15# Create a comprehensive dashboard with multiple visualizations
16fig, axes = plt.subplots(3, 3, figsize=(18, 15))
17fig.suptitle('Customer Churn Analysis Dashboard', fontsize=20, fontweight='bold', y=0.995)
18
19# 1. Churn Rate Overview
20churn_counts = df_pd[churn_label].value_counts()
```



GA

Nov, 21st

pre-re:Invent

9 Regions



[AWS News Blog Post](#)



Metadata rules for SageMaker Catalog

ENFORCING RULES FOR GLOSSARY TERMS FOR ASSETS PUBLISHING

Amazon SageMaker Catalog now supports metadata enforcement rules for glossary terms, requiring data producers to apply approved business vocabulary when publishing assets. This helps consistent data classification and improves discoverability across organizational catalogs.

Discover ▾Build ▾Govern ▾My_Project_mhomlto2 ▾

Home > Domain units > domain-11-07-2025-142909 > Create required metadata rule

Rule configuration

Name

Enter name

Action

Workflow action to require the selected metadata rules.

☒ Data asset and product publishing

Completed form is required before publishing data assets and products to the catalog.

☐ Subscription request

Completed form is required for subscription requests for data assets and/or products.

Rule requirements

Select the type of requirements for this rule

☐ Metadata forms

requires users to complete metadata forms that you attach

☒ Glossary association

Requires users to associate glossary terms to the asset


Required glossary terms

Select up to 5 glossary terms per rule.

term3 ✕

glosterm6 ✕

add terms



GA

Nov, 19th

pre-re:Invent

All Regions



[AWS News Blog Post](#)



View Kafka topics in Amazon MSK console

INSPECT YOUR KAFKA TOPICS WITHOUT SETTING KAFKA ADMIN CLIENTS

aws

Search

[Option+S]

Amazon MSK

Clusters

United States (N. Virginia)

Amazon MSK

MSK Clusters

Cluster configurations

Managed VPC connections

Replicators

MSK Connect

Connectors

Custom plugins

Worker configurations

MSK Integrations

S3 delivery

EventBridge pipes

Real-time vector embedding

Resources

AWS Streaming Data Solutions

AWS Glue Schema Registry

Customer survey

Clusters (1)

Last updated November 29, 2025 at 10:36 UTC

Create real-time vector embedding - new

Actions

Create cluster


List of clusters in the account. User action status shows disk capacity alerts on Standard brokers. To connect a provisioned MSK cluster, use private or public bootstrap brokers with either of the following sets of requirements. To connect a serverless MSK cluster, the cluster must meet private bootstrap broker requirements.

Find clusters

All types

< 1 >

Cluster name	Status	User action status	Cluster type	Authentication	Apache Kafka version	Metadata mode	Brokers
demo-cluster-1	Active	-	Provisioned	IAM	3.8.x	ZooKeeper	kafka



GA

Nov, 19th

pre-re:Invent

All Regions



AWS Announcement



Amazon Redshift writes to Apache Iceberg

APACHE ICEBERG TABLES WRITE CAPABILITY FROM WITHIN AMAZON REDSHIFT

CREATE TABLE

```
CREATE TABLE [IF NOT EXISTS] <external_schema>.<table_name> (  
    column_name data_type [, ...]  
)  
USING ICEBERG  
[LOCATION 's3://your-bucket-name/prefix/']  
[PARTITIONED BY [[column_name | transform_function]], ...]  
[TABLE PROPERTIES ('compression_type'='<compression_value>')]
```

You can also use three-part notation for S3 table buckets:

```
CREATE TABLE "<table_bucket_name>@s3tablescatalog".<database_name>.<table_name> (  
    column_name data_type [, ...]  
)  
USING ICEBERG  
[PARTITIONED BY [[column_name | transform_function]], ...]  
[TABLE PROPERTIES ('compression_type'='<compression_value>')]
```

Note that `<external_schema>` must be an existing external schema name in which the external table will be created. For more information about how to create and manage external schemas, see [CREATE EXTERNAL SCHEMA](#) in the Amazon Redshift documentation.

The `LOCATION` clause defines the table location for this newly created Iceberg table. For Amazon S3 tables, `LOCATION` cannot be specified as the table location is determined by Amazon S3 tables catalog (`s3tablescatalog`).



GA

Nov, 17th

pre-re:Invent

All Regions




[AWS Announcement](#)

Serverless deployment option for MWAA

AMAZON MANAGED WORKFLOWS FOR APACHE AIRFLOW (MWAA) HAS SERVERLESS OPTION NOW

With this announcement, you can create **Amazon MWAA** workflows via declarative **YAML format (DAG Factory)** that are deployed and billed in a truly serverless manner (*actual task execution time*).



GA

Nov, 17th

pre-re:Invent

All Regions



Payload size increase for Amazon KDS

AMAZON KINESIS DATA STREAMS (KDS) NOW SUPPORTS RECORD SIZES UP TO 10 MiB

You can create new Amazon Kinesis Data Streams or enable for existing ones an increased payload size (via UpdateMaxRecordSize API) up to **10 MiB** with no additional costs involved.

This is a **10x increase** from the previously supported limits (1 MiB). Update affects PutRecord and PutRecords APIs.

In addition to that, *AWS Lambda* from now supports larger payloads (up to 6 MiB) coming from *KDS*.



GA

Oct, 28th

pre-re:Invent

All Regions



[AWS Announcement](#)

Other Selected Announcements (1/2)



1. [Announcing the Apache Spark upgrade agent for Amazon EMR](#) (re:Invent)
2. [Amazon EMR Serverless eliminates local storage provisioning for Apache Spark workloads](#) (re:Invent)
3. [Amazon SageMaker Catalog now exports asset metadata as queryable dataset](#) (re:Invent)
4. [AWS Glue now supports Apache Iceberg based materialized views](#) (re:Invent)
5. [AWS Glue introduces version 5.1](#) (pre-re:Invent)
6. [Amazon EMR and AWS Glue now support audit context support with Lake Formation](#) (pre-re:Invent)
7. [Amazon EMR and AWS Glue has write operations with fine-grained access controls](#) (pre-re:Invent)
8. [AWS Glue Data Quality now supports rule labeling for enhanced reporting](#) (pre-re:Invent)
9. [AWS Glue Data Quality now supports pre-processing queries](#) (pre-re:Invent)
10. [AWS Glue launches Amazon DynamoDB connector with Spark DataFrame support](#) (pre-re:Invent)
11. [Amazon Athena launches auto-scaling solution for Capacity Reservations](#) (pre-re:Invent)
12. [Amazon Athena for Apache Spark is now available in Amazon SageMaker notebooks](#) (pre-re:Invent)
13. [AWS Glue supports AWS CloudFormation and AWS CDK for zero-ETL integrations](#) (pre-re:Invent)
14. [Amazon Athena adds cost and performance controls for Capacity Reservations](#) (pre-re:Invent)

Other Selected Announcements (2/2)



- 15. Amazon SageMaker Unified Studio adds EMR on EKS support with SSO capabilities (pre-re:Invent)
- 16. Amazon Kinesis Data Streams now supports up to 50 enhanced fan-out consumers (pre-re:Invent)
- 17. Amazon Redshift supports Just-In-Time (JIT) ANALYZE for Apache Iceberg (pre-re:Invent)
- 18. Amazon Redshift supports SUPER type with case-insensitivity (pre-re:Invent)
- 19. Amazon SageMaker Catalog now supports read and write access to Amazon S3 (pre-re:Invent)
- 20. Amazon SageMaker Unified Studio adds support for catalog notifications (pre-re:Invent)
- 21. Amazon SageMaker launches custom tags for project resources (pre-re:Invent)
- 22. Amazon MSK Express brokers now support Intelligent Rebalancing (pre-re:Invent)
- 23. Amazon Kinesis Data Streams launches On-Demand Advantage mode (pre-re:Invent)
- 24. AWS Clean Rooms launches advanced configurations to optimize SQL (pre-re:Invent)

Machine Learning, AI, and Generative AI

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



Reinforcement fine-tuning in Bedrock

66% ACCURACY GAINS ON AVERAGE OVER BASE MODELS WITH REINFORCEMENT FINE-TUNING

Create Reinforcement fine-tuning job [Info](#)

▼ Customization details

Customization name

Enter a customization name. This identifies the model and continuous customization when applicable, with the job using a modified version.

RFT-Model-8656bb

Maximum of 63 alphanumeric characters. Cannot include underscores (_) and spaces. Must be unique within your account in AWS Region.

☐ KMS key selection | [Info](#)

▼ Model details

Source model

Choose from a list of models that you wish to customize with using your own data.

 Amazon Nova 2... v1 [Info](#) [Edit](#)

←

Amazon Bedrock received a new model customization capability that creates smarter, more cost-effective models that learn from feedback and deliver higher-quality outputs for specific business needs.

Reinforcement fine-tuning uses a feedback-driven approach where models improve iteratively based on reward signals, delivering 66% accuracy gains on average over base models.



GA

Dec, 3rd

re:Invent

All Regions



[AWS News Blog Post](#)



Updates to Strands Agents

TYPESCRIPT, EDGE DEVICES, BI-DIRECTIONAL STREAMING, AND ADDITIONAL PROVIDERS

Strands Agents open source project got significant updates:

1. Dedicated **SDK for TypeScript**, with idiomatic language support working in browsers, *Node.js*, and runtimes like *AWS Lambda* or *Bedrock AgentCore*.
2. Edge devices got better support with **bi-directional streaming** and additional **local model providers** like *llama.cpp* for small-scale devices.
3. An experimental support for **Steering**, to support modular prompting and driving agents without rigid workflows.
4. **Evaluations** for validating behaviors, measuring improvements, and general observability are available in preview.



PP

Dec, 3rd

re:Invent

Open Source



AWS Announcement



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

SageMaker serverless model customization

QUICKLY CUSTOMIZE POPULAR MODELS WITH MODERN TECHNIQUES IN A SERVERLESS WAY




GA

Dec, 3rd

re:Invent

4 Regions



Amazon SageMaker AI got a new serverless customization for popular AI models. This new customization capability provides an easy-to-use interface for the latest fine-tuning techniques like reinforcement learning, so you can accelerate the AI model customization process.

It supports *Supervised Fine-Tuning* and the latest model customization techniques including *Direct Preference Optimization*, *Reinforcement Learning from Verifiable Rewards (RLVR)*, and *Reinforcement Learning from AI Feedback (RLAIF)*.



[AWS News Blog Post](#)

Checkpointless training for HyperPod

NEW CAPABILITY MITIGATING CHECKPOINT-BASED JOB-LEVEL RESTART FOR FAULT RECOVERY



GA


Dec, 3rd

re:Invent

All Regions

Amazon SageMaker HyperPod now supports checkpointless training, a new foundational model training capability that mitigates the need for a checkpoint-based job-level restart for fault recovery.

How? Checkpointless training maintains forward training momentum despite failures, reducing recovery time from hours to minutes.



[AWS News Blog Post](#)



Elastic training for HyperPod

ACCELERATE TRAINING FMS BY AUTOMATICALLY SCALING TRAINING WORKLOADS

Amazon SageMaker HyperPod now supports elastic training, enabling organizations to accelerate foundation model training by automatically scaling training workloads based on resource availability and workload priorities.

This is a fundamental shift from training with a fixed set of resources, as it saves hours of engineering time spent reconfiguring training jobs based on compute availability.



GA

Dec, 3rd

re:Invent

All Regions



[AWS News Blog Post](#)



Amazon Nova 2 Lite

FAST, COST-EFFECTIVE REASONING MODEL FROM NOVA FAMILY FOR EVERYDAY WORKLOAD

	Nova 2 Lite	Claude Haiku 4.5	GPT 5 Mini	Gemini 2.5 Flash
CORE INTELLIGENCE				
MMLU-Pro Knowledge-centric Reasoning	80.9	80.0	83.7	83.2
GPQA Diamond Scientific Reasoning	79.6	73.0	82.3	82.8
AIME 2025 Mathematics	91.0	80.7	91.1	72.0
IF-Bench Instruction Following	70.8	54.3	75.4	50.3
MultiChallenge Multi-turn Instruction Following	76.6	57.1	67.8	62.6
LongCodeBench 1M Long Context	84.0	—	—	78.0
MULTIMODAL PERCEPTION				
MMMU-Pro Complex Visual Reasoning	61.8	58.0	74.1	69.0
OCRBench-v2 Document OCR	56.1	48.3	55.4	58.2
RealKIE Key Information Extraction	62.1	52.2	53.1	50.1
QVHighlights Query-Based Video Highlight	77.2	—	—	52.4
AGENTIC WORKFLOWS				
r²-bench - Telecom Tool Use	76.0	54.7	71.1	31.6
r²-bench - Retail Verified Tool Use	76.5	69.1	73.7	57.7
r²-bench - Airline Verified Tool Use	64.8	54.0	68.8	44.0
BFCL Overall FC V4 Tool Use	60.3	61.8	56.0	55.9
Scale MCP Atlas Tool Use (Blind)	24.6	—	22.2	—
ScreenSpot User Interface Grounding	83.3	—	24.8	27.2
SOFTWARE ENGINEERING				
SWE-Bench Verified Agentic Coding	64.5	73.3	71.0	60.4
Terminal-Bench 1 Overall Terminal Coding	32.5	41.8	30.8	16.8
LiveCodeBench Coding Code Generation	71.0	61.5	83.8	69.5

Bold values indicate the best result on the benchmarks. Nova 2 values are bolded when the model achieves either the highest score on the benchmark or where the best result overlaps with the 95% confidence interval of Nova 2's result.

Amazon Nova 2 Lite is an industry-leading price performance model for everyday AI workloads.

It provides step-by-step reasoning with configurable budget levels for cost-efficiency.

It has 1M token context window for expanded reasoning and richer in-context learning. Also, it has built-in web grounding and code interpreter.



GA

Dec, 2nd

re:Invent

27 Regions



[AWS News Blog Post](#)



Amazon Nova 2 Pro

MOST INTELLIGENT MODEL FROM NOVA FAMILY FOR HIGHLY COMPLEX, MULTISTEP TASKS

	Nova 2 Pro	Claude Sonnet 4.5	GPT 5.1	Gemini 2.5 Pro	Gemini 3 Pro Preview
CORE INTELLIGENCE					
MMLU-Pro Knowledge-centric Reasoning	81.6	87.5	87.0	86.2	89.8
GPQA Diamond Scientific Reasoning	81.4	83.4	88.1	86.4	91.9
AIME 2025 Mathematics	92.3	87.0	94.0	88.0	95.0
IF-Bench Instruction Following	80.2	57.3	72.9	48.7	70.4
MultiChallenge Multi-turn Instruction Following	77.7	63.1	72.2	62.3	73.6
LongCodeBench 1M Long Context	84.0	82.0	—	74.0	82.0
MULTIMODAL PERCEPTION					
MMMU-Pro Complex Visual Reasoning	63.5	69.0	76.0	68.0	81.0
OCRBench-v2 Document OCR	64.5	52.3	57.2	59.3	69.6
RealKIE Key Information Extraction	67.0	62.4	53.6	54.9	71.7
QVHighlights Query-Based Video Highlight	76.7	—	—	75.0	48.1
Common Voice ASR Accuracy (100-WER%)	94.5	—	—	92.8	92.6
AGENTIC WORKFLOWS					
r²-bench - Telecom Tool Use	92.7	78.1	81.9	54.1	87.1
r²-bench - Retail Verified Tool Use	77.7	77.2	77.5	71.3	77.7
r²-bench - Airline Verified Tool Use	65.2	66.8	72.4	60.0	70.8
BFCL Overall FC V4 Tool Use	61.6	68.7	60.1	52.3	—
ScreenSpot User Interface Grounding	88.1	—	28.5	44.7	66.8
SOFTWARE ENGINEERING					
SWE-Bench Verified Agentic Coding	70.0	82.0	76.3	67.2	76.2
Terminal-Bench 1 Overall Terminal Coding	41.3	51.0	53.8	25.3	56.3
LiveCodeBench Coding Code Generation	74.6	71.4	86.8	80.1	91.7

Bold values indicate the best result on the benchmarks. Nova 2 values are bolded when the model achieves either the highest score on the benchmark or where the best result overlaps with the 95% confidence interval of Nova 2's result.

Amazon Nova 2 Pro is the most intelligent reasoning model with extended thinking capabilities.

It has three thinking intensity levels for speed, intelligence, and cost control (also available with *Nova 2 Lite*) and it can serve as *teacher* model for knowledge distillation to smaller models. It has built-in web grounding and code interpreter.

Available in the preview for the customers of *Amazon Nova Forge*, upon earlier request via AWS Account team.



Dec, 2nd

re:Invent



AWS Announcement



Amazon Nova 2 Omni

ALL-IN-ONE MODEL FROM NOVA FAMILY FOR MULTIMODAL REASONING AND IMAGE GENERATION



Dec, 2nd

re:Invent

All Regions

Nova 2 Lite A fast cost-effective reasoning model	Input: Text, image, video Output: Text	Use cases: Everyday AI tasks such as chatbots, code generation	Context window: 1M	Developer controls for thinking effort: Yes
Nova 2 Pro (preview) Our most intelligent reasoning model	Input: Text, image, video, speech Output: Text	Use cases: Complex AI tasks such as agentic coding	Context window: 1M	Developer controls for thinking effort: Yes
Nova 2 Omni (preview) All-in-one model for multimodal reasoning and image generation	Input: Text, image, video, speech Output: Text, image	Use cases: Image generation, multimodal understanding	Context window: 1M	Developer controls for thinking effort: Yes
Nova 2 Sonic A speech-to-speech model for conversational AI	Input: Text, speech Output: Text, speech	Use cases: Customer support, voice-enabled personal assistant	Context window: 1M	
Nova Multimodal Embeddings A state-of-the-art multimodal embedding model	Input: Text, image, video, audio Output: Embedding	Use cases: Semantic search, agentic RAG	Context window: 8K	



Amazon Nova 2 Sonic

SPEECH-TO-SPEECH MODEL FROM NOVA FAMILY FOR NATURAL, REAL-TIME CONVERSATIONAL AI



Dec, 2nd

re:Invent

4 Regions

Capability	Nova 2 Sonic	Gemini 2.5 Flash Live API	GPT Realtime
ASR Accuracy (100-WER% on Common Voice) ↑ Speech Understanding	93.5	84.1	91.6
Big Bench Audio (reported by Artificial Analysis) ↑ Speech Reasoning	87.0	71.0	83.0
Berkeley function calling leaderboard (overall accuracy) ↑ Task Completion	74.5	69.4	80.4
IFBench Prompt ↑ Instruction Following	33.3	40.4	33.3
IFBench Instruction Following ↑ Instruction Following	37.5	41.5	36.5

Amazon Nova 2 Sonic is a new, best-in-class streaming speech understanding with robustness to background noise.

It has expanded language support including *Portuguese* and *Hindi* with polyglot voices and provides configurable *turn-taking* and cross-modal voice-text interaction.



[AWS News Blog Post](#)



Amazon Nova Forge

BUILD YOUR OWN FRONTIER MODELS USING AMAZON NOVA FAMILY VIA AMAZON NOVA FORGE



GA

Dec, 2nd

re:Invent

US East
(N. Virginia)

Build with Nova 2.0 Lite

Custom model name *

my-cpt-run-e1

A training job with the name my-cpt-run-e1- will be created

Customization technique *

Continued Pre-Training (Base Model)

Train a base model further on domain-specific data to specialize its knowledge

Upload dataset Existing datasets

Dataset files must be JSONL and match Add up to

Mix with Nova Data Nova Forge Blend your proprietary data with Amazon Nova curated training data to mitigate catastrophic forgetting of general capabilities.

Enable data mixing Reset to defaults

Continued Pre-Training checkpoints Pre-trained Mid-trained Post-trained

Your training dataset 25 %

Amazon curated data for training Auto-balance Amazon Nova curated training data comprises of pre-training, mid-training, and post-training data from Nova's training. Use the recommended ratios across Amazon Nova's taxonomy as a starting point and modify the ratios to best suit your requirements. You can use the Auto-balance button to automatically redistribute the Amazon Nova curated data ratios across all domains while keeping your training dataset percentage fixed.

Agents 15 %

Business and Finance 3 %

Scientific 7.5 %

Code 3.75 %

We are excited to announce **Amazon Nova Forge**, a new service to build your own frontier models using Amazon Nova.

You can develop models on *SageMaker AI* from early Nova checkpoints across pre-training, mid-training, or post-training phases.

You can blend proprietary or industry-specific data with *Amazon Nova*-curated data to train the model, that truly understands your domain.



[AWS News Blog Post](#)



General availability of Amazon Nova Act

BUILD AND MANAGE FLEETS OF HIGHLY RELIABLE AGENTS FOR AUTOMATING UI WORKFLOWS



GA

Dec, 2nd

re:Invent

US East
(N. Virginia)

Amazon Nova

New Session

Generate Image

Generate Video

Speak

Side by Side

LABS

Act

Preview

Public Gallery

Nova API

RECENT SESSIONS

donnie

Nova Act

Untitled agent

Enter a URL

Describe an agent action

+ Add an Action

▶

↻

→

⬇

1. Open a Sample Agent

Click the Drop Down Arrow

Open Booking Agent Sample

2. Run the Agent

Edit the Steps

Or Just Click Play

3. Watch the Agent Go!

Read Agent Thinking

and Watch it Take Action

Nova Act Playground

a no-code, web-based agent builder that shows the power of the Nova Act SDK

This website collects information on interactions including prompts and screenshots. We do this to enable great features and to improve our services. You can request to delete your Nova Act data by emailing us at nova-act@amazon.com.

name	Amazon AGI
author	AGI Lab
url	nova.amazon.com/act
ver	1.2
loc	Luci
ver	V1.0
ver	12/2
ver	Playground Tutorial

Agent Output Window

Ready to execute automated workflow.

Agent browser disconnected

Agent browser connecting...

Ready to execute automated workflow.

Nova Act may not always get it right. Information collected from your interactions with Nova Act includes in-browser screenshots. Email us at nova-act@amazon.com to request deletion of your Nova Act data.



[AWS News Blog Post](#)



Policy in Amazon Bedrock AgentCore

INTERCEPT TOOL CALLS IN REAL-TIME TO ENSURE AGENTS STAY WITHIN DEFINED BOUNDARIES



Dec, 2nd
re:Invent
9 Regions



[AWS News Blog Post](#)

Policy

Policy provides deterministic, real-time enforcement to ensure agents stay within defined boundaries without slowing them down.

How it works

Create policy engine

A policy engine is a collection of policies that evaluates and authorizes agent tool calls. When associated with a gateway, the policy engine intercepts all agent requests and determines whether to allow or deny each action based on the defined policies.

Create policy engine

Associate and enforce policy engine to a Gateway

Associate the policy engine to a Gateway (one engine per Gateway) and choose an enforcement mode: log only mode to evaluate and trace decisions without enforcement, or enforcement mode to allow/deny AI agent tool calls based on policy decisions.

Associate Gateway

Audit policy enforcement decisions

Monitor policy enforcement metrics and traces for every Policy decision in Observability.

Observe with Gateways

Policy engines (1)

Policy engines offer unified control over a collection of policies including governance and association management.

Find policy engine

Name	Description	Status	Last updated date
PolicyEngine_90ncj	-	Active	November 25, 2025, 14:49 (UTC+00:00)

Load examples>Create policy(s)

Cedar preview (3) Policies generatedEdit Cedar (Code)

▼ policy_c0e2qValid

```
1 forbid (  
2   principal,  
3   action,  
4   resource ==  
5     AgentCore::Gateway::"arn:aws:bedrock-agentcore:us-east-1:  
6 )  
7 when  
8 {  
9   !((principal.hasTag("username")) &&  
10    ((principal.getTag("username")) like "ADMIN-*"))  
11 };
```

▼ policy_ldmakValid

```
1 forbid (  
2   principal,  
3   action,  
4   resource ==  
5     AgentCore::Gateway::"arn:aws:bedrock-agentcore:us-east-1:  
6 )  
7 when  
8 {  
9   !((principal.hasTag("scope")) &&  
10    ((principal.getTag("scope")) like "*messaging:write*"))  
11 };
```

▼ policy_rdbx3Not translatable

```
1
```

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

123

PP

re:Invent



124

Amazon Bedrock AgentCore enhancements

EPISODIC MEMORY FOR AGENTCORE MEMORY AND BI-DIRECTIONAL STREAMING IN RUNTIME



GA

Dec, 2nd

re:Invent

9 Regions

Episodic

The **episodic** (EpisodicStrategy) memory strategy captures interactions as structured episodes consisting of scenarios, intents, thoughts, actions taken, outcomes, and artifacts. With this strategy, reflections are also made across episodes to extract broader insights, letting an agent learn and apply successful patterns from prior interactions to new interactions.

- **Example use case:** A customer support agent logs interactions as episodes. The system captures which phrases and actions lead to successful interactions

Configuration example:

```
from bedrock_agentcore_starter_toolkit.operations.memory.manager import MemoryManager
from bedrock_agentcore_starter_toolkit.operations.memory.models.strategies import SummaryStrategy

# Create memory manager
memory_manager = MemoryManager(region_name="us-west-2")

# Create memory resource with summary strategy
memory_manager.get_or_create_memory(
    name="MyMemory",
    strategies=[
        EpisodicStrategy(
            name="EpisodicStrategy",
            namespaces=["strategy/{memoryStrategyId}/actors/{actorId}/sessions/{sessionId}"],
            reflection={
                "namespaces": ["strategy/{memoryStrategyId}/actors/{actorId}/"]
            }
        )
    ]
)
```

We have added new features to **Amazon Bedrock AgentCore Memory, Runtime, and Identity**:

1. *AgentCore Memory* now includes **episodic memory**, enabling agents to learn and adapt from experience, building knowledge over time.
2. *AgentCore Runtime* supports *bi-directional streaming* for natural conversations where agents listen and respond while handling interruptions.
3. *AgentCore Identity* now supports custom claims for enhanced authentication rules for multi-tenant environments.



[AWS News Blog Post](#)



SageMaker AI serverless MLflow capability


SERVERLESS MLFLOW CAPABILITY FOR FASTER AI DEVELOPMENT



Dec, 2nd

re:Invent

16 Regions



Amazon SageMaker AI offers now serverless MLflow capability.

It dynamically scales to support AI model development tasks without infrastructure setup, eliminates infrastructure management with automatic scaling and zero capacity planning.

It provides cross-account access via AWS Resource Access Manager (RAM) for simplified collaboration, and native integration with SageMaker AI JumpStart, Model Registry, and Pipelines.



[AWS News Blog Post](#)



New models from Mistral family on Bedrock

MISTRAL LARGE 3 AND MINISTRAL 3 FAMILY NOW AVAILABLE, FIRST ON AMAZON BEDROCK




GA

Dec, 2nd

re:Invent

5 Regions



Mistral Large 3 and Ministral 3 family are available first on **Amazon Bedrock**, spanning frontier intelligence to efficient edge computing.

Mistral Large 3 features 41B active parameters with 675B total in Mixture-of-Experts architecture. Ministral 3 family offers 14B, 8B, and 3B models for flexible deployment options. They provide 256k context window with powerful agentic capabilities for enterprise workflows.



Multimodal retrieval for Knowledge Bases

RETRIEVE RELEVANT INFORMATION ACROSS TEXT, IMAGES, AUDIO, AND VIDEO CONTENT

AWS announces the general availability of multimodal retrieval in **Amazon Bedrock Knowledge Bases**.

Supporting multimodal retrieval in Knowledge Bases enables developers to build AI-powered search and question-answering applications that retrieve information across documents, images, audio, and video files.

Build a knowledge base for multimodal content

[PDF](#) [RSS](#) ☐ Focus mode

Amazon Bedrock Knowledge Bases supports multimodal content including images, audio, and video files. You can search using images as queries, retrieve visually similar content, and process multimedia files alongside traditional text documents. This capability enables you to extract insights from diverse data types—standalone images, audio recordings, and video files stored across your organization.

Amazon Bedrock Knowledge Bases enable you to index and retrieve information from text, visual, and audio content. Organizations can now search product catalogs using images, find specific moments in training videos, and retrieve relevant segments from customer support call recordings.

Features and capabilities

Multimodal knowledge bases provide the following key capabilities:

Image-based queries

Submit images as search queries to find visually similar content when using Nova Multimodal Embeddings. Supports product matching, visual similarity search, and image retrieval.

Audio content retrieval

Search audio files using text queries. Retrieve specific segments from recordings with timestamp references. Audio transcription enables text-based search across spoken content including meetings, calls, and podcasts.

Video segment extraction

Locate specific moments within video files using text queries. Retrieve video segments with precise timestamps.

Cross-modal search

Search across different data types including text documents, images, audio, and video. Retrieve relevant content regardless of original format.

Source references with timestamps

Retrieval results include references to original files with temporal metadata for audio and video. Enables precise navigation to relevant segments within multimedia content.

Flexible processing options

Choose between native multimodal embeddings for visual similarity or text conversion for speech-based content. Configure processing approach based on content characteristics and application requirements.



GA

Nov, 30th

re:Invent

US East
(N. Virginia)



Training Plans for SageMaker AI inference

FLEXIBLE TRAINING PLANS (FTPs) ARE NOW SUPPORTING INFERENCE ENDPOINTS AS WELL



GA

Nov, 26th

pre-re:Invent

3 Regions



aws

Search

[Option+S]

United States (N. Virginia)

Amazon SageMaker AI

Training plans

create

Step 1

Search training plan offerings

Step 2

Add plan details

Step 3

Review and purchase

Search training plan offerings

Search the optimal training plan offerings for your ML workload requirements. Training Plan provides immediate feedback on the plan that matches your specific needs such as workload duration, instance type, and instance count.

Training plan requirements

Target

Select a target service for the training plan.

☐ Training Job

☐ HyperPod Cluster

☒ Inference Endpoint

Instance group must have at least one subnet in the training plan's Availability Zone.

Endpoint must have at least one subnet in the training plan's Availability Zone.

Capacity purchased for one target cannot be used for the other.

Compute type

Select a compute type for the training plan

☒ Instance

☐ UltraServer

Individual ML-optimized compute instances.

Interconnected set of AI accelerators optimized for large-scale AI workloads.

Instance type

m5p4d.24xlarge

Instance count

1

Date and duration

Choose a start date from today up to 8 weeks in the future.

Start date

2025/11/29

End date - optional

YYYY/MM/DD

Duration

1

Day(s)

Find training plan



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Announcement

129

EAGLE decoding supported by SageMaker AI

EAGLE SPECULATIVE DECODING IS SUPPORTED TO SPEED UP GENERATIVE AI INFERENCE BY 2.5X




GA

Nov, 25th

pre-re:Invent

7 Regions



Amazon SageMaker AI introduces new enhancements to its inference optimization toolkit that bring EAGLE based adaptive speculative decoding to more model architectures. An example benchmark above covers qwen3-32B model.



[AWS ML Blog Post](#)



Bidirectional streaming for SageMaker AI

AMAZON SAGEMAKER AI INFERENCE NOW SUPPORTS BIDIRECTIONAL STREAMING

Invoke to Bidirectionally Stream an Inference Request and Response

If you want to invoke a model endpoint in your application code to support bidirectional streaming, you can use the [new experimental SDK for Python](#) that supports bidirectional streaming capability with HTTP/2 support. This SDK enables real-time, two-way communication between your client application and the SageMaker endpoint, allowing you to send inference requests incrementally while simultaneously receiving streaming responses as the model generates them. This is particularly useful for interactive applications where both the client and server need to exchange data continuously over a persistent connection.

Note

The new experimental SDK is different from the standard Boto3 SDK and supports persistent bidirectional connections for data exchange. While using the experimental Python SDK we strongly advise strict pinning to a version of the SDK for any non-experimental use cases.

To invoke your endpoint with bidirectional streaming, use the `invoke_endpoint_with_bidirectional_stream` method. This method establishes a persistent connection that allows you to stream multiple payload chunks to your model while receiving responses in real-time as the model processes data. The connection remains open until you explicitly close the input stream or the endpoint closes the connection, supporting up to 30 minutes of connection time.

Amazon SageMaker AI Inference now supports bidirectional streaming for real-time speech-to-text transcription, enabling continuous speech processing instead of batch input.



GA

Nov, 25th

pre-re:Invent

29 Regions




[AWS ML Blog Post](#)



Embedded Chat for Amazon Quick Suite

EMBED QUICK SUITE'S CONVERSATIONAL AI DIRECTLY AND SECURELY INTO YOUR APPLICATIONS




GA

Nov, 24th

pre-re:Invent

4 Regions



[AWS BI Blog Post](#)



Test generation for Automated Reasoning

TEST GENERATION FOR AUTOMATED REASONING CHECKS IN AMAZON BEDROCK GUARDRAILS




GA

Nov, 21st

pre-re:Invent

6 Regions



From now on, when you create an **Automated Reasoning** check in **Amazon Bedrock Guardrails**, you can rely on natural language automated generation of tests examples for Q&A.



AWS Announcement



Service tiers for Amazon Bedrock

RESERVED, PRIORITY, STANDARD, AND FLEX INFERENCE SERVICE TIERS FOR AMAZON BEDROCK



GA

Nov, 18th

Nov, 26th

pre-re:Invent

All Regions

Priority	Lower latency responses for user-facing apps
Standard	Responsive performance for important workloads
Flex	Cost-efficient for less urgent workloads
Reserved	Predictable performance with guaranteed capacity



[AWS News Blog Post](#)



Direct code deployment for AgentCore

AMAZON BEDROCK AGENTCORE NOW SUPPORTS DIRECT PYTHON CODE DEPLOYMENT

Starting from this update, **Amazon Bedrock AgentCore** allows for direct code upload without the need to create Docker container first.

It supports **Strands Agents** and Python-based agents with runtime versions from 3.10 up to 3.13.

```
(.env) 30-amazon-bedrock-agentcore [main] make deploy-agentcore
```



GA

Nov, 4th

pre-re:Invent

9 Regions



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Four new image editing tools in **Bedrock**

EDITING, UPSCALING, CONTROL TOOLS ARE AVAILABLE AS STABILITY AI IMAGE SERVICES



Amazon Bedrock announces the availability of 4 new image editing tools to **Stability AI Image Services**: outpaint, fast upscale, conservative upscale, and creative upscale.



GA

Oct, 29th

pre-re:Invent

3 Regions



[AWS News Blog Post](#)

Web Grounding for Amazon Nova models

USING REAL-TIME INFORMATION TO REDUCE HALLUCINATIONS AND ACCURATE OUTPUT




GA

Oct, 28th

pre-re:Invent

3 Regions



[AWS News Blog Post](#)



Amazon Nova Multimodal Embeddings

FIRST UNIFIED EMBEDDING MODEL SUPPORTING TEXT, DOCUMENTS, IMAGES, VIDEO, AND AUDIO



GA

Oct, 28th

pre-re:Invent

US East
(N. Virginia)

	Amazon Nova Multimodal Embeddings	TwelveLabs Marengo 2.7	Google Vertex AI Multimodal Embeddings	Cohere Embed 4	Amazon Titan Multimodal Embeddings	Amazon Titan Text Embeddings V2	Google Gemini Embedding
VIDEO RETRIEVAL							
ActivityNet Avg (recall@1, recall@5, recall@10)	81.2	51.0	55.7	—	—	—	—
DiDeMo Avg (recall@1, recall@5, recall@10)	80.6	64.6	54.3	—	—	—	—
AUDIO RETRIEVAL							
AudioCaps Avg (recall@1, recall@5, recall@10)	69.9	55.9 ¹	—	—	—	—	—
VISUAL DOCUMENT RETRIEVAL							
ViDoRe V2 NDCG@5	58.7	—	—	53.6	—	—	—
IMAGE RETRIEVAL							
TextCaps Avg (recall@1, recall@5, recall@10)	88.9	75.7 ¹	*	—	84.7	—	—
MSCOCO Avg (recall@1, recall@5, recall@10)	76.7	77.6 ¹	70.1	—	74.8	—	—
TEXT RETRIEVAL							
MTEB (Multilingual) NDCG@10	63.8	—	—	*	—	58.3	67.7 ¹

— Unsupported use-case
* Benchmark not available as of 10/2025

Evaluation Methodology
Benchmark numbers for non-Amazon models are a mix of self-reported numbers and Amazon conducted evaluations performed on Amazon Bedrock and Google Vertex. Embedding dimensions used: Amazon Nova Multimodal Embeddings: 3072; TwelveLabs Marengo 2.7: 1024; Google Vertex Multimodal Embeddings: 1408; Cohere Embed 4: 1536; Amazon Titan Multimodal Embeddings: 1024; Amazon Titan Text Embeddings V2: 1024; Google Gemini Embedding: 3072. Note: results may differ depending on model inputs and API configurations.

[1] Results sourced from providers' self-reported benchmark numbers.



[AWS News Blog Post](#)



Other Selected Announcements



1. [Amazon SageMaker supports migration of Notebook instances to latest versions](#) (re:Invent)
2. [Amazon Bedrock OpenAI GPT OSS models now supports Responses API](#) (re:Invent)
3. [Amazon Bedrock adds 18 fully managed open weight models](#) (re:Invent)
4. **Amazon SageMaker HyperPod:**
 - a) [Support for programmatic node reboot and replacement](#) (pre-re:Invent)
 - b) [Support for Managed tiered KV cache and intelligent routing](#) (pre-re:Invent)
 - c) [Manage HyperPod clusters with the new Amazon SageMaker AI MCP Server](#) (pre-re:Invent)
 - d) [Support for Spot Instances](#) (pre-re:Invent)
 - e) [Support for NVIDIA Multi-Instance GPU \(MIG\) for generative AI tasks](#) (pre-re:Invent)
 - f) [Support for running IDEs and Notebooks to accelerate AI development](#) (pre-re:Invent)
5. [Claude Opus 4.5 now available in Amazon Bedrock](#) (pre-re:Invent)
6. [Amazon Polly expands generative TTS engine with more languages](#) (pre-re:Invent)
7. [Amazon Bedrock Data Automation supports 10 additional languages](#) (pre-re:Invent)
8. [Amazon Bedrock AgentCore Browser now reduces CAPTCHAs with Web Bot Auth](#) (pre-re:Invent)
9. [TwelveLabs' Marengo Embed 3.0 for advanced video understanding now in Bedrock](#) (pre-re:Invent)

App Integration and Serverless

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



MCP proxy support for API Gateway

TRANSFORM REST APIS INTO MCP-COMPATIBLE ENDPOINTS WITH AGENTCORE GATEWAY

[Documentation](#) > [Amazon API Gateway](#) > [Developer Guide](#)

Add an API Gateway REST API as a target for Amazon Bedrock AgentCore Gateway

[PDF](#) [RSS](#) ☐ Focus mode

An Amazon Bedrock AgentCore Gateway provides AI agent developers a secure way to expose your API Gateway REST APIs as Model Context Protocol (MCP)-compatible tools. AgentCore Gateway uses targets to define tools. When you add your stage as a target, your Gateway becomes a single MCP URL that enables access to the tools for an agent. For more information, see [API Gateway REST API stages as targets](#) in the *Amazon Bedrock AgentCore Gateway Developer Guide*.

API Gateway targets connect your AgentCore Gateway to stages of your REST APIs. You can include the entire stage as a target, or select resources. After you create the API Gateway target, AgentCore Gateway translates incoming MCP requests into HTTP requests and handles the response formatting. MCP clients can retrieve API documentation using the `tools/list` method and invoke APIs using the `tools/call` method.

This new capability of **Amazon API Gateway** enables organizations to make their APIs accessible to AI agents and MCP clients.

Through integration with *Amazon Bedrock AgentCore Gateway* service, you can securely convert your REST APIs into agent-compatible tools while enabling intelligent tool discovery through semantic search.



GA

Dec, 2nd

re:Invent

9 Regions



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Lambda durable functions

MULTI-STEP RESUMABLE SERVERLESS FUNCTIONS POWERED BY LAMBDA AND STEP FUNCTIONS



GA

Dec, 2nd

re:Invent

US East (Ohio)

▼ Function overview Info

Export to Infrastructure Composer

Download

Diagram

Template

demo-durable-function-3

Layers (0)

+ Add trigger

+ Add destination

Type

Durable

Description

-

Last modified

23 seconds ago

Function ARN

arn:aws:lambda:us-east-1:123456789012:function:demo-durable-function-3

Function URL

Info

Code

Test

Monitor

Durable executions - new

Configuration

Aliases

Versions

Code source Info

Open in Visual Studio Code

Upload from

demo-durable-function-3

lambda_function.py

```
1 from aws_durable_execution_sdk_python.config import Duration
2 from aws_durable_execution_sdk_python.context import DurableContext, StepContext, logger
3 from aws_durable_execution_sdk_python.execution import durable_execution
4
5 @durable_step
6 def my_step(step_context: StepContext, my_arg: int) -> str:
7     step_context.logger.info("Hello from my_step")
8     return f"from my_step: {my_arg}"
9
10 @durable_execution
11 def lambda_handler(event, context) -> dict:
12     msg: str = context.step(my_step(123))
13
14     context.wait(Duration.from_seconds(10))
15
16     context.logger.info("Waited for 10 seconds without consuming CPU.")
17
18     return {
19         "statusCode": 200,
20         "body": msg,
21     }
```

DEPLOY

Current

Deploy

Test

TEST EVENTS [NONE SELECTED]

+ Create new test event

Event ID	Event type	Durable operation	Start	Time from first event
1	ExecutionStarted	-	Nov 28, 2025, 13:42:07.331 (UTC-08:00)	-
2	StepStarted	validate_order	Nov 28, 2025, 13:42:10.298 (UTC-08:00)	2 secs 967 ms
3	StepSucceeded	validate_order	Nov 28, 2025, 13:42:10.298 (UTC-08:00)	2 secs 967 ms
4	CallbackStarted	awaiting-approval	Nov 28, 2025, 13:42:10.487 (UTC-08:00)	3 secs 156 ms
5	StepStarted	send_for_approval	Nov 28, 2025, 13:42:10.678 (UTC-08:00)	3 secs 347 ms
6	StepSucceeded	send_for_approval	Nov 28, 2025, 13:42:10.678 (UTC-08:00)	3 secs 347 ms
7	InvocationCompleted	-	Nov 28, 2025, 13:42:10.804 (UTC-08:00)	3 secs 473 ms
8	CallbackSucceeded	awaiting-approval	Nov 28, 2025, 13:44:45.206 (UTC-08:00)	2 mins 37 secs 87 ms
9	StepStarted	process_order	Nov 28, 2025, 13:44:45.906 (UTC-08:00)	2 mins 38 secs 57 ms
10	StepFailed	process_order	Nov 28, 2025, 13:44:45.906 (UTC-08:00)	2 mins 38 secs 57 ms
11	InvocationCompleted	-	Nov 28, 2025, 13:44:46.039 (UTC-08:00)	2 mins 38 secs 70 ms
12	StepStarted	process_order	Nov 28, 2025, 13:44:55.451 (UTC-08:00)	2 mins 48 secs 12 ms
13	StepSucceeded	process_order	Nov 28, 2025, 13:44:55.451 (UTC-08:00)	2 mins 48 secs 12 ms
14	InvocationCompleted	-	Nov 28, 2025, 13:44:55.587 (UTC-08:00)	2 mins 48 secs 25 ms
15	ExecutionSucceeded	-	Nov 28, 2025, 13:44:55.587 (UTC-08:00)	2 mins 48 secs 25 ms

AWS Lambda announces durable functions, enabling developers to build reliable multi-step applications within the Lambda developer experience.

Durable functions automatically checkpoint progress, suspend execution for up to one year during long-running tasks, and recover from failures - all without requiring you to manage additional infrastructure or write custom state management and error handling code.



[AWS News Blog Post](#)



AWS Lambda Managed Instances

RUN FUNCTIONS ON AMAZON EC2 INSTANCES WHILE MAINTAINING OPERATIONAL SIMPLICITY

aws [Option+S] [Search] [United States (N. Virginia)]

Lambda > Functions > Create function

▼ **Additional configurations**
Use additional configurations to set up networking, security, and governance for your function. These settings help secure and customize your Lambda function deployment.

Compute type
Select a compute type: Serverless with pay-per-use pricing and automatic scaling, or Enhanced Scaling on EC2 instances with dedicated capacity and different pricing. This setting can't be changed after function creation.

☐ Lambda (default)
Automatically scales and runs code without server management. Pay only for the compute time you use.

☒ **Lambda managed instances - new**
Run on EC2 instances managed by Lambda with enhanced compute scaling. You must create a capacity provider first to use this option. Different pricing model applies based on EC2 usage. [View pricing](#).

Capacity provider ARN
Select or [create a capacity provider](#) for your Lambda function. Capacity providers offer extended capabilities for large-scale applications that need to go beyond standard Lambda limits.

[C]

Memory size
Your function is allocated vCPU proportional to the memory configured.

GB
Set memory between 2 GB and 32 GB.

Execution environment memory (GiB) per vCPU ratio

☒ 2:1
☐ 4:1
☐ 8:1

AWS Lambda Managed Instances

lets you run functions on your *Amazon EC2* instances while maintaining serverless operational simplicity.

It fully manages all infrastructure tasks, including instance lifecycle, OS and runtime patching, built-in routing, load balancing, and auto-scaling based on configurable parameters - so you can focus on writing code.

This is ideal for customers requiring specialized hardware, as well as those with steady-state or predictable workloads seeking to optimize costs, while maintaining serverless experience.



GA

Nov, 30th

re:Invent

5 Regions




[AWS News Blog Post](#)



Private integration with ALB for API Gateway

AMAZON API GATEWAY REST API NOW SUPPORTS PRIVATE INTEGRATION WITH ALB




GA

Nov, 21st

pre-re:Invent

All Regions



[AWS Compute Blog Post](#)



Enhanced TestState for Step Functions

ENHANCED DEVELOPER EXPERIENCE FOR TESTSTATE API TO SUPPORT LOCAL TESTING



GA

Nov, 20th

pre-re:Invent

All Regions

\$ teststate-api-for-aws-step-functions.sh ×

\$ teststate-api-for-aws-step-functions.sh

```
1  aws stepfunctions test-state --region us-east-1 \
2  --definition '{
3      "Type": "Task",
4      "Resource": "arn:aws:states:::lambda:invoke",
5      "Parameters": { "FunctionName": "process-order" },
6      "End": true
7  }' \
8  --mock '{ "errorOutput": { "error": "Lambda.ServiceException", "cause": "Function failed" } }' \
9  --inspection-level DEBUG
10
11 # Command output:|
12 #
13 # {
14 #   "error": "Lambda.ServiceException",
15 #   "cause": "Function failed",
16 #   "inspectionData": {
17 #     "input": "{}",
18 #     "afterInputPath": "{}",
19 #     "afterParameters": "{$\"FunctionName\": \"process-order\"}"
20 #   },
21 #   "status": "FAILED"
22 # }
```



[AWS News Blog Post](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Tenant isolation mode for AWS Lambda

ISOLATE REQUEST PROCESSING FOR INDIVIDUAL TENANTS OR END-USERS INVOKING FUNCTION

AWS Lambda introduces a new tenant isolation mode that extends the existing isolation capabilities.

This new mode extends isolation to the individual tenant within a function. It processes function invocations in separate execution environments for each tenant, enabling you to meet strict isolation requirements without additional implementation.

▼ Additional configurations

Use additional configurations to set up networking, security, and governance for your function. These settings help secure and customize your Lambda function deployment.

Networking

Function URL [Info](#)

Use function URLs to assign HTTP(S) endpoints to your Lambda function.

☐ Enable

❗ Function URL cannot be used with tenant isolation mode. To enable function URL, deselect tenant isolation mode.

VPC [Info](#)

Connect your function to a VPC to access private resources during invocation.

☐ Enable

Security & governance

Tenant isolation mode - new [Info](#)

Enable tenant isolation mode to isolate execution environments between tenants during function invocations. Tenant isolation mode cannot be modified after function creation. [View pricing](#)

☒ Enable

Code signing [Info](#)

Use code signing configurations to ensure that the code has been signed by an approved source and has not been altered since signing.

☐ Enable

Encryption with an AWS KMS customer managed key [Info](#)

By default, Lambda encrypts the .zip file archive using an AWS owned key.

☐ Enable

Tags [Info](#)

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources, track your AWS costs, and enforce attribute-based access control.

☐ Enable

[Cancel](#)

[Create function](#)



GA

Nov, 19th

pre-re:Invent

All Regions




[AWS News Blog Post](#)



REST API Streaming in Amazon API Gateway

PROGRESSIVE RESPONSE PAYLOAD STREAMING CAPABILITY IN AMAZON API GATEWAY REST API



GA

Nov, 19th

pre-re:Invent

All Regions




[AWS Compute Blog Post](#)



Developer Portals in Amazon API Gateway

FULLY MANAGED AND AWS NATIVE DEVELOPER PORTALS FOR IMPROVING API DISCOVERABILITY



GA

Nov, 19th

pre-re:Invent

All Regions

Pet adoption portal

pet adoption

Documentation

Pet adoption

API reference

AdoptAnimals

View cat

View cats

View dogs

View dog

AdoptProcess

Request visit

Request visit

Request visit

Request visit

Request visit

Request visit

Request visit

View cat

GET https://ra8obxcevq.execute-api.us-west-1.amazonaws.com/Prod/cats/{catId}

View a cat for adoption.

Arguments

catId (string): The cat ID. Required: Yes

Request body

The request does not have a request body.

Responses

200

{ "catId": "string", "catAge": number }

View cats

GET https://ra8obxcevq.execute-api.us-west-1.amazonaws.com/Prod/cats

View all the cats available for adoption.

Request body

The request does not have a request body.

Last update published on: November 17, 2025, 20:59 (UTC-08:00)

API Gateway

APIs

Custom domain names

Domain name access associations

VPC links

Usage plans

API keys

Client certificates

Settings

API publishing

Developer portals

Products

API Gateway

Developer portals

ACME Developer Center

Successfully created the portal ACME Developer Center

Actions

Edit portal

Generate preview

Publish portal

Portal details

Name

ACME Developer Center

Description

A developer portal for internal ACME asset management API tools.

Date created

February 6, 2025, 10:24 (UTC+3:30)

Authentication

Cognito

Domain name

-

Documentation status

132/257 endpoints across 4 products have been documented 31%

Status

Unpublished

Changes

Unpublished changes

Date published

-

Portal home page

Product documentation

Portal design

Tags

Portal home page

The portal home page is the first page your customers will see when they visit your developer portal.

Edit home page

Welcome to ACME Developer Center

A developer portal for internal ACME asset management API tools.

Products (4)

Find Products or APIs

ACME Transfer

ACME Inventory Wizard

ACME DataVis

Powersearch Connect

Last update published on: (unpublished) © 2025



[AWS Compute Blog Post](#)



AWS Lambda adds GA support for Rust

GENERAL AVAILABILITY OF THE RUST RUNTIME ON AWS LAMBDA WITH SUPPORT AND SLA

You can develop your AWS Lambda functions with Rust as you would do previously (e.g., using Cargo Lambda tool or AWS CDK construct), but now it gains official SLA and support by AWS.

aws

Search

[Option+S]

Europe (Ireland)

Lambda > Functions > Create function

1 2

Create function

Info

Choose one of the following options to create your function.

☒ Author from scratch

Start with a simple Hello World example.

☐ Use a blueprint

Build a Lambda application from sample code and configuration presets for common use cases.

☐ Container image

Select a container image to deploy for your function.

Basic information

Function name

Enter a name that describes the purpose of your function.

myFunctionName

Function name must be 1 to 64 characters, must be unique to the Region, and can't include spaces. Valid characters are a-z, A-Z, 0-9, hyphens (-), and underscores (_).

Runtime

Info

Choose the language to use to write your function. Note that the console code editor supports only Node.js, Python, and Ruby.

Amazon Linux 2023

Q Rust

Latest supported

Amazon Linux 2023

OS-only runtime for Go, Rust, C++, custom

Other supported

Amazon Linux 2

OS-only runtime for Go, Rust, C++, custom

Permissions

Info

By default, Lambda will create an execution role with permissions to upload logs to Amazon CloudWatch Logs. You can customize this default role later when adding triggers.

Change default execution role



GA

Nov, 14th

pre-re:Invent

All Regions



[AWS News Blog Post](#)



Visual Rule Builder for Amazon EventBridge

ENHANCED, DRAG & DROP CANVAS BASED VISUAL RULE BUILDER WITH EVENT CATALOG



GA

Nov, 14th

pre-re:Invent

All Regions

aws

Search

[Option+S]

Europe (Ireland)

Application Integration

Amazon EventBridge

A serverless service for building event-driven applications

Amazon EventBridge is a serverless service that uses events to connect application components together, making it easier for developers to build scalable event-driven applications.

Get started

☒ EventBridge Rule with event pattern

A rule matches incoming events and sends them to targets for processing.

☐ EventBridge Scheduled rule

A rule that will invoke a target at a scheduled time.

☐ EventBridge Pipe

A pipe connects an event source to a target with optional filtering and enrichment.

☐ EventBridge Schedule

A schedule invokes a target one-time or at regular intervals defined by a cron or rate expression.

☐ EventBridge Schema registry

Schema registries collect and organize schemas.

Create rule

Pricing

There is no up-front commitment or minimum fee. You simply pay for what you use and will be charged at the end of the month for your usage. Pricing varies by region.

Cost calculator

Learn

Introduction to Amazon EventBridge

Amazon EventBridge Blog

How it works

Serverless 101: Amazon EventBridge

Serverless 101

Amazon EventBridge

Kopiuuj link

CloudShell

Feedback

Console Mobile App

© 2025, Amazon Web Services, Inc. or its affiliates.

Privacy

Terms

Cookie preferences



AWS Announcement



AWS Serverless MCP Server supports ESM

GUIDE YOUR AI AGENTS ABOUT EVENT SOURCE MAPPINGS (ESM) FOR AWS LAMBDA

Create a new Kafka cluster and a VPC named demo-vpc in us-east-1. The cluster should be in the VPC's private subnets. Then, create a Lambda function to consume from the stream within the same VPC cluster. Prefix all created resources with esm-demo



Kiro

I'll help you create a Kafka cluster with VPC and Lambda function for event streaming. Let me start by getting guidance on setting up this infrastructure and then generate the necessary templates.

>  Called MCP tool `esm_guidance`

Now let me generate the secure IAM policy for MSK and create the complete SAM template for your Kafka streaming setup:

>  Called MCP tool `secure_esm_msk_policy`

AWS Serverless MCP Server now also guides AI-powered assistance about event source mappings (ESM), in addition to the current set of serverless application development capabilities – including architecture guidance, deployment automation, and infrastructure provisioning.



GA

Oct, 30th

pre-re:Invent

Open Source



[AWS News Blog Post](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Other Selected Announcements



1. [AWS Lambda adds support for Node.js 24](#) (pre-re:Invent)
2. [AWS Lambda announces enhanced error handling for Kafka](#) (pre-re:Invent)
3. [Additional TLS security policies for Amazon API Gateway REST APIs](#) (pre-re:Invent)
4. [AWS Lambda adds support for Python 3.14](#) (pre-re:Invent)
5. [Amazon MQ now supports LDAP authentication for RabbitMQ](#) (pre-re:Invent)
6. [AWS Lambda announces Provisioned Mode for SQS event source mapping \(ESM\)](#) (pre-re:Invent)
7. [AWS Lambda adds support for Java 25](#) (pre-re:Invent)
8. [Amazon EventBridge now supports targeting SQS fair queues](#) (pre-re:Invent)
9. [AWS Health enhances Amazon EventBridge to give more flexibility and higher resilience](#) (pre-re:Invent)
10. [AWS Step Functions announces a new metrics dashboard](#) (pre-re:Invent)

Developer Tools, Management, and Governance

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



Transformation of AWS Support

ENHANCED TIERS AND AI-POWERED OPERATIONS WITH THE HUMAN EXPERTISE YOU TRUST



GA

Dec, 2nd

re:Invent

All Regions

AWS Support				
Overview Support Plans Support Solutions Partners Pricing More				
Compare Support plans				
Plan category	Plan details	Business Support+	Enterprise Support	Unified Operations
Recommended for	Best for	Minimum recommended plan for production workloads by AWS.	Recommended plan for business-critical workloads across organizations looking for expert guidance.	Recommended plan for mission-critical workloads that require enhanced resilience and application-specific expertise.
AI-led troubleshooting	Unlimited 24/7 contextual recommendations	✓	✓	✓
Human response times*	Business/Mission-critical system down	< 30 mins	< 15 mins	< 5 mins from Incident Management Engineer ****
	Production system down	< 1 hour	< 1 hour	< 1 hour

Business+

30-minute response times
AI-powered assistance 24/7
\$29/month starting price

Enterprise Support

Designated TAMs
Security Incident Response included
\$5,000/month starting price

Unified Operations

5-minute response times
Global team of experts
\$50,000/month starting price



[AWS News Blog Post](#)



AWS DevOps Agent

FRONTIER AI AGENT FOR OPERATIONAL EXCELLENCE

AWS DevOps Agent

Preview

AWS News Blog - Lambda App

DevOps Center

Incident Response

Pre

Investigation 2025-11-21T07:37:52.391Z

Created 2025-11-21 07:40:08 by AWS DevOps Agent

Ask for human support

Analyze Lambda function errors and execution details

Topology

Full Screen

us-east-1

486622066892 - [us-east-1]

AWS DevOps Lambda

AWS CloudWatch Alarm

AWS-AIDevOps-Lambda-Error-Test

AWS Lambda Function

AWS-AIDevOps-test-lambda

AWS IAM Role

AWS-AIDevOps-LambdaTestRole

Standalone Resources

AWS Lambda Log Group

/aws-logs-123456789012-us-east-1/lambda-logs-AIDevOps-test-lambda

AWS-AIDevOps-Lambda-Error-Test alarm triggered due to Lambda errors

Update

5 minutes ago

User Request

TITLE: Investigation 2025-11-21T07:37:52.391Z, DESCRIPTION: Investigation details: There are errors on my API in us-east-1 Investigation starting point: Check if the Lambda functions works as expected Date and time of incident: 2025-11-21T07:30:52.391Z AWS Account ID for incident: [redacted] Region where the incident occurred: us-east-1

PP

Dec, 2nd

re:Invent

US East
(N. Virginia)

AWS DevOps Agent is a new frontier agent that resolves and prevents incidents autonomously.

It reduces MTTR by triaging incidents and guiding teams to resolution. It integrates with *Amazon CloudWatch*, *Datadog*, *Dynatrace*, *GitHub Actions*, and *GitLab CI/CD*, and it analyses historical incidents to provide actionable recommendations.

[AWS News Blog Post](#)

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

155

Unified management in CloudWatch

UNIFIED DATA MANAGEMENT AND ANALYTICS FOR OPERATIONS, SECURITY, AND COMPLIANCE



GA

Dec, 2nd

re:Invent

All Regions

CloudWatch > Log management

Log groups | **Data sources - new** | Summary

How it works

Your log data is seamlessly discovered by data source, and data type. To further enrich or transform your data, attach a pipeline. For enhanced analytics and efficient querying, using field indexes and facets in Logs Insights. To enable new data sources, go to [Ingestion](#). Learn more about [Data sources](#), [Facets](#), and [S3 Tables integration](#). To map additional data sources, go to "Unmapped log data" tab below.

Enable data sources

Data sources

Unmapped log data

Data sources (30) Info

Search by source or type

Filter by category

Any

Associate with S3 Tables integration

View in Logs Insights

☐	Data source	Category	Log groups	Pipeline	Event f...	Source ...	Source type
☐	amazon_route53_private_dns_quer...	AWS	1 log group	No pipeline	Default	amazon_r...	private_dns_query
☒	amazon_vpc_flow	AWS	4 log groups	No pipeline	Default	amazon_vpc	flow
☒	apache_access	Custom	1 log group	No pipeline	Default	apache	access
☒	aws_cloudtrail_data	AWS	1 log group	No pipeline	Default	aws	
☐	aws_cloudtrail_management	AWS	1 log group	No pipeline	Default	aws	
☐	aws_nlb_access	AWS	1 log group	No pipeline	Default	aws	
☐	Bedrock_AC	Custom	1 log group	No pipeline	Default	Bed	
☐	microsoft_entraid_account_change	Custom	2 log groups	No pipeline	Default	microsoft	

Amazon CloudWatch now provides unified management for operational, security, and compliance data across AWS and third-party sources.

It eliminates need for multiple separate log stores and complex ETL pipelines, it provides organization-wide enablement for AWS log sources like CloudTrail, VPC, and WAF, and Apache Iceberg-compatible access through Amazon S3 Tables at no additional storage charge.

Step 1 of 5

General settings Info

Select data source

AWS CloudTrail logs

Pipeline Name

cloudtrail-log-pipeline-test

Must start with a lowercase letter and contain only lowercase letters, numbers, and hyphens (3-28 characters)

Tags

No tags associated with the resource.

Add new tag

You can add up to 50 tags.

Cancel

Next

Step 2 of 5

Configure data source Info

AWS CloudTrail logs

Monitor AWS API calls and user activity across your AWS infrastructure.

IAM Role

Q '498258:role/CloudTrailLogsRole X

Create IAM role

Log source name

aws_cloudtrail

Log source type

Data

Management

Cancel

Previous

Next

Step 3 of 5

Configure destination Info

Select destination

Cloudwatch Logs

Log group

Original log group

Cancel

Previous

Next

Step 4 of 5

Configure processors - optional Info

Add more processors - optional

Processors perform an action on your data, such as filtering, transforming or enriching.

Select processor

Q |

Add

Parse

OCSF processor

This processor converts logs into Open Cybersecurity Schema Framework (OCSF) events.

Mutate Events

Enrich

Mutate String

Processors - optional

processors, including 1 parsers, Copy Values.

Previous

Next



AWS News Blog Post



AWS IAM Policy Autopilot

OPEN SOURCE MODEL CONTEXT PROTOCOL (MCP) SERVER AND CLI TOOL FOR AUTHORIZING IAM

AWS also announced an **AWS IAM Policy Autopilot**, which is an open source *Model Context Protocol (MCP)* server and command-line tool that helps your AI coding assistants quickly create baseline IAM policies that you can refine as your application evolves, so you can build your applications faster.

```
~ kiro-cli
✓ iam-policy-autopilot loaded in 0.06 s
```

KIRO

Did you know?

Ask me directly about my capabilities! Try questions like "What can you do?" or "Can you save conversations?"

```
Model: Auto (/model to change)
```

```
> █
```



GA

Nov, 30th

re:Invent

Open Source



Public preview of AWS MCP Server

FULLY MANAGED REMOTE MODEL CONTEXT PROTOCOL (MCP) SERVER FOR API AND DOCS



PP

Nov, 30th

re:Invent

US East
(N. Virginia)

[Documentation](#) > [AWS MCP Server](#) > [User Guide](#)

What is the AWS MCP Server?

PDF

RSS



Focus mode

The AWS MCP Server is a fully managed remote Model Context Protocol (MCP) server that provides AI assistants and agents with secure, authenticated access to AWS services through natural language interactions. You can use the AWS MCP Server to perform complex, multi-step AWS tasks by combining real-time access to AWS documentation, syntactically correct API calls, and pre-built workflows called Agent SOPs that follow AWS best practices.

With the AWS MCP Server, you can ask AI assistants to provision infrastructure, troubleshoot issues, configure services, and manage AWS resources without needing to know specific API syntax or remember complex procedures. The server handles authentication through standard AWS Identity and Access Management (IAM) controls and provides comprehensive audit logging through AWS CloudTrail.

The AWS MCP Server consolidates capabilities from existing MCP servers (AWS Knowledge MCP and AWS API MCP) into a single, unified interface that reduces configuration complexity while improving AI agent effectiveness across multi-service AWS workflows.

AWS MCP Server consolidates capabilities from the existing *AWS API MCP* and *AWS Knowledge* servers into remotely accessible and fully managed *Model Context Protocol (MCP)* server, and a unified interface, providing access to documentation, executing AWS API calls, and following pre-built workflows called *Agent Standard Operating Procedures (SOPs)*, that guide AI agents through common tasks on AWS.



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

[AWS Announcement](#)

Delete protection for logs in CloudWatch

YOU CAN NOW PROTECT YOUR AMAZON CLOUDWATCH LOG GROUPS FROM DELETION

Amazon CloudWatch now offers configuring deletion protection on your log groups, helping customers safeguard their critical logging data from accidental or unintended deletion.

This feature provides an additional layer of protection for logs maintaining audit trails, compliance records, and operational logs that must be preserved.

Create log group

Log group details [Info](#)

CloudWatch Logs offers three log classes: Standard, Infrequent Access, and Delivery. [Learn more about the features offered by each log class.](#)

Log group name
new-log-group-with-deletion-protection

Retention setting
1 week (7 days)

Log class [Info](#)
Standard

KMS key ARN - optional

☒ **Deletion protection**
Deletion protection is turned off by default. Deletion protection helps prevent accidental or unintended deletion of log groups that store critical operational or compliance data.

Tags
A tag is a label that you assign to an Amazon Web Services resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your Amazon Web Services costs.

No tags are associated with this log group.

[Add new tag](#)
You can add up to 50 more tag(s).

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




GA

Nov, 26th

pre-re:Invent

All Regions



AWS Announcement



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Automated Management in Service Quotas

ADJUST SERVICE QUOTAS BASED ON CUSTOMER'S USAGE AND RECEIVE NOTIFICATIONS



GA

Nov, 25th

pre-re:Invent

All Regions

[Documentation](#) > [Service Quotas](#) > [User Guide](#)

Service Quotas Automatic Management

[PDF](#)

[RSS](#)

☐ Focus mode

Service Quotas Automatic Management monitors your service quotas usage and notifies you before you run out of your allocated quotas. You gain better visibility and proactive awareness, enabling you to run your applications without interruptions.

Key features of Automatic Management

Opt-in options

Enable Automatic Management using the Service Quotas console, AWS CLI, or API.

Usage notifications

Receive notifications when quota usage reaches the following thresholds:

- 80% utilization
- 95% utilization

Notification channels

Configure notifications through multiple channels:

- AWS Console Mobile Application
- Email
- Slack

Integration options

- Subscribe to [Amazon EventBridge events](#) for automation workflows
- View notifications in the [AWS Health](#) personal health dashboard




[AWS Announcement](#)



AWS CodeCommit returns to GA

BASED ON THE CUSTOMER FEEDBACK, CODECOMMIT RETURNS TO GA EFFECTIVE IMMEDIATELY



GA

Nov, 24th

pre-re:Invent

All Regions

AWS DevOps & Developer Productivity Blog

The Future of AWS CodeCommit

by Anthony Hayes | on 24 NOV 2025 | in [Announcements](#), [AWS CodeCommit](#), [Developer Tools](#) | [Permalink](#) | [Share](#)

Back in July 2024, we announced plans to de-emphasize AWS CodeCommit based on adoption patterns and our assessment of customer needs. We never stopped looking at the data or listening to you, and what you've shown us is clear: you need an AWS-managed solution for your code repositories. Based on this feedback, CodeCommit is returning to full General Availability, effective immediately.

We Listened, and We Heard You

After the de-emphasis announcement last year, we heard from many of you. Your feedback was direct and revealing. You told us that CodeCommit isn't just another code repository for you—it's a critical piece of your infrastructure. Its deep IAM integration, VPC endpoint support, CloudTrail logging, and seamless connectivity with CodePipeline and CodeBuild provide significant value, especially for teams operating in regulated industries or those who want all their development infrastructure within AWS boundaries. In short, we learned that CodeCommit is essential for many of you, so we're bringing it back.



[AWS DevOps Blog Post](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Deployment ordering for StackSets

AWS CLOUDFORMATION STACKSETS NOW ALLOW YOU TO DEFINE SEQUENCE OF DEPLOYMENTS



GA

Nov, 21st

pre-re:Invent

All Regions

Auto-deployment options

Configure automatic deployment settings for your stack sets.

Automatic deployment

With automatic deployment enabled, if an account is added to an OU, StackSets automatically deploys additional stack instances to this account. If an account is removed from an OU, StackSets automatically deletes stack instances in this account.

- ☒ Activated
Stack instances will be automatically deployed or deleted when accounts are added or removed from OUs.
- ☐ Deactivated
You will need to manually deploy or delete stack instances when accounts are added or removed from OUs.

Account removal behavior

When an account is removed from a target OU, should stack instances in the account be deleted or retained?

- ☒ Delete stacks
Stack instances will be automatically deleted when an account is removed from a target OU.
- ☐ Retain stacks
Stack instances will remain in the account when it's removed from a target OU.

Stack set dependencies - optional

Stack sets with dependencies will automatically deploy stack instances only after all dependent stack sets complete their deployments in the same account and partition. This enables sequencing automatic deployments across stack sets. You can set up to 10 dependencies for each stack set and up to 100 dependencies in the account.

 If a dependency stack set does not exist or has automatic deployments disabled, it will be ignored. However, any transitive dependencies through that stack set will still be enforced.

Stack set ARN

arn:aws:cloudformation:us-west-2:123456789012:stackset/example-stack-set

Remove

Add another stack set ARN

You can add 9 more stack set ARNs.



[AWS DevOps Blog Post](#)



Dedicated-controls for AWS Control Tower

APPLY AWS MANAGED CONTROLS WITHOUT FULL AWS CONTROL TOWER DEPLOYMENT



Nov, 21st

pre-re:Invent

All Regions

The setup of your landing zone is complete.

Your landing zone is now available.

AWS Control Tower has set up the following:

- 2 organizational units, one for your shared accounts and one for accounts that will be provisioned by your users.
- 3 shared accounts, which are the management account and isolated accounts for log archive and security audit.
- Your selected identity and access management configuration.
- 23 preventive controls to enforce policies and 3 detective controls to detect configuration violations.

Enroll existing accounts in AWS Control Tower

You can enroll existing accounts from your AWS Organizations organization in AWS Control Tower. Some additional work is required for enrollment.

[Learn more about enrolling existing accounts in AWS Control Tower](#)

Recommended actions

Programmatically manage controls with APIs

AWS Control Tower customers can manage controls programmatically across their organization, at scale.

Add or register organizational units

Add or register organizational units (OUs) to organize accounts and projects.

Enable more controls

Enable additional controls to meet your security, operational, and compliance requirements.

Review users and access

Review your user identity store and single sign-on access for your users across accounts.

Review shared accounts

Review the settings of the shared accounts that AWS Control Tower has set up for you.

Environment summary

0

Organizational units

2

Accounts

Controls summary (10)

10

Preventive controls

0

Detective controls

0

Proactive controls

[View enabled controls](#)

Noncompliant resources

< 1 >

⊗

Resource ID	Resource type	Service	Region	Account name	Organizational unit	Control
No noncompliant resources found						
No noncompliant resources detected on controls with a Clear status.						

Registered organizational units

Find organizational units

< 1 >

Name	Parent organizational unit	Control status	AWS Control Tower baseline status	AWS Config baseline status	Compliance
Root	-	-	Enabled	Not enabled	Compliant
MyProdOU	Root	Controls enabled	Not enabled	Not enabled	Compliant

[View all organizational units](#)

Enrolled accounts

Find accounts

< 1 >

Account name	Account email	Organizational unit	Control status	Owner	Compliance status	AWS Control Tower baseline status	AWS Config baseline status
awsb-utopia-console-only-private-1-us-east-1	awsb-utopia-1-conf@amazon.com	MyProdOU	Controls enabled	AWS Control Tower	Compliant	Not enabled	Not enabled
awsb-utopia-console-only-private-1-us-east-1	awsb-utopia-console-only-private-1@amazon.com	Root	-	AWS Control Tower	Compliant	Not enabled	Not enabled

With this feature, you can use AWS managed controls without the need to set up resources you don't need, which means you get started faster if you already have an established multi-account environment and want to use **AWS Control Tower** only for its managed controls



[AWS News Blog Post](#)



AWS CloudTrail Insights for Data Events

IDENTIFY AND RESPOND TO UNUSUAL ACTIVITIES ASSOCIATED WITH DATA PLANE EVENTS



GA

Nov, 20th

pre-re:Invent

All Regions

[Documentation](#) > [AWS CloudTrail](#) > [User Guide](#)

Working with CloudTrail Insights

PDF

RSS

Focus mode

AWS CloudTrail Insights help AWS users identify and respond to unusual activity associated with API call rates and API error rates by continuously analyzing CloudTrail management and data events. CloudTrail Insights analyzes your past management and data events to establish your normal patterns of API call rates and API error rates, also called the *baseline*. CloudTrail then generates Insights events when the current API call rates or error rates deviate from the baseline.

You can collect two types of Insights, each for management and data events.

Management events Insights

- **API call rate** – A measurement of write-only management API calls that occur per minute against a baseline API call volume. To log Insights events on the API call rate for management events, the trail or event data store must enable Insights and log `write` management events.
- **API error rate** – A measurement of management API calls that result in error codes. The error is shown if the API call is unsuccessful. To log Insights events on API error rate, the trail or event data store must enable Insights and log `read` or `write` management events, or both `read` and `write` management events.

Data events Insights

- **API call rate** – A measurement of data API calls that occur per minute against a baseline API call volume. To log Insights events on the API call rate, the trail must enable Insights and log data events.
- **API error rate** – A measurement of data API calls that result in error codes. The error is shown if the API call is unsuccessful. To log Insights events on API error rate, the trail must enable Insights and log `read` or `write` data events, or both `read` and `write` data events.

Note

Insights events on data events are only supported on trails and not on event data stores.



Required Tags for Organizations Tag Policies

PROACTIVE VALIDATION CHECK THAT YOUR INFRASTRUCTURE AS CODE APPLIES CORRECT TAGS



GA

Nov, 20th

pre-re:Invent

All Regions

[Documentation](#) > [AWS Organizations](#) > [User Guide](#)

Enforce "Required tag key" with IaC

[PDF](#)

[RSS](#)

☐ Focus mode

Tag policies help you maintain consistent tagging across your infrastructure as code (IaC) deployments. With "Required tag keys", you can ensure that all resources created through IaC tools like CloudFormation, Terraform, and Pulumi include the mandatory tags defined by your organization.

This capability checks your IaC deployments against your organization's tag policies before resources are created. When a deployment is missing required tags, you can configure your IaC settings to either warn your development teams or prevent the deployment entirely. This proactive approach maintains tagging compliance from the moment resources are created, rather than requiring manual remediation later. The enforcement works across multiple IaC tools using a single tag policy definition, eliminating the need to configure separate tagging rules for each tool your organization uses.

AWS Organizations Tag Policies announces *Reporting for Required Tags*, a new validation check that proactively ensures your *AWS CloudFormation*, *Terraform*, and *Pulumi* deployments include the required tags critical to your business.




[AWS Announcement](#)




Billing Transfer for AWS Organizations

CENTRALLY MANAGE AND PAY BILLS ACROSS MULTIPLE AWS ORGANIZATIONS



With the general availability of **Billing Transfer**, a new capability to centrally manage and pay bills across multiple organizations by transferring payment responsibility to other billing administrators, it introduces a new way for centralizing billing procedures for AWS Customers.



GA

Nov, 19th

pre-re:Invent

All Regions
excluding
AWS GovCloud and
Chinese regions



[AWS News Blog Post](#)



AWS CloudFormation Language Server

INTELLIGENT AUTHORING, EARLY VALIDATION, TROUBLESHOOTING DIRECTLY IN YOUR IDE



GA

Nov, 19th

pre-re:Invent

All Regions

[Documentation](#) > [AWS CloudFormation](#) > [User Guide](#)

AWS CloudFormation language server



Focus mode

The AWS CloudFormation language server provides capabilities to accelerate authoring infrastructure-as-code (IaC) and deploying AWS resources safely and with confidence. It follows the [language server protocol](#)¹ (LSP) to provide documentation on hover, auto-completion, diagnostics via static validation, go to definition and code actions. In addition to these traditional language server capabilities, the server adds online features to explore and deploy AWS resources via CloudFormation. This includes the ability to validate and deploy templates using change sets; view stack diffs, events, resources, and outputs; list stacks and browse resources by type; and insert live resource state directly into CloudFormation templates. Developer tooling and integrated development environments (IDE) can leverage the language server capabilities to enhance the experience of IaC developers on CloudFormation.

AWS CloudFormation announces the launch of the *AWS CloudFormation Language Server*, a new capability that brings intelligent authoring, early validation, troubleshooting, and drift management directly into IDEs through the *AWS Toolkit* plug-in.



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Drift-aware Change Sets in CloudFormation

CHANGE SETS CAN NOW COMPARE TEMPLATE WITH THE ACTUAL STATE AND ALIGN RESOURCES



GA

Nov, 18th

pre-re:Invent

All Regions

aws

Search

[Option+S]

Europe (Frankfurt)

CloudFormation > Stacks > testing-cloudformation-drift-aware-changesets

testing-cloudformation-drift-aware-changesets

Delete

Update stack

Stack actions

Create stack

Stack info

Events

Resources

Outputs

Parameters

Template

Change sets

Git sync

Overview

Stack ID

arn:aws:cloudformation:eu-central-1:381492054569:stack/testing-cloudformation-drift-aware-changesets/9d6822e0-cf7b-11f0-82c0-063dc5a155b7

Description

-

Status

CREATE_COMPLETE

Detailed status

-

Status reason

-

Root stack

-

Parent stack

-

Created time

2025-12-02 13:37:07 UTC+0100

Updated time

-

Deleted time

-

Drift status

DRIFTED

Last drift check time

2025-12-02 13:42:00 UTC+0100

Termination protection

Deactivated

IAM role

-

Latest operations

Operation 1

Operation ID

d779a043-d169-4596-ab2b-73052cae8f7f

Operation type

-

CloudShell Feedback

© 2025, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences



AWS Announcement



Early feedback in AWS CloudFormation

CHANGE SETS NOW PROVIDE EARLY VALIDATION AND MORE INSIGHT INTO TROUBLESHOOTING



GA

Nov, 18th

pre-re:Invent

All Regions

How pre-deployment validation works

Pre-deployment validation involves these phases:

1. **Create your change set** – Generate a change set as you normally would for your CloudFormation stack updates. Pre-deployment validation is enabled by default when creating your change set.
2. **Validation execution** – CloudFormation runs multiple validation checks against your template and target environment. Currently 3 types of validation are supported: Property syntax validation against resource schemas, resource name conflict detection with existing resources, and S3 bucket emptiness validation for deletion operations.
3. **Review validation results** – CloudFormation provides detailed feedback on any issues found, including precise path pinpointing the issue location in template, eliminating manual template debugging.
4. **Resolve issues** – Address identified problems by updating your templates or resolving conflicts before proceeding with deployment.
5. **Execute with confidence** – Deploy your change set knowing that common failure scenarios have been validated upfront.

Considerations


As you use pre-deployment validation, keep the following in mind:

- Pre-deployment validation focuses on the three common deployment failure scenarios. It doesn't guarantee that your deployment will succeed, but reduces the likelihood of common failures.
- Validation modes behave differently:
 - **FAIL mode** prevents change set execution when validation detects errors, ensuring problematic templates cannot proceed to deployment. This applies to property syntax errors and resource naming conflicts.
 - **WARN mode** allows change set creation to succeed despite validation failures, providing warnings that developers can review and address before execution. This applies to constraint violations like S3 bucket emptiness that may be resolvable through manual intervention.
- Validation results are tied to the specific change set. If you modify your template, you'll need to create a new change set to get updated validation results.
- S3 bucket validation only checks for object presence, not for bucket policies or other constraints that might prevent deletion.



Kiro is Generally Available

BUILD WITH YOUR TEAM IN THE IDE AND TERMINAL




GA

Nov, 17th

pre-re:Invent

All Regions



[Kiro Blog](#)



Automatic Enrollment in AWS Control Tower

ENROLLING AWS ACCOUNTS TO AWS CONTROL TOWER VIA MOVE TO ORGANIZATIONAL UNIT (OU)



GA

Nov, 10th

pre-re:Invent

All Regions

[Documentation](#) > [AWS Control Tower](#) > User Guide

Move and enroll accounts with auto-enrollment

PDF

RSS

Focus mode

The account auto-enrollment feature is available for landing zones of version 3.1 and above.

If you optionally enable this feature, you can utilize the AWS Organizations APIs and console to move accounts into AWS Control Tower, without creating *inheritance drift*. The account automatically receives baseline resources and control configurations from the destination organizational unit (OU) in AWS Control Tower. This optional capability also allows you to move accounts between OUs within AWS Control Tower, without creating inheritance drift, if the two OUs have the same baseline configuration and the same controls enabled.

To activate auto-enrollment: You can select auto-enrollment of accounts on the landing zone **Settings** page in the AWS Control Tower console, or by calling the AWS Control Tower `CreateLandingZone` or `UpdateLandingZone` APIs, with the value of the `RemediationType` parameter set to **Inheritance Drift**.

To apply auto-enrollment: After selecting this option in your **Settings** page, you can move an account by means of the AWS Organizations console, the AWS Organizations `MoveAccount` API, or the AWS Control Tower console.

To unenroll an account with auto-enrollment: If you move an account outside an OU that's registered, AWS Control Tower removes all deployed baseline resources and controls automatically.



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Shared Memory Metrics in CloudWatch

CLOUDWATCH AGENT SUPPORTS COLLECTION OF SHARED MEMORY METRICS FROM LINUX HOSTS

- `mem` – Optional. Specifies that memory metrics are to be collected. This section is valid only for Linux instances. This section can include the following fields:
 - `drop_original_metrics` – Optional. If you are using the `aggregation_dimensions` field in the `metrics` section to roll up metrics into aggregated results, then by default the agent sends both the aggregated metrics and the original metrics that are separated for each value of the dimension. If you don't want the original metrics to be sent to CloudWatch, you can specify this parameter with a list of metrics. The metrics specified along with this parameter don't have their metrics by dimension reported to CloudWatch. Instead, only the aggregated metrics are reported. This reduces the number of metrics that the agent collects, reducing your costs.
 - `measurement` – Specifies the array of memory metrics to be collected. Possible values are `active`, `available`, `available_percent`, `buffered`, `cached`, `free`, `inactive`, `shared`, `total`, `used`, and `used_percent`. This field is required if you include `mem`.

To see the default units for each `mem` metric, see [Metrics collected by the CloudWatch agent on Linux and macOS instances](#).

Within the entry for each individual metric, you might optionally specify one or both of the following:

- `rename` – Specifies a different name for this metric.
- `unit` – Specifies the unit to use for this metric, overriding the default unit of `None` for the metric. The unit that you specify must be a valid CloudWatch metric unit, as listed in the `Unit` description in [MetricDatum](#).



GA

Nov, 10th

pre-re:Invent



All Regions



[AWS Announcement](#)

Mapping AWS SDK operation to IAM action

AWS SERVICE REFERENCE INFORMATION PROVIDES AN API TO MAP OPERATION TO IAM ACTION



GA

Nov, 4th

pre-re:Invent

All Regions



[AWS Announcement](#)



MCP Proxy for AWS is Generally Available

CLIENT-SIDE PROXY THAT ENABLES MCP CLIENTS TO CONNECT VIA AWS SIGV4 AUTHENTICATION



GA

Oct, 31st

pre-re:Invent
Open Source

README Code of conduct Contributing More

MCP Proxy for AWS


Overview

The **MCP Proxy for AWS** package provides two ways to connect AI applications to MCP servers on AWS:

- Using it as a proxy** - It becomes a lightweight, client-side bridge between MCP clients (AI assistants like Claude Desktop, Amazon Q Developer CLI) and MCP servers on AWS. (See [MCP Proxy](#))
- Using it as a library** - Programmatically connect popular AI agent frameworks (LangChain, LlamaIndex, Strands Agents, etc.) to MCP servers on AWS. (See [Programmatic Access](#))

When Do You Need This Package?

- You want to connect to **MCP servers on AWS** (e.g., using Amazon Bedrock AgentCore) that use AWS IAM authentication (SigV4) instead of OAuth
- You're using MCP clients (like Claude Desktop, Amazon Q Developer CLI) that don't natively support AWS IAM authentication
- You're building AI agents with popular frameworks like LangChain, Strands Agents, LlamaIndex, etc., that need to connect to MCP servers on AWS
- You want to avoid building custom SigV4 request signing logic yourself



Other Selected Announcements (1/2)



1. Amazon CloudWatch:

- a) GenAI observability now supports Amazon AgentCore Evaluations (re:Invent)
- b) Incident reports now support Five Whys analysis (re:Invent)
- c) Application Signals adds GitHub Action and MCP server improvements (pre-re:Invent)
- d) Logs now supports scheduled queries in Logs Insights (pre-re:Invent)
- e) Logs now supports Network Load Balancer access logs (pre-re:Invent)
- f) Composite Alarms adds threshold-based alerting (pre-re:Invent)
- g) Application Signals adds AI-powered Synthetics debugging (pre-re:Invent)

2. AWS Knowledge MCP Server now supports topic-based search (pre-re:Invent)

3. Announcing AWS Compute Optimizer automation rules (pre-re:Invent)


4. AWS Cost Explorer now provides 18-month forecasting and AI-powered forecasts (pre-re:Invent, **PP**)

5. AWS Cost Optimization Hub introduces Cost Efficiency metric (pre-re:Invent)

6. AWS Cost Anomaly Detection expands AWS managed monitoring (pre-re:Invent)

7. AWS CloudTrail adds data event aggregation to simplify security monitoring (pre-re:Invent)

Other Selected Announcements (2/2)



8. Enhanced cost management capabilities in Amazon Q Developer (pre-re:Invent)
9. AWS Organizations introduces direct account transfers between organizations (pre-re:Invent)
10. AWS CloudFormation Hooks invocation summary improvements (pre-re:Invent)
11. AWS CUR 2.0 features: EC2 ODCR and Capacity Blocks for ML monitoring (pre-re:Invent)
12. AWS Fault Injection Service (FIS) launches test scenarios for partial failures (pre-re:Invent)
13. Split Cost Allocation Data for Amazon EKS supports Kubernetes labels (pre-re:Invent)

Security, Identity, and Compliance

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



AWS Security Agent

FRONTIER AI AGENT FOR PROACTIVE APPLICATION SECURITY

Sample design security review

Clone this design review

Download report

Details

Review status

Completed

Date completed

12/2/2025, 8:33:16 AM

Files reviewed

1 uploads

Findings summary

2

Non-compliant

Security requirements are violated or not adequately addressed.

0

Insufficient data

A security requirement's relevance criteria indicate it doesn't apply to this system design.

1

Compliant

Security requirements are met based on the uploaded files.

0

Not applicable

The uploaded files don't contain enough information.

Files reviewed (1)

Download original

Search for file

File name	Type	Size
High level design sample design security review	PDF	137.53 KB

Review findings (3)

Select a finding below to view more details and suggested remediation guidance.

Search by security requirement

Filter security requirements by name

Filter by compliance status

View all

Finding by security requirement	Status
Short Session Timeouts for Privileged and PII Access	Non-compliant
Customer-Managed Encryption Keys Required	Non-compliant
Network Segmentation Strategy Defined	Compliant



PP

Dec, 2nd

re:Invent

US East
(N. Virginia)

AWS Security Agent is a frontier AI agent that proactively secures applications throughout development lifecycle.

It conducts automated security reviews tailored to organizational requirements, it delivers context-aware penetration testing on demand, and continuously validates security preventing vulnerabilities early.




[AWS News Blog Post](#)



Near real-time risk analytics in Security Hub

CORRELATE AND ENRICH SECURITY SIGNALS NEAR REAL-TIME FROM AWS SECURITY SERVICES




GA

Dec, 2nd

re:Invent

All Regions



AWS Security Hub is a unified cloud security solution that prioritizes critical security issues at scale.

It provides near real-time risk analytics with advanced trends and streamlined pricing, by correlating signals from *Amazon GuardDuty*, *Amazon Inspector*, and *AWS Security Hub Cloud Security Posture Management (CSPM)* automatically.



[AWS News Blog Post](#)



Extended Threat Detection for EC2 and ECS


AMAZON GUARDDUTY EXTENDED THREAT DETECTION NOW SUPPORTS AMAZON EC2 AND ECS



Dec, 2nd

re:Invent

All Regions



New Detection Capabilities

EC2 instance group compromise detection
ECS cluster compromise detection

Enhanced Findings

Critical-severity attack sequences
MITRE ATT&CK mapping included

AI-Powered Analysis

Correlates security signals automatically
Detects sophisticated attack patterns



[AWS News Blog Post](#)



aws login

ENABLING BUILDERS TO USE AWS MANAGEMENT CONSOLE CREDENTIALS FOR AWS CLI AND SDK

Getting credentials for local development with AWS is now simpler and more secure.

aws login, lets you start building immediately after signing up for AWS without creating and managing long-term access keys.

You use the same sign-in method you already use for the AWS Management Console up to this point.

Login with console credentials (Recommended)

This authentication method lets you use your console credentials with the AWS CLI, making it easy to get started with AWS programmatically within minutes of account set up. You can get temporary credentials that work seamlessly across local development tools like the AWS CLI, AWS SDKs and AWS Tools for PowerShell.

Prerequisites

- Install the AWS CLI. For more information, see [Installing or updating to the latest version of the AWS CLI](#). A minimum version of 2.32.0 is required to use the `aws login` command.
- Access to sign into the AWS Management Console as a root user, IAM user, or through federation with IAM. If you use IAM Identity Center, go to [Login with IAM Identity Center credentials](#) instead.
- Ensure the IAM identity has the appropriate permissions. Attach the [SignInLocalDevelopmentAccess](#) managed policy to your IAM user, role, or group. If you sign in as a root user, no additional permissions are required.

To login with console credentials

1. Run the following command to start the browser-based authentication process:

```
$ aws login
```

The `aws login` command supports several optional parameters:

- `aws login --remote` - For cross-device authentication when your device doesn't support a browser
 - `aws login --profile profile-name` - To authenticate with a specific profile
 - `aws login --region region` - To authenticate in a specific region
2. Follow the prompts in your terminal. The command will automatically open your default browser and guide you through the authentication process. After successful authentication, your AWS CLI session will be valid for up to 12 hours.
 3. To end your session, use:

```
$ aws logout
```

If you are accessing AWS services programmatically by using AWS Tools for PowerShell, please see [Authenticating the AWS Tools for PowerShell with AWS](#). If you are using AWS SDKs, please see [Authentication and access using AWS SDKs and tools](#).



GA

Nov, 19th

pre-re:Invent

All Regions



[AWS Security Blog Post](#)



AWS IAM Outbound Identity Federation

SECURELY FEDERATE AWS IDENTITIES TO EXTERNAL SERVICES VIA JSON WEB TOKEN (JWT)




GA

Nov, 19th

pre-re:Invent

All Regions



This release provides a new capability called **AWS IAM Outbound Identity Federation**, that customers can use to securely federate their AWS identities to external services without storing long-term credentials.



[AWS News Blog Post](#)



Malware Protection for AWS Backup

AMAZON GUARDDUTY MALWARE PROTECTION EXTENDS SUPPORT WITH AWS BACKUP ARTEFACTS




GA

Nov, 19th

pre-re:Invent

All Regions



[AWS Storage Blog Post](#)



Post-quantum certificates in AWS Private CA

SUPPORT FOR MODULE LATTICE-BASED DIGITAL SIGNATURE ALGORITHM (ML-DSA)



GA

Nov, 10th

pre-re:Invent

All Regions

AWS Private CA now enables you to create *certificate authorities (CAs)* and issue certificates that use Module Lattice-based *Digital Signature Algorithm (ML-DSA)*.

This feature enables you to begin transitioning your *public key infrastructure (PKI)* towards post-quantum cryptography, allowing you to put protections in place now against future quantum computing threats.

What is post-quantum cryptography at AWS?

At AWS, the confidentiality, integrity, and authenticity of our customers' data is a top priority. Today's widely-used public-key cryptographic schemes rely on mathematical problems - like integer factoring and discrete logarithms - that could be efficiently solved by future quantum computers. To address this challenge, AWS is deploying new NIST-standardized post-quantum cryptographic algorithms that are designed to resist both classical and quantum computing attacks. These algorithms, including the Module-Lattice-based Key-Encapsulation Mechanism (ML-KEM) and Module-Lattice-based Digital Signature Algorithm (ML-DSA), are based on different mathematical foundations that are believed to be resistant to quantum computing attacks.

AWS has already deployed post-quantum cryptography across several key services. [AWS Key Management Service \(AWS KMS\)](#), [AWS Secrets Manager](#), and [AWS Certificate Manager](#) have implemented post-quantum hybrid key establishment combining Elliptic Curve Diffie-Hellman (ECDH) with ML-KEM to protect against "harvest now, decrypt later" attacks. At the foundation of these implementations is [AWS-LC](#), our [FIPS-140-3-validated cryptographic library](#), which was the first open-source cryptographic module to include ML-KEM in its FIPS validation.



AWS Announcement



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Simplified M2M pricing in Amazon Cognito

AMAZON COGNITO REMOVES MACHINE-TO-MACHINE (M2M) APP CLIENT PRICE DIMENSION



GA

Nov, 3rd

pre-re:Invent

All Regions

Pricing table for Amazon Cognito add-ons

API quotas

M2M authorization

Cognito Sync

Amazon Cognito supports machine-to-machine (M2M) use cases using the OAuth 2.0 specification's client credentials flow. You can use Amazon Cognito to set up your service (software or an API service represented as an "app client"), establish the app client credentials, and issue access tokens in exchange for these credentials (known as token requests). These access tokens can then be used to communicate with your services. You can configure the validity of the access token for each service. You can also determine token usage per app client.

Amazon Cognito charges you monthly per successful token response. There is no additional charge for the number of app clients you have registered in your account.

Region:

US East (Ohio)

	Number of token requests per month		Price
Tier 1	1	250,000	\$2.250 per 1000 token requests
Tier 2	250,001	5,000,000	\$1.500 per 1000 token requests
Tier 3	5,000,001	And above	\$1.125 per 1000 token requests

* Please contact your account team if you require over 2,500 app clients.

Amazon Cognito announces simplified *Machine-to-Machine (M2M)* pricing model.

Previously, customers were charged for each *M2M app client* registered, regardless of usage. With this change, customers will only pay for their successful token requests.



AWS Announcement




Other Selected Announcements



1. [AWS Security Token Service now supports IPv6](#) (pre-re:Invent)
2. [Amazon Inspector supports organization-wide management](#) (pre-re:Invent)
3. [AWS Secrets Manager announces managed external secrets](#) (pre-re:Invent)
4. [AWS IAM launches aws:SourceVpcArn condition key](#) (pre-re:Invent)

Business Applications

5 announcements (3 *pre-re:Invent*)



Bedrock Knowledge Base support in Connect

BRING AMAZON BEDROCK KNOWLEDGE BASES WITH SUPPORT OF MULTIPLE BASES PER AI AGENT



GA

Nov, 30th

re:Invent

All Regions

Step 3: Create an integration (knowledge base)

1. On the **Amazon Q** page, choose **Add integration**.
2. On the **Add integration** page, choose **Create a new integration**, and then select a source.



The steps for creating an integration vary, depending on the source that you choose. Expand the following sections as needed to finish creating an integration.

► Create an Amazon Simple Storage Service integration

► Create a web crawler integration

Amazon Connect now allows you to bring *Amazon Bedrock Knowledge Bases* and supports multiple knowledge bases per AI agent.

You can now connect your existing Bedrock Knowledge Bases directly to Amazon Connect AI agents in just a few clicks - with no additional setup or data duplication required – to leverage your current data sources and the connectors, including Adobe Experience Manager, Confluence, SharePoint, and OneDrive.



AWS Announcement



MCP support in Amazon Connect

AI AGENTS FOR END-CUSTOMER SELF-SERVICE NOW CAN USE STANDARDIZED TOOLS WITH MCP



GA

Nov, 30th

re:Invent

All Regions

With this launch, **Amazon Connect** provides out-of-the-box MCP tools for common tasks such as updating contact attributes and retrieving case information.

You can also use flow modules as MCP tools to reuse the same business logic across both deterministic and generative AI workflows.

Additionally, you can integrate custom tools or third-party services through flow modules or the *Amazon Bedrock AgentCore Gateway*.

[Documentation](#) > [Amazon Connect](#) > Administrator Guide

Enable AI agents to retrieve information and complete actions with MCP tools

PDF

RSS

Focus mode

Amazon Connect supports Model Context Protocol (MCP), enabling AI agents for both end-customer self-service and employee assistance to use standardized tools for retrieving information and completing actions. With MCP support, you can enhance your AI agents with extensible tool capabilities that reduce contact handle time and increase issue resolution across customer and agent interactions.

MCP provides AI agents with the ability to automatically perform tasks such as looking up order status, processing refunds, and updating customer records during interactions without requiring human intervention. This standardized protocol enables AI agents to access and execute tools from multiple sources while maintaining consistent security and governance controls.

Tool types and integration options

Amazon Connect provides multiple ways to add tools to AI agent configurations:

Out-of-the-box tools

Amazon Connect includes prebuilt tools for common tasks such as updating contact attributes and retrieving case information, enabling immediate functionality without additional configuration.

Flow module tools

You can create new or convert existing flow modules into MCP tools, enabling you to reuse the same business logic across both static and generative AI workflows. Flow modules can connect to third-party sources and integrate with existing business systems.

Third-party MCP tools

You can use third-party integrations through AgentCore Gateway. By registering AgentCore Gateways in the AWS Management Console, similar to how third-party applications are registered to Amazon Connect today, you gain access to whatever tools are available on those servers, including remote MCP servers.



AWS Announcement



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Multi-skill agent scheduling in Connect

OPTIMIZE SCHEDULING BASED ON AGENT'S MULTIPLE SPECIALIZED SKILLS




GA

Nov, 20th

pre-re:Invent

All Regions



Persistent agent connections in Connect

REDUCE TIME TO ESTABLISH CONNECTION WITH END CUSTOMERS WITH PERSISTENT CHANNELS



GA

Nov, 20th

pre-re:Invent

All Regions

[Documentation](#) > [Amazon Connect](#) > Administrator Guide

Enable persistent connection for Amazon Connect agents

↓ PDF

↓ RSS



Focus mode

When **Enable persistent connection** is selected for an agent, after a call ends the agent's softphone maintains its media connection to Amazon Connect for a few minutes. This enables subsequent calls to connect faster. In case the agent remains idle for an extended period of time, the softphone media connection is dropped to reduce agent workstation and network resource consumption. It is re-established upon the next call for this agent.

This functionality doesn't apply to chats or tasks.

Amazon Connect now offers the ability to maintain an open communication channel between your agents and *Connect* instance, helping reduce the time it takes to establish a connection with a customer.



[AWS Announcement](#)



© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Conversational analytics in Amazon Connect

CONVERSATIONAL ANALYTICS FOR CUSTOMER INTERACTIONS ACROSS VOICE AND CHAT BOTS




GA

Nov, 19th

pre-re:Invent

All Regions



Other Selected Announcements (1/2)




1. Amazon SES adds VPC support for API endpoints (re:Invent)
2. **Amazon Connect:**
 - a) Launches WhatsApp channel for Outbound Campaigns (re:Invent)
 - b) Customer Profiles launches new segmentation capabilities (re:Invent, **PP**)
 - c) Enhances its agent assistance capabilities (re:Invent)
 - d) Launches AI-powered predictive insights (re:Invent, **PP**)
 - e) Introduces agentic self-service with enhanced voice interactions (re:Invent)
 - f) Adds support for third-party speech-to-text and text-to-speech AI models (re:Invent)
 - g) Enables business users to create custom UIs to adjust configurations (re:Invent)
 - h) Streams messages for AI-powered interactions (re:Invent)
3. Amazon WorkSpaces Applications now supports IPv6 (pre-re:Invent)
4. Amazon Connect supports ring time configuration for unanswered calls (pre-re:Invent)
5. Amazon Connect now supports enhanced Instance-to-Instance communication (pre-re:Invent)
6. Amazon WorkSpaces add new instance types and configurable storage options (pre-re:Invent)


Other Selected Announcements (2/2)



7. Amazon Connect Cases adds conditional field visibility and dependent options (pre-re:Invent)
8. Amazon Connect now supports configuration of email address aliases (pre-re:Invent)
9. Amazon Connect now supports scheduling of individual agents (pre-re:Invent)



Join **Builder Center** and get access to a **200+ pages PDF** with key re:Invent announcements



builder.aws.com

Developer: Stay connected with AWS!




The AWS Developers Podcast



AWS Skill Builder



AWS Developers
YouTube Channel



[@aws-developers](https://www.linkedin.com/company/aws-developers)



[@awsdevelopers](https://twitter.com/awsdevelopers)



[@awsdevelopers](https://www.instagram.com/awsdevelopers)



[collectives/aws](https://collectives.aws)

Session Videos from AWS re:Invent 2025 (1/3)

1. [Slides 34 & 35: Breaking 25 years of tech debt using AWS Transform for .NET](#) (MAM410)
2. [Slide 30: Amazon S3 Tables architecture, use cases, and best practices](#) (STG334)
3. [Slide 31: Building multi-region data lakes with Replication for Amazon S3 Tables](#) (STG358)
4. [Slide 32: Transforming AI storage economics with Amazon S3 Vectors](#) (STG318)
5. [Slide 39: Simplify backup for stateful Amazon EKS workloads](#) (CNS209)
6. [Slide 46: AWS Graviton: The best price performance for your AWS workloads](#) (CMP307)
7. [Slides 47 & 48: End-to-end foundation model lifecycle on AWS Trainium](#) (AIM351)
8. [Slide 51: Simplify your Kubernetes journey with Amazon EKS Capabilities](#) (CNS378)
9. [Slide 54: Launch web applications in seconds with Amazon ECS Express Mode](#) (CNS379)
10. [Slides 55 & 56: Build AI-powered developer experiences with MCP on ECS and EKS](#) (CNS358)
11. [Slide 62: Amazon Linux: latest features and AWS optimizations](#) (CMP303)
12. [Slide 66: Deep dive into AWS Database Savings Plans](#) (DAT326)
13. [Slides 67 & 68: What's new in search, observability, and vector databases with OpenSearch](#) (ANT201)
14. [Slides 70 & 71: Deep dive into Amazon Aurora DSQL and its architecture](#) (DAT439)
15. [Slide 72: DynamoDB: Resilience & lessons from the Oct 2025 service disruption](#) (DAT453)
16. [Slide 80: Secure Global DNS resolution: Introducing Route 53 Global Resolver](#) (NET217)
17. [Slide 85: Advanced VPC design and new capabilities](#) (NET340)
18. [Slide 87: Load balancing evolved: ALB Target Optimizer](#) (NET336)
19. [Slide 88: Deep dive into advanced routing policy with AWS Cloud WAN](#) (NET401)
20. [Slide 102: Scaling Amazon Redshift with a multi-warehouse architecture](#) (ANT318)

Session Videos from AWS re:Invent 2025 (2/3)

21. [Slide 115: Build multi-modal AI agents with Strands Agents and Amazon S3 Vectors](#) (DEV332)
22. [Slide 114: Unlock Advanced Model Training: Reinforcement Fine-tuning on Bedrock](#) (AIM3327)
23. [Slides 117 & 118: SageMaker HyperPod: Checkpointless & elastic training for AI models](#) (AIM3338)
24. [Slides 119 and 120: Build useful, reliable agents with Amazon Nova](#) (AIM372)
25. [Slide 121: Amazon Nova 2 Omni: A new frontier in multimodal AI](#) (AIM3324)
26. [Slide 122: Create hyper-personalized voice interactions with Amazon Nova Sonic](#) (AIM374)
27. [Slide 123: Amazon Nova Forge: Build your own frontier models using Amazon Nova](#) (AIM3325)
28. [Slide 126: Improve agent quality in production with Bedrock AgentCore Evaluations](#) (AIM3348)
29. [Slide 127: Make agents remember with Amazon Bedrock AgentCore Memory](#) (AIM331)
30. [Slides 143, 146, 149, and 150: A decade of API Gateway: Architecting the future of API Management](#) (API314)
31. [Slide 144: Deep Dive on AWS Lambda durable functions](#) (CNS380)
32. [Slide 145: Lambda Managed Instances: EC2 Power with Serverless Simplicity](#) (CNS382)
33. [Slide 148: Secure Multi-tenant SaaS with AWS Lambda: A Tenant Isolation Deep Dive](#) (CNS381)
34. [Slide 151: Unleash Rust's potential on AWS](#) (DEV307)
35. [Slide 156: AWS Support: New Tiers Making Customer Operations Effortless](#) (SPS325)
36. [Slide 157: Move beyond reactive: Transform cloud ops with AWS DevOps Agent](#) (COP362)
37. [Slides 158, 161, and 174: What's new with Amazon CloudWatch](#) (COP363)
38. [Slide 159: From Code to Policies: Accelerate Development with IAM Policy Autopilot](#) (SEC351)
39. [Slides 164, 169, 170, and 171: AWS infrastructure as code: a year in review](#) (DVT203)
40. [Slide 172: Spec-driven development with Kiro](#) (DEV314)

Session Videos from AWS re:Invent 2025 (3/3)

- 41. [Slide 180: AWS Security Agent: Proactive AppSec from Design to Deployment](#) (SEC348)
- 42. [Slide 183: Programmatically access AWS with your console credentials](#) (SEC354)
- 43. [Slide 191: Agentic AI advancements in customer experience with Amazon Connect](#) (BIZ221)
- 44. Bonus: [Best practices for serverless developers](#) (CNS403)
- 45. Bonus: [Integration patterns for multi-tenant systems](#) (SAS312)
- 46. Bonus: [Serverless full-stack in action: AI-driven developer experience](#) (DEV309)
- 47. Bonus: [Under the hood: Architecting Amazon EKS for scale and performance](#) (CNS429)
- 48. Bonus: [Build modern applications with Amazon Aurora DSQL](#) (DEV308)
- 49. Bonus: [Red Team vs Blue Team: Securing AI Agents](#) (DEV317)
- 50. Bonus: [Accelerate development with Kiro's new agentic AI capabilities](#) (DVT228)
- 51. Bonus: [Integration patterns for distributed systems](#) (API315)
- 52. Bonus: [What's new in fullstack AWS app development](#) (DVT204)
- 53. Bonus: [Compile blazing-fast MCP servers in Rust](#) (DEV405)
- 54. Bonus: [How Amazon Teams Use AI Assistants to Accelerate Development](#) (DEV403)
- 55. Bonus: [Build a multi-agent role-playing Game Master with Strands Agents](#) (DEV330)
- 56. Bonus: [AI Agents for Databases: Discover, Recommend, Optimize](#) (DEV315)
- 57. Bonus: [Optimize AWS Costs: Developer Tools and Techniques](#) (DEV318)
- 58. Bonus: [Building Scalable, Self-Orchestrating AI Workflows with A2A and MCP](#) (DEV415)
- 59. Bonus: [Designing local Generative AI inference with AWS IoT Greengrass](#) (DEV316)
- 60. Bonus: [Modernize Legacy Applications with Agentic AI](#) (DEV333)