

AWS Recent Announcements RSS



Recent Announcements

The AWS Cloud platform expands daily. Learn about announcements, launches, news, innovation and more from Amazon Web Services.



[AWS Firewall Manager is now available in the AWS Canada West \(Calgary\) region](#)

AWS Firewall Manager is now available in the AWS Canada West (Calgary) region, enabling customers to create policies for AWS WAF and manage web application security for applications running in this region. Support for other policy types will be available in the coming months. Firewall Manager is now available in a total of 31 AWS commercial regions, 2 GovCloud regions, and all Amazon CloudFront edge locations.

[AWS IoT TwinMaker announces Knowledge Graph optimization for efficient entity-metadata query capabilities](#)

AWS IoT TwinMaker makes it easier to create digital twins of real-world systems such as buildings, factories, industrial equipment, and production lines. Today, AWS announced enhancements to AWS IoT TwinMaker Knowledge Graph that enable faster and more flexible entity metadata query capabilities for industrial customers. These optimizations provide faster entity-metadata queries using full-text search and wildcard search, addressing the data engineering needs of customers requiring efficient entity metadata access.

[Amazon EMR Serverless announces detailed performance monitoring of Apache Spark jobs with Amazon Managed Service for Prometheus](#)

Amazon EMR Serverless is a serverless option in Amazon EMR that makes it simple for data engineers and data scientists to run open-source big data analytics frameworks without configuring, managing, and scaling clusters or servers. Today, we are excited to announce detailed performance monitoring of Apache Spark jobs with Amazon Managed Service for Prometheus, allowing you to analyze, monitor, and optimize your jobs using job-specific engine metrics and information about Spark event timelines, stages, tasks, and executors.

[Amazon Connect Cases now provides APIs for managing attachments](#)

Amazon Connect Cases now provides APIs that make it easy to upload files, check file details, and delete files from cases. Contact center administrators can use these APIs to automate the attachment of files to cases. In addition, these APIs also enable you to use case attachments in a custom agent desktop.

[Amazon EC2 R7i instances are now available in AWS GovCloud \(US-East\) Region](#)

Starting today, Amazon Elastic Compute Cloud (Amazon EC2) R7i instances powered by custom 4th Gen Intel Xeon Scalable processors (code-named Sapphire Rapids) are available in AWS GovCloud (US-East) region. These custom processors, available only on AWS, offer up to 15% better performance over comparable x86-based Intel processors utilized by other cloud providers.

[AWS Amplify Gen 2 is now generally available](#)

AWS Amplify Gen 2, the code-first developer experience for building full-stack apps using TypeScript, is now generally available. Amplify Gen 2 enables developers to express app requirements like the data models, business logic, and authorization rules in TypeScript. The necessary cloud infrastructure is then automatically provisioned, without needing explicit infrastructure definitions. This streamlined approach accelerates full-stack development for teams of all sizes.

[Amazon EMR Studio is now available in the two additional AWS regions](#)

Starting today, you can use Amazon EMR Studio in the Asia Pacific (Melbourne) and Israel (Tel Aviv) regions to run interactive workloads on EMR.

[Amazon FSx for Windows File Server is now available in the AWS Canada West \(Calgary\) Region](#)

Customers can now create Amazon FSx for Windows File Server file systems in the AWS Canada West (Calgary) Region.

[Amazon RDS for Oracle now supports April 2024 Release Update](#)

Amazon Relational Database Service (Amazon RDS) for Oracle now supports the April 2024 Release Update (RU) for Oracle Database versions 19c and 21c.

[Amazon Connect Contact Lens now provides support for PII redaction in Spanish](#)

Amazon Connect Contact Lens now provides support for personally identifiable information (PII) redaction in the Spanish (US) language, enabling contact centers to help identify and redact sensitive information on contact transcripts such as social security numbers, credit card details, bank account information, and personal contact information (i.e. name, email address, phone number, and mailing address).

[Amazon FSx for OpenZFS is now available in the AWS Canada West \(Calgary\) Region](#)

Customers can now create Amazon FSx for OpenZFS file systems in the AWS Canada West (Calgary) Region.

[Amazon FSx for Lustre is now available in the AWS Canada West \(Calgary\) Region](#)

Customers can now create Amazon FSx for Lustre file systems in the AWS Canada West (Calgary) Region.

[Amazon FSx for NetApp ONTAP is now available in the AWS Canada West \(Calgary\) Region](#)

Customers can now create Amazon FSx for NetApp ONTAP file systems in the AWS Canada West (Calgary) Region.

[Amazon DynamoDB introduces configurable maximum throughput for On-demand tables](#)

Amazon DynamoDB on-demand is a serverless, pay-per-request billing option that can serve thousands of requests per second without capacity planning. Previously, the on-demand request rate was only limited by the default throughput quota (40K read request units and 40K write request units), which uniformly applied to all tables within the account, and could not be customized or tailored for diverse workloads and differing requirements. Since on-demand mode scales instantly to accommodate varying traffic patterns, a piece of hastily written or unoptimized code could rapidly scale up and consume resources, making it difficult to keep costs and usage bounded.

[IP prefix visibility on Amazon CloudWatch Internet Monitor console](#)

Amazon CloudWatch Internet Monitor, an internet traffic monitoring service for AWS applications that gives you a global view of traffic patterns and health events, now displays IPv4 prefixes in its console dashboard. Using this data in Internet Monitor, you can get more details about your application traffic and health events. For example, you can do the following:

- View the IPv4 prefixes associated with a client location that is impacted by a health event
- View IPv4 prefixes associated with a client location
- Filter and search traffic data by the network associated with an IPv4 prefix or IPv4 address

AWS Control Tower is now available in AWS Canada West (Calgary) Region

Starting today, customers can use AWS Control Tower in the AWS Canada West (Calgary) Region. With this launch, AWS Control Tower is available in 29 AWS Regions and the AWS GovCloud (US) Regions. AWS Control Tower offers the easiest way to set up and govern a secure, multi-account AWS environment. It simplifies AWS experiences by orchestrating multiple AWS services on your behalf while maintaining the security and compliance needs of your organization. You can set up a multi-account AWS environment within 30 minutes or less, govern new or existing account configurations, gain visibility into compliance status, and enforce controls at scale.

Amazon RDS for SQL Server Supports SSAS Multidimensional for SQL Server 2019

Amazon RDS for SQL Server now supports SQL Server Analysis Services (SSAS) in Multidimensional mode for SQL Server 2019. There is no additional cost to install SSAS directly on your Amazon RDS for SQL Server DB instance.

Amazon Personalize now makes it easier than ever to delete users from datasets

Amazon Personalize now makes it easier than ever to remove users from your datasets with a new deletion API. Amazon Personalize uses datasets provided by customers to train custom personalization models on their behalf. This new capability allows you to delete records about users from your datasets including user metadata and user interactions. This helps to maintain data for your compliance programs and keep your data current as your user base changes. Once the deletion is complete, Personalize will no longer store information about the deleted users and therefore will not consider the user for model training.

Amazon Pinpoint introduces country rules to precisely control SMS message delivery

Amazon Pinpoint now offers country rules, a new feature that allows developers to control the specific countries they send SMS and voice messages to. This enhancement helps organizations align their message sending activities to the precise list of countries where they operate.

Amazon Route 53 Resolver DNS Firewall now available in the Canada West (Calgary) Region

Starting today, you can use Amazon Route 53 Resolver DNS Firewall in the Canada West (Calgary) Region.

AWS announces a new Amazon EC2 API to retrieve the public endorsement key from NitroTPM

Today, AWS introduces a new EC2 API to retrieve the public endorsement key (EkPub) for the Nitro Trusted Platform Module (NitroTPM) of an Amazon EC2 instance.

AWS Transfer Family is now available in the AWS Canada West (Calgary) Region

Customers in AWS Canada West (Calgary) Region can now use AWS Transfer Family.

Announcing automated 997 and TA1 acknowledgements for AWS B2B Data Interchange

AWS B2B Data Interchange now automatically generates 997 functional acknowledgements and TA1 interchange acknowledgements in response to all relevant inbound X12 electronic data interchange (EDI) transactions. These acknowledgements are used to confirm receipt of individual transactions and to report errors. With this launch, you can now automate delivery of 997 and TA1 acknowledgements to trading partners that require them.

Amazon Connect Contact Lens now provides generative AI-powered agent performance evaluations (preview)

Amazon Connect Contact Lens now provides managers with generative AI-powered recommendations for answers to questions in agent evaluation forms, enabling them to perform evaluations faster and more accurately. Managers now receive additional agent behavioral insights (e.g., did the agent show empathy while delivering bad news?) and will get context and justification for the recommended answers (reference points from the transcript that were used to provide answers).

Amazon Personalize launches new recipes supporting larger item catalogs with lower latency

Today, Amazon Personalize announces the general availability of two new recipes, User-Personalization-v2 and Personalized-Ranking-v2 (v2 recipes). Built on [Transformers](#) architecture, these new recipes support catalogs with up to 5 million items with lower inference latency. Amazon Personalize testing showed that v2 recipes improved recommendation accuracy by up to 9% and recommendation coverage by up to 1.8x compared to previous versions. A higher coverage means Amazon Personalize recommends more of your catalog. These new recipes also support item metadata like genres and descriptions in inference responses, allowing customers to easily enrich recommendations in their user interfaces.

Amazon Chime SDK Voice Connector now supports audio streaming G.711 A-Law encoded audio

Amazon Chime SDK Voice Connector audio streaming now supports G.711 A-law encoded audio. With this update, companies can stream audio to AWS from their phone systems using session initiation protocol recording (SIPREC) with G.711 A-law encoded audio.

Amazon Bedrock now available in the Asia Pacific (Mumbai) Region

Beginning today, customers can use Amazon Bedrock in the Asia Pacific (Mumbai) region to easily build and scale generative AI applications using a variety of foundation models (FMs) as well as powerful tools to build generative AI applications.

Knowledge Bases for Amazon Bedrock now supports MongoDB Atlas for vector storage

Knowledge Bases for Amazon Bedrock securely connects foundation models (FMs) to internal company data sources for Retrieval Augmented Generation (RAG), to deliver more relevant and accurate responses. Today, we are announcing vector storage support for MongoDB Atlas in Knowledge Bases (KB) for Amazon Bedrock.

Introducing file commit history in Amazon CodeCatalyst

Today, AWS announces the general availability of file commit history in Amazon CodeCatalyst. Customers can now view the file git commit history in the CodeCatalyst's console. Amazon CodeCatalyst helps teams plan, code, build, test, and deploy applications on AWS. Viewing the history of commits, pertaining to a file, is a way to reduce the effort for developers when trying to understand the history of changes for a codebase.

AWS Trusted Advisor now supports API to exclude resources

AWS Trusted Advisor introduces new API to enable you to programmatically exclude resources from recommendations associated with Trusted Advisor best practice checks. On Nov 17 2023, Trusted Advisor announced the launch of new APIs available to Business, Enterprise On-Ramp, or Enterprise Support customers. This launch adds the BatchUpdateRecommendationResourceExclusions API to this suite of APIs to allow users to exclude specific resources from specific checks programmatically. This capability is only available to via AWS Trusted Advisor APIs, and not available through AWS Support API (SAPI).

Amazon EC2 simplifies visibility into your active AMIs

Starting today, you can check when your Amazon Machine Images (AMIs) were last used to launch EC2 instances by simply describing your AMIs, enabling you to efficiently filter and track your active AMIs.

Amazon EMR Serverless introduces Shuffle-optimized disks delivering improved performance for I/O intensive workloads

Amazon EMR Serverless is a serverless option in Amazon EMR that makes it simple for data engineers and data scientists to run open-source big data analytics frameworks without configuring, managing, and scaling clusters or servers. An EMR Serverless application uses workers to execute workloads, allowing users to configure ephemeral storage per worker based on the workload's needs. Today, we are excited to introduce Shuffle-optimized disks on Amazon EMR Serverless, offering increased storage capacity (up to 2TB) and higher IOPS delivering better performance for I/O-intensive Spark and Hive workloads.

New feature for Amazon Connect Contact Lens agent screen recording

You can now enable agent screen recording when your VDI environment is configured to allow multiple agents to connect concurrently to the same Windows instance (multi-session VDI). This makes it even easier and more cost effective for you to help agents improve their performance when using Amazon Connect in a multi-session VDI environment.

Amazon EC2 now protects your AMIs from accidental deregistration

Starting today, you can prevent Amazon Machine Images (AMIs) from accidental deregistration by marking them as protected. A protected AMI cannot be deregistered until you explicitly disable deregistration protection.

Amazon EFS increases maximum per-client throughput to 1.5 GiB/s

Amazon EFS file systems now support up to 1.5 GiB/s of throughput per client, a 3x increase over the previous limit of 500 MiB/s. With this launch, you now have a simpler way to run throughput-intensive file workloads on AWS.

Amazon Route 53 Resolver DNS Firewall now supports Domain Redirection

Starting today, you can enable Route 53 Resolver DNS Firewall to automatically skip the inspection of domains included in a domain redirection chain, such as Canonical Name (CNAME) and Delegation Name (DNAME), thus avoiding the need to explicitly specify each domain from the chain in your Route 53 DNS Firewall rules when allow-listing domains.

Amazon Pinpoint now supports multimedia messaging service (MMS)

Today, Amazon Pinpoint announces support for outbound multimedia messaging service (MMS) within the United States and Canada. The addition of MMS allows Amazon Pinpoint users to enhance their customer communication experiences with the ability to combine traditional text-only messaging with rich media content such as images, audio, and video files.

Amazon CloudWatch launches resource filtering for cross-account observability

Amazon CloudWatch is excited to announce a resource filtering capability for cross-account observability, providing customers with the flexibility to share a subset of their logs or metrics across multiple AWS accounts using configurable filters.

Amazon Q Developer is now generally available

Today, AWS announces the general availability of Amazon Q Developer, a generative AI-powered assistant that reimagines your experience across the entire software development lifecycle (SDLC). Amazon Q Developer includes unique, game-changing capabilities that allow developers to offload time-consuming, manual tasks inside or outside of AWS. Amazon Q Developer capabilities include Q&A and diagnosing common errors in the AWS Management Console, Amazon Q data integration which enables you to build data integration pipelines using natural language, conversational coding and inline code generation in the IDE, and Amazon Q Developer Agent for software development in the IDE and in Amazon CodeCatalyst. Amazon Q Developer also includes Amazon Q Developer Agent for code transformation, a feature that can accelerate your application maintenance, upgrades, and migration in minutes. Additionally, Amazon Q Developer lists and describes resources in your AWS account (preview), and can now help you retrieve and analyze cost data from AWS Cost Explorer (preview).

AWS Security Hub announces the AWS Resource Tagging Standard

Today, AWS Security Hub announces the release of the [AWS Resource Tagging standard](#). The standard contains 85 new controls which can be used to identify if any of your AWS Resources are missing tag keys required by your organization. With the release of this standard, Security Hub now offers 386 security controls that automatically check the compliance of your AWS resources against pre-defined security principles and best practices.

Amazon Q data integration in AWS Glue is now generally available

Today, AWS announces general availability of Amazon Q data integration, a new generative AI-powered capability of Amazon Q Developer that enables you to build data integration pipelines using natural language. Amazon Q Developer is the AWS expert to assist you with all of your development tasks. Amazon Q data integration is a new chat experience specifically for [AWS Glue](#), design for authoring and troubleshooting data integration pipelines.

AWS CodeArtifact now supports RubyGems

Today, AWS announces the general availability of RubyGems support in CodeArtifact. Gems, which are used to distribute Ruby libraries, can now be stored in CodeArtifact.

AWS Config simplifies usage analysis with Amazon CloudWatch metrics

Starting today, the Amazon CloudWatch metrics for monitoring AWS Config data usage will display only billable usage. With this enhancement, non-billable usage will no longer be displayed in both the Amazon CloudWatch Config metrics and AWS Config console. This allows you to validate AWS Config setup and usage using Amazon CloudWatch metrics and correlate billable usage with associated costs.

AWS HealthOmics now supports dynamic run storage for private workflows

We are excited to announce [AWS HealthOmics](#) private workflows now support file systems that dynamically scale with your workflow. AWS HealthOmics is a fully managed service that empowers healthcare and life science organizations to store, query, analyze omics data to generate insights to improve health and drive scientific discoveries. With this release, customers can now choose between static and dynamic run storage options.

Automate deployment of SAP Web Dispatcher using AWS Launch Wizard

AWS Launch Wizard now allows you to automate deployment of [SAP Web Dispatcher](#). This launch expands on existing Launch Wizard capabilities that allow you to automate deployment of SAP HANA, SAP NetWeaver based applications on HANA and ASE databases, SAP BW/4HANA, SAP S/4HANA, and SAP S/4HANA foundations using APIs or a console-based approach.

Announcing Timestream Compute Unit (TCU) for Amazon Timestream for LiveAnalytics

Today, Amazon Timestream for LiveAnalytics announces Timestream Compute Unit (TCU), serverless compute capacity, for customers to predict and control query costs. With TCUs, you are charged for the duration of compute units (TCU) used by your queries, and there are no minimum bytes metered for queries.

Amazon EventBridge Pipes now supports event delivery through AWS PrivateLink

Amazon EventBridge Pipes now supports event delivery through [AWS PrivateLink](#), allowing you to send events from an event source located in an Amazon Virtual Private Cloud (VPC) to a Pipes target without traversing the public internet. With today's launch, you can use Pipes to poll from Amazon Managed Streaming for Apache Kafka (MSK), self-managed Kafka, and Amazon MQ sources residing in a private subnet without the need to deploy a NAT gateway, configure firewall rules, or set up proxy servers.

Announcing the general availability of Amazon Q Business and Amazon Q Apps (Preview)

Today, AWS announces the general availability of Amazon Q Business and the preview of Amazon Q Apps, a new Amazon Q Business capability. Amazon Q Business revolutionizes the way that employees interact with organizational knowledge and enterprise systems. It helps users get comprehensive answers to complex questions and take actions in a unified, intuitive web-based chat experience—all using an enterprise's existing content, data, and systems. Amazon Q Business connects seamlessly to over 40 popular enterprise systems, including Amazon Simple Storage Service (Amazon S3), Microsoft 365, and Salesforce. It ensures that users access content securely with their existing credentials using single sign-on, according to their permissions, and enterprise-level access controls.

Amazon OpenSearch Service now supports Route 53 alias record for custom endpoint

Amazon OpenSearch Service now supports Amazon Route 53 alias records for defining custom domain endpoints. Alias records provide better flexibility when configuring routing to AWS resources. For more information about Route 53 alias records, please see [documentation](#).

AWS announces Amazon Q in QuickSight

Amazon Q in QuickSight is now generally available. The Generative BI capabilities of Amazon Q in QuickSight help business analysts and business users easily build and consume insights using natural language.

AWS HealthOmics now supports cross-account sharing for private workflows

We are excited to announce [AWS HealthOmics](#) now supports the ability to share HealthOmics workflows across AWS accounts. AWS HealthOmics is a fully managed service that empowers healthcare and life science organizations to store, query, analyze omics data, and generate insights to improve health and drive scientific discoveries. With this release, customers can now develop analyses within a single AWS account and share them with other AWS accounts in the same organization or across organizations.

AWS WAF is now available in the Canada West (Calgary) Region

Starting today, AWS WAF is available in the AWS Canada West (Calgary) Region. This is the second Region in Canada where AWS WAF is available, joining the AWS Canada (Central) Region, and giving customers more choice and flexibility.

[AWS AppFabric now supports Salesforce, Azure Monitor and Google Analytics](#)

Today, [AWS AppFabric](#) announces support for three new data sources: Salesforce, Azure Monitor, and Google Analytics. Starting now, IT administrators and security analysts can use AppFabric to quickly integrate with [29 supported SaaS applications](#), aggregate enriched and normalized SaaS audit logs, and audit end-user access across their SaaS apps.

[AWS announces new AWS Direct Connect location in Aurora, Illinois](#)

Today, AWS announced the opening of a new [AWS Direct Connect](#) location within the CyrusOne Aurora data center in Aurora, Illinois. By connecting your network to AWS at the new Illinois location, you gain private, direct access to all public AWS Regions (except those in China), AWS GovCloud Regions, and [AWS Local Zones](#). This is the third AWS Direct Connect site within Chicago Metropolitan area and the 43rd site in the United States.

[Amazon DynamoDB now supports an AWS FIS action to pause global table replication](#)

[Amazon DynamoDB](#) now supports an [AWS Fault Injection Simulator \(FIS\)](#) action to pause replication for global tables. FIS is a fully managed service for running controlled fault injection experiments to improve an application's performance, observability, and resilience. Global tables replicate your Amazon DynamoDB tables automatically across your choice of AWS Regions to achieve fast, local read and write performance. Customers can use the new FIS action to observe how their application responds to a pause in regional replication, and tune their monitoring and recovery process to improve resiliency and application availability.

[Amazon Location Service releases Places integration plugin for MapLibre GL JS Geocoder](#)

[Amazon Location Service](#) has released a Places plugin for MapLibre GL JS Geocoder, simplifying the integration of Amazon Location Places Service with MapLibre. Along with the authentication libraries, this plugin provides a frictionless experience to use Amazon Location Places with MapLibre GL JS.

[AWS Mainframe Modernization File Transfer introduces enhanced data set discovery workflow](#)

We are excited to announce public availability of an enhanced dataset discovery workflow for AWS Mainframe Modernization File Transfer with BMC. This new capability makes it easier for users to discover and select mainframe data sets for conversion and transfer to AWS. Users can now browse the mainframe data set catalog directly from the AWS Mainframe Modernization service console, select data sets from one or multiple mainframe volumes for transfer, view extended data set metadata, and choose data set codepage conversions. The enhanced workflow provides an intuitive and unified mechanism for mainframe data access and transfer.

[Amazon Titan Text Embeddings V2 now available in Amazon Bedrock](#)

[Amazon Titan Text Embeddings V2](#), a new embeddings model in the [Amazon Titan family of models](#), is now generally available in [Amazon Bedrock](#). Using Titan Text Embeddings V2, customers can perform various natural language processing (NLP) tasks by representing text data as numerical vectors, known as embeddings. These embeddings capture the semantic and contextual relationships between words, phrases, or documents in a high-dimensional vector space. This model is optimized for Retrieval-Augmented Generations (RAG) use cases and is also well suited for a variety of other tasks such as information retrieval, question and answer chatbots, classification, and personalized recommendations.

[Amazon Redshift announces support for Multi-AZ deployment with zero-ETL integration](#)

[Amazon Redshift](#) now supports Multi-AZ deployment for zero-ETL integration on RA3 clusters, enabling customers to run near real-time analytics on a highly available data warehouse. With a Multi-AZ deployment, your zero-ETL integration can automatically recover from any infrastructure or Availability Zone (AZ) failures ensuring your workloads remain uninterrupted.

[Amazon Q launches subscription management with AWS IAM Identity Center integration](#)

Today, Amazon Q launched a subscription management service enabling customers to manage subscriptions for Amazon Q plans like Amazon Q Business Pro, Amazon Q Business Lite, and Amazon Q Developer Pro. The new subscription management service offers administrators access to dashboards that provide subscription details, including the specific users and groups assigned to each subscription. This centralized visibility enables tracking Amazon Q subscriptions across the entire organization.

[Amazon Transcribe announces general availability of generative AI-powered call summarization](#)

Today, Amazon Transcribe announces the general availability of generative AI-powered call summarization available through the Amazon Transcribe Call Analytics API. Generative call summarization delivers a concise summary of contact center interactions, capturing key components such as why the customer called, how the issue was addressed, and what follow-up actions were identified.

[Amazon RDS for MySQL supports Innovation Release version 8.3 in Amazon RDS Database Preview Environment](#)

[Amazon RDS for MySQL](#) now supports MySQL Innovation Release 8.3 in the [Amazon RDS Database Preview Environment](#), allowing you to evaluate the latest Innovation Release on Amazon RDS for MySQL. You can deploy MySQL 8.3 in the Amazon RDS Database Preview Environment that has the benefits of a fully managed database, making it simpler to set up, operate, and monitor databases.

[AWS Firewall Manager now supports central deployment and management of VPC NACLs with common NACL policies](#)

Starting today, [AWS Firewall Manager](#) allows customers to define policies for centrally creating, deploying and managing Amazon Virtual Private Cloud Network Access Control List (NACL) rules for their VPC subnets across accounts in their AWS Organization. Additionally, NACL policies provide customers with visibility into policy adherence by reporting the compliance status for each policy.

[Cohere Command R and Command R+ now available in Amazon Bedrock](#)

You can now access Cohere's newest state-of-the-art enterprise foundation model family, Command R+ and Command R, in [Amazon Bedrock](#). These generative AI models are highly scalable, optimized for long context tasks like advanced retrieval-augmented generation (RAG) with citations to mitigate hallucinations, multi-step tool use to automate complex business tasks, and are multilingual in 10 languages to support global business operations.

[Amazon Aurora supports PostgreSQL 16.2, 15.6, 14.11, 13.14, and 12.18](#)

[Amazon Aurora PostgreSQL-Compatible Edition](#) now supports PostgreSQL versions 16.2, 15.6, 14.11, 13.14, and 12.18. These releases contain product improvements and bug fixes made by the PostgreSQL community, along with Aurora-specific improvements. These releases also offer improved performance through faster COPY operations (for PostgreSQL 16.2), [Query Plan Management](#) enhancements for queries with aggregate operations, and optimizations for further reducing the logical replication lag.

[Announcing the general availability of AWS Local Zones in Honolulu, Hawaii](#)

Today, AWS announces the general availability of AWS Local Zones in Honolulu, Hawaii. You can now use Local Zones in Honolulu to deliver applications that require single-digit millisecond latency or local data processing. In this new Local Zone, you can access Amazon Elastic Block Store (Amazon EBS), Amazon Elastic Compute Cloud (Amazon EC2), Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), Application Load Balancer, and AWS Direct Connect to support a broad set of workloads at the edge.

[Amazon Q generative SQL \(preview\) is now available in AWS Europe \(Frankfurt\) region](#)

[Amazon Q generative SQL](#) in Amazon Redshift Query Editor is available for public preview in AWS Europe (Frankfurt) region. Amazon Q generative SQL is available in Amazon Redshift Query Editor, an out-of-the-box web-based SQL editor for Redshift, to simplify query authoring and increase your productivity by allowing you to express queries in natural language and receive SQL code recommendations. Furthermore, it allows you to get insights faster without extensive knowledge of your organization's complex database metadata.

[Amazon CloudWatch Lambda Insights available in the AWS GovCloud \(US\) Regions](#)

[Amazon CloudWatch Lambda Insights](#) is now available in AWS GovCloud (US) Regions. Lambda Insights enables you to monitor, troubleshoot, and optimize the performance of AWS Lambda functions. Lambda Insights provides you access to automated dashboards summarizing the performance and health of your Lambda functions that provide visibility into issues such as memory leaks or performance changes caused by new function versions.

[Network Load Balancer now supports Resource Map in AWS Management Console](#)

[Network Load Balancer \(NLB\)](#) now supports Resource Map, a tool in the console that displays all your NLB resources and their relationships in a visual format on a single page, providing you a clear understanding of your NLB architecture.

Amazon Data Lifecycle Manager default policies now supports AWS Organizations

Customers can now create and manage default policies across their entire organization or organizational unit (OU) with AWS CloudFormation StackSets. Default policies work in conjunction with customers' existing backup mechanisms to only create EBS-backed AMIs and EBS Snapshots of instances and volumes without recent backups. This helps administrators ensure that all member accounts have comprehensive backup protection without creating duplicate backups or increasing management overhead and cost.

Restore running applications to pre-update state in Amazon Managed Service for Apache Flink

You can now restore your Amazon Managed Service for Apache Flink application to the previous running version and application state from the most recent, successful snapshot. This feature will work when your application is running and is most useful when you want to immediately rollback to the previous application version to mitigate downstream impact of an application update. Prior to this launch, you could only rollback applications that were in updating or autoscaling statuses.

AWS CodePipeline supports stage level manual and automated rollback

AWS CodePipeline V2 type pipelines now support stage level rollback to help customers to confidently deploy changes to their production environment. When a pipeline execution fails in a stage due to any action(s) failing, customers can quickly get that stage to a known good state by rolling back to a previously successful pipeline execution in that stage. Customers can roll back changes in any stage, whether succeeded or failed, except the Source stage.

Local time zone support for Amazon RDS for Db2

[Amazon Relational Database Service \(Amazon RDS\) for Db2](#) now supports local time zones. You can now set the time zone for your Amazon RDS for Db2 instances to the local time zone of your choice.

Amazon RDS now supports M6gd database instances in four additional AWS regions

[Amazon Relational Database Service \(Amazon RDS\)](#) for PostgreSQL, MySQL, and MariaDB now supports AWS Graviton2-based M6gd database instances in Asia Pacific (Hyderabad), Europe (Spain, Zurich), and Middle East (UAE) regions.

AWS supports dynamically removing and adding auto assigned public IPv4 address

Amazon VPC announces a network interface setting to dynamically remove and add an [auto assigned public IPv4 address](#) on EC2 instances. With this capability, customers that no longer require an auto assigned public IPv4 address on their EC2 instance can remove the public IPv4 address, and if needed attach back a new public IPv4 address, by modifying the public IP setting on the network interface. Before today, once a public IPv4 address was auto assigned to EC2 instance it was not possible to remove it. It remained on the network interface for the lifetime of the EC2 instance.

Amazon SageMaker Clarify now supports foundation model evaluations

Foundation model evaluations with SageMaker Clarify is now generally available. This capability helps data scientists and machine learning engineers evaluate, compare, and select foundation models based on a variety of criteria across different tasks within minutes.

AWS AppFabric now supports SentinelOne Singularity Cloud

Today, AWS AppFabric announces support for SentinelOne Singularity Cloud as data source and compatible security destination. Starting now, IT administrators and security analysts can use AppFabric to quickly integrate with 27 supported SaaS applications, aggregate enriched and normalized SaaS audit logs, and audit end-user access across their SaaS apps.

AWS Global Accelerator now supports endpoints in the AWS Canada West (Calgary) Region

Starting today, [AWS Global Accelerator](#) supports application endpoints in the AWS Canada West (Calgary) Region, expanding the number of [supported AWS Regions](#) to twenty-nine.

NoSQL Workbench for Amazon DynamoDB launches a revamped operation builder user interface

NoSQL Workbench for Amazon DynamoDB is a client-side application that provides data modeling and query development features to help you design, create, and work with DynamoDB tables. You can use the NoSQL Workbench operation builder to build and save operations for viewing, updating, and exploring your live datasets. You can even generate code in multiple programming languages. Today, we are pleased to announce quality enhancements to the operation builder to help customers better navigate, run operations, and browse their DynamoDB tables.

AWS Direct Connect adds 25 Gbps hosted connection capacities

Today, AWS Direct Connect now enables AWS Direct Connect Service Delivery Partners (AWS Direct Connect Partners) to support 25 Gbps hosted connection capacity.

AWS DataSync supports disabling task schedules

[AWS DataSync](#) now lets you enable and disable task schedules. Using this new feature, you can temporarily disable scheduled executions of your task to accommodate events such as maintenance on your storage systems. Once the event is complete, you can enable your task schedule to resume execution of your task at the next scheduled interval.

AWS CodeBuild now supports managed GitHub Action runners

AWS CodeBuild now supports managed GitHub Action self-hosted runners. Customers can configure their CodeBuild projects to receive GitHub Actions workflow job events and run them on CodeBuild ephemeral hosts. AWS CodeBuild is a fully managed continuous integration service that compiles source code, runs tests, and produces software packages ready for deployment.

Amazon GameLift now includes containers support (Preview)

We are excited to announce that Amazon GameLift now supports containers for building, deploying, and running game server packages. [Amazon GameLift](#) is a fully managed service that allows developers to quickly manage and scale dedicated game servers for multiplayer games. With this preview release, GameLift now supports end-to-end development of containerized workloads, including deployment and scaling on-premises, in the cloud, or hybrid configurations.

Amazon Bedrock Agents simplifies agent creation and launches Return of Control capability

Agents for Amazon Bedrock enable generative AI applications to automate multi-step tasks across company systems and data sources. Agents removes the undifferentiated lifting of orchestration, infrastructure hosting and management, and we're making building Agents easier than ever.

Amazon Titan Image Generator model in Amazon Bedrock now generally available

[Amazon Titan Image Generator](#) enables content creators with rapid ideation and iteration resulting in high efficiency image generation. The Amazon Titan Image Generator model is now generally available in [Amazon Bedrock](#), helping you easily build and scale generative AI applications with new image generation and image editing capabilities.

Watermark detection for Amazon Titan Image Generator now available in Amazon Bedrock

Amazon Titan Image Generator's new watermark detection feature is now generally available in [Amazon Bedrock](#). All Amazon Titan-generated images contain an invisible watermark, by default. The watermark detection mechanism allows you to identify images generated by [Amazon Titan Image Generator](#), a foundation model that allows users to create realistic, studio-quality images in large volumes and at low cost, using natural language prompts.

Model evaluation on Amazon Bedrock is now Generally Available

Model Evaluation on Amazon Bedrock allows you to evaluate, compare, and select the best foundation models for your use case. Amazon Bedrock offers a choice of automatic evaluation and human evaluation. You can use automatic evaluation with predefined algorithms for metrics such as accuracy, robustness, and toxicity. Additionally, for those metrics or subjective and custom metrics, such as friendliness, style, and alignment to brand voice, you can set up a human evaluation workflow with a few clicks. Human evaluation workflows can leverage your own employees or an AWS-managed team as reviewers. Model evaluation provides built-in curated datasets or you can bring your own datasets.

AWS AppFabric now supports 1Password

Today, AWS AppFabric announces support for 1Password. Starting now, IT administrators and security analysts can use AppFabric to quickly integrate with [26 supported SaaS applications](#), aggregate enriched and normalized SaaS audit logs, and audit end-user access across their SaaS apps.

AWS Transfer Family now lists files from remote SFTP servers using SFTP connectors

AWS Transfer Family customers can now use [SFTP connectors](#) to list files stored in remote SFTP servers, enabling visibility into the contents of directories in remote SFTP file systems and transfer files when file names are not known in advance.

Guardrails for Amazon Bedrock is generally available with new safety & privacy controls

Today, we are announcing the general availability of Guardrails for Amazon Bedrock that enables customers to implement safeguards across large language models (LLMs) based on their use cases and responsible AI policies. Customers can create multiple guardrails tailored to different use cases and apply them on multiple LLMs, providing a consistent user experience and standardizing safety controls across generative AI applications.

Amazon ECS now integrates with CloudWatch alarms for deployment monitoring in the AWS GovCloud (US) Regions

[Amazon Elastic Container Service \(Amazon ECS\)](#) now lets you add automated safeguards for rolling updates of Amazon ECS services in the AWS GovCloud (US) Regions. You can now monitor and automatically react to changes during an Amazon ECS rolling update by using Amazon CloudWatch alarms. This allows you to more easily automate discovery and remediation for failed deployments and minimize the impact of a bad change.

Amazon RDS Performance Insights provides execution plan for RDS SQL Server

[Amazon RDS \(Relational Database Service\) Performance Insights](#) now collects the query execution plans of the resource-intensive SQL queries in Amazon RDS for SQL Server, and stores them over time. It helps you identify if a change in the query execution plan is the cause of performance degradation or stalled query.

Amazon QuickSight launches cross sheet filters and controls

Amazon QuickSight now includes cross-sheet filters and controls. This enables authors to create and manage filters and controls across an entire analysis or dashboard.

Knowledge Bases for Amazon Bedrock now simplifies asking questions on a single document

Knowledge Bases for Amazon Bedrock allows you to connect foundation models (FMs) to internal company data sources to deliver more relevant, context-specific, and accurate responses. Knowledge Bases (KB) now provides a real-time, zero-setup, and low-cost method to securely chat with single documents.

Knowledge Bases for Amazon Bedrock now supports multiple data sources

Knowledge Bases for Amazon Bedrock is a fully managed Retrieval-Augmented Generation (RAG) capability that allows you to connect foundation models (FMs) to internal company data sources to deliver more relevant and accurate responses. Knowledge Bases now supports adding multiple data sources, across accounts.

Custom Model Import for Amazon Bedrock

We are excited to announce the preview of Custom Model Import for Amazon Bedrock. Now you can import customized models into Amazon Bedrock to accelerate your generative AI application development. This new feature allows you to leverage your prior model customization investments within Amazon Bedrock and consume them in the same fully-managed manner as Bedrock's existing models. For supported architectures such as Llama, Mistral, or Flan T5, you can now import models customized anywhere and access them on-demand.

Amazon RDS for Oracle supports Kerberos authentication in additional regions

Amazon RDS for Oracle now supports external authentication of database users using Kerberos and Microsoft Active Directory in additional regions. This feature provides the benefits of single sign-on and centralized authentication of Oracle Database users. Keeping all of your user credentials in the same Active Directory will save you time and effort as you will now have a centralized place for storing and managing them for multiple DB instances.

AWS IoT Core for LoRaWAN public network support now available in Spain

Today, AWS IoT Core for LoRaWAN announced the expansion of the public network support in the Spain region. With this expansion, Internet of Things (IoT) customers that offer LoRaWAN-based systems and solutions in Spain can seamlessly connect their LoRaWAN-powered devices to AWS over a public network infrastructure. This public infrastructure is provided as a service and supported by Everynet - a global LoRaWAN network operator offering networks in the United States, United Kingdom, and Spain. Thanks to the publicly available network infrastructure, customers in Spain can realize improved savings in time and costs associated with managing a private network infrastructure for LoRaWAN-based solutions.

Agents for Amazon Bedrock add support for Anthropic Claude 3 Haiku and Sonnet

Agents for Amazon Bedrock enable developers to create generative AI-based applications that can complete complex tasks for a wide range of use cases and deliver answers based on company knowledge sources. In order to complete complex tasks, with high accuracy, reasoning capabilities of the underlying foundational model (FM) play a critical role.

Meta Llama 3 now available in Amazon Bedrock

You can now access Meta's Llama 3 models, Llama 3 8B and Llama 3 70B, in [Amazon Bedrock](#). Meta Llama 3 is designed for you to build, experiment, and responsibly scale your generative artificial intelligence applications. You can now use these two new Llama 3 models in Amazon Bedrock enabling you to easily experiment with and evaluate even more top foundation models for your use case.