

AWS re:Invent

NOV. 30 – DEC. 18, 2020



Data Analytics & AI/ML Announcements

4th January 2021

Intellify

aws partner
network

Advanced
Consulting
Partner

Machine Learning
Consulting Partner
AWS Public Sector
Partner



[Home](#)

Table of Contents

Data Analytics Services	6
AWS Glue DataBrew.....	7
AWS Glue Elastic Views	8
AWS Lake Formation New Features	9
Amazon Redshift Data Sharing.....	10
Amazon AQUA for Amazon Redshift	11
Amazon Redshift federated querying MySQL	12
Amazon Redshift console partner integration	13
Amazon Redshift native JSON and semi-structured data support.....	14
Amazon Managed Workflows for Apache Airflow (MWAA)	15
Amazon QuickSight Q.....	16
Amazon EMR Studio.....	17
Amazon EMR on Amazon EKS	18
Machine Learning Services	19
Amazon SageMaker Data Wrangler	20
Amazon SageMaker Feature Store	21
Amazon SageMaker Clarify.....	22
Amazon SageMaker JumpStart	23
Amazon SageMaker Distributed training.....	24
Amazon SageMaker Debugger and Profiling.....	25
Amazon SageMaker Model Monitor	26
Amazon SageMaker Edge Manager.....	27
Amazon SageMaker Pipelines.....	28
Amazon Neptune ML.....	29
Amazon Redshift ML	30

[Home](#)

- Amazon Athena ML..... 31**
- Amazon Aurora ML 32**
- AWS Trainium 33**
- AI Services 34***
 - Amazon Monitron 35**
 - Amazon Lookout for Equipment 36**
 - Amazon Lookout for Vision..... 37**
 - Amazon Lookout for Metrics 38**
 - AWS Panorama Appliance 39**
 - Amazon DevOps Guru 40**
 - Amazon Kendra Incremental Learning 41**
 - Amazon HealthLake 42**
 - Amazon Forecast Weather Index 43**
- Sydney Region Availability 44***

[Home](#)

Acknowledgment that you can't fight gravity.

If you step back and have conviction that something is going to change because it's a better experience for the customers, it is going to change whether you want it or not whether it's convenient for you or not, it is going to change.

Andy Jassy

re:Invent 2020 Keynote

[Home](#)

Important Note:

re:Invent 2020 breakout session recordings have not been uploaded to the AWS YouTube channel at time of writing this document.

To watch the breakout session recordings listed in each of the below announcements, you must have a registered account on the re:Invent 2020 portal. If you have not registered previously, you can do so using [this link](#).

Since we cannot provide a direct link to the videos, please use the below guide to find the breakout session recording mentioned throughout this document.

1. Login to the [AWS re:Invent 2020 portal](#) using the account you have created previously
2. In the search box on top right corner, copy and paste the name of the session listed in the announcement pages below
3. Watch the recording

[Home](#)

Data Analytics Services

[Home](#)

AWS Glue DataBrew

AWS Glue



AWS Glue DataBrew is a new visual data preparation tool that makes it easy for data analysts and data scientists to clean and normalize data to prepare it for analytics and machine learning. You can choose from over 250 pre-built transformations to automate data preparation tasks, all without the need to write any code. You can automate filtering anomalies, converting data to standard formats, and correcting invalid values, and other tasks. After your data is ready, you can immediately use it for analytics and machine learning projects. You only pay for what you use - no upfront commitment.

Availability Is available today in US East (N. Virginia), US East (Ohio), US West (Oregon), Europe (Ireland), Europe (Frankfurt), Asia Pacific (Tokyo), Asia Pacific (Sydney).

Links [Blog post](#)
[Landing page](#)

Sessions [Brief demo](#)
[Introduction Video](#)
re:Invent portal video:
[NEW LAUNCH!] Zero-code data preparation with AWS Glue DataBrew

Tags AI, Feature, Glue, ML

[Home](#)

AWS Glue Elastic Views

AWS Glue



AWS Glue Elastic Views is a new capability of AWS Glue that makes it easy to build materialized views that combine and replicate data across multiple data stores without you having to write custom code. With AWS Glue Elastic Views, you can use familiar Structured Query Language (SQL) to quickly create a virtual table—a materialized view—from multiple different source data stores. AWS Glue Elastic Views copies data from each source data store and creates a replica in a target data store. AWS Glue Elastic Views continuously monitors for changes to data in your source data stores and provides updates to the materialized views in your target data stores automatically, ensuring data accessed through the materialized view is always up to date.

Availability Is preview today in US East (N. Virginia), US East (Ohio), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland). Please [sign up here](#)

Links [Announcement](#)
[Landing page](#)

Sessions [Introduction video](#)
re:Invent portal video:
[NEW LAUNCH!] Deep dive on AWS Glue Elastic Views

Tags Materialized view, Elastic Views, Glue

[Home](#)

AWS Lake Formation New Features

AWS Lake Formation



AWS Lake formation new features - transactions, row-level security, and acceleration are now available for preview. These capabilities are available via new, open, and public update and access APIs for data lakes. These APIs extend AWS Lake Formation's governance capabilities with row-level security. In addition, this preview introduced governed tables - a new Amazon S3 table type that supports atomic, consistent, isolated, and durable (ACID) transactions. AWS Lake Formation transactions simplify ETL script and workflow development, and allow multiple users to concurrently and reliably insert, delete, and modify rows across multiple governed tables. AWS Lake Formation automatically compacts and optimizes storage of governed tables in the background to improve query performance.

Availability Is preview today in US East (N. Virginia). Please [Sign up here](#)

Links [Announcement](#)
[Landing page](#)

Sessions re:Invent portal video:
[NEW LAUNCH!] Data lakes: Easily build, secure, and share with AWS Lake Formation

Tags ACID, Security, Lake formation, Governance

[Home](#)

Amazon Redshift Data Sharing

Amazon Redshift



Amazon Redshift Data Sharing can securely and easily share live data across Amazon Redshift clusters for read purposes. Data sharing improves the agility of your organization by giving you instant, granular, and high-performance access to data across Amazon Redshift clusters without you needing to manually copy or move it. With data sharing, you have live access to data so that your users can see the most up-to-date and consistent information as it's updated in Amazon Redshift clusters.

Availability Is preview today on Amazon Redshift RA3 node types in the following regions in US East (Ohio) US East (N. Virginia) US West (N. California) US West (Oregon) Asia Pacific (Seoul) Asia Pacific (Sydney) Asia Pacific (Tokyo) Europe (Frankfurt) Europe (Ireland).

Links [Blog post](#)
[Landing page](#)

Videos [Amazon Redshift Data Sharing workflow](#)
[Amazon RedShift Use Cases](#)
re:Invent 2020 portal video:
What's new with Amazon Redshift

Tags Redshift, Data Sharing

[Home](#)

Amazon AQUA for Amazon Redshift

Amazon Redshift



AQUA (Advanced Query Accelerator) for Amazon Redshift is available in preview. AQUA provides a new distributed and hardware accelerated cache that brings compute to the storage layer for Amazon Redshift and delivers up to 10x faster query performance than other cloud data warehouses.

AQUA is a high-speed cache on top of Redshift Managed Storage that can scale out and process data in parallel across many AQUA nodes. AQUA uses AWS designed analytics processors that dramatically accelerate data compression, encryption, and data processing on queries that scan, filter, and aggregate large data sets.

Availability AQUA is available for preview in US East (Ohio), US East (N. Virginia), and US West (Oregon) regions.

Links [Announcement](#)

Sessions AWS re:Invent 2019:
[\[NEW LAUNCH!\] Amazon Redshift reimaged: RA3 and AQUA](#)
re:Invent 2020 portal video:
What's new with Amazon Redshift

Tags Redshift, Cache, Accelerator

[Home](#)

Amazon Redshift federated querying MySQL

Amazon Redshift



Redshift federated query capability has been extended to Amazon RDS for MySQL and Amazon Aurora MySQL (currently in preview).

Amazon Redshift federated query allows you to incorporate live data from the transactional databases as part of your business intelligence (BI) and reporting applications to enable operational analytics. The intelligent optimizer in Amazon Redshift pushes down and distributes a portion of the computation directly into the remote operational databases to speed up performance by reducing data moved over the network.

Amazon Redshift complements subsequent execution of the query by leveraging its massively parallel processing capabilities for further speed up. Federated query also makes it easy to ingest data into Amazon Redshift by letting you query operational databases directly, applying transformations on the fly, and loading data into the target tables without requiring complex ETL pipelines.

Availability Available to all Amazon Redshift customers for preview

Links [Announcement](#)
[Developer guide](#)

Sessions [Amazon Redshift Federated Query](#)
re:Invent 2020 portal video:
What's new with Amazon Redshift

Tags Redshift, Federated query, MySQL

[Home](#)

Amazon Redshift console partner integration

Amazon Redshift



Amazon Redshift now supports native integration with select AWS partners from within the Amazon Redshift Console (in preview).

With these solutions, you can bring data from applications like Salesforce, Google Analytics, Facebook Ads, Slack, Jira, Splunk, and Marketo into your Amazon Redshift data warehouse in an efficient and streamlined way. It also enables you to join these disparate datasets and analyze them together to produce actionable insights.

Availability Available to all Amazon Redshift customers for preview

Links [Announcement](#)
[Developer guide](#)

Sessions re:Invent portal video:
What's new with Amazon Redshift

Tags Redshift, data integration

[Home](#)

Amazon Redshift native JSON and semi-structured data support

Amazon Redshift



Amazon Redshift announced preview of native support for JSON and semi-structured data. It is based on the new data type 'SUPER' that allows you to store the semi-structured data in Redshift tables.

Redshift also adds support for the PartiQL query language to seamlessly query and process the semi-structured data.

This functionality enables you to achieve advanced analytics that combine the classic structured SQL data (such as strings, numerics, and timestamps) with the semi-structured SUPER data with superior performance, flexibility, and ease-of-use.

Availability Available as public preview in SQL_PREVIEW track

Links [Announcement](#)
[Developer guide](#)

Sessions re:Invent portal video:
What's new with Amazon Redshift

Tags Redshift, Semi-structured data, PartiQL, schemaless

[Home](#)

Amazon Managed Workflows for Apache Airflow (MWAA)

Amazon Managed Workflows



Amazon Managed Workflows for Apache Airflow (MWAA) is a fully managed service that makes it easy to run open-source versions of Apache Airflow on AWS, and to build workflows to execute your extract-transform-load (ETL) jobs and data pipelines.

Apache Airflow is an open-source tool used to programmatically author, schedule, and monitor sequences of processes and tasks referred to as “workflows.” With Managed Workflows, you can use Airflow and Python to create workflows without having to manage the underlying infrastructure for scalability, availability, and security.

Availability US East (Northern Virginia), US West (Oregon), US East (Ohio), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Europe (Ireland), Europe (Frankfurt), and Europe (Stockholm).

Links [Blog post](#)
[Landing page](#)

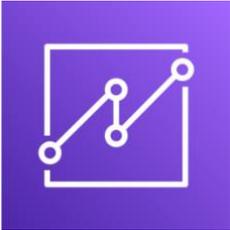
Sessions [Getting started video](#)
re:Invent portal video:
[NEW LAUNCH!] Data pipelines with Amazon Managed Workflows for Apache Airflow

Tags Apache Airflow, Managed service, Workflow management

[Home](#)

Amazon QuickSight Q

Amazon QuickSight



Amazon QuickSight Q is a machine learning powered capability that uses natural language processing to answer your business questions instantly, saving weeks of effort from BI teams having to build pre-defined data models and dashboards.

Availability Q is available in preview for US East (N. Virginia), US West (Oregon), US East (Ohio) and Europe (Ireland).

Links [Blog post](#)
[Landing page](#)

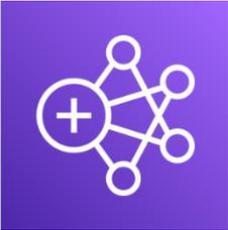
Sessions [Understand QuickSight Q in 5 minutes](#)
re:Invent portal video:
[NEW LAUNCH!] Introducing Amazon QuickSight Q: Ask questions on data & get answers in seconds

Tags AI, NLP, ML, QuickSight

[Home](#)

Amazon EMR Studio

Amazon EMR



Amazon EMR Studio is an integrated development environment (IDE) that makes it easy for data scientists and data engineers to develop, visualize, and debug data engineering and data science applications written in R, Python, Scala, and PySpark.

EMR Studio provides fully managed Jupyter Notebooks, and tools like Spark UI and YARN Timeline Service to simplify debugging. EMR Studio uses AWS Single Sign-On and allows you to log in directly with your corporate credentials without logging into the AWS console. Data scientists and analysts can install custom kernels and libraries, collaborate with peers using code repositories such as GitHub and Bitbucket, or execute parameterized notebooks as part of scheduled workflows using orchestration services like Apache Airflow or Amazon Managed Workflows for Apache Airflow.

Availability In preview today in US East (N. Virginia), US East (Ohio), US West (Oregon), Europe (Ireland)

Links [Blog post](#)
[Landing page](#)

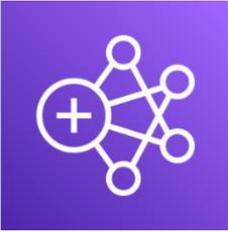
Sessions re:Invent portal video:
What's new with Amazon EMR

Tags EMR, Notebook, AWS Big Data, Analytics

[Home](#)

Amazon EMR on Amazon EKS

Amazon EMR / Amazon EKS



Amazon EMR on Amazon EKS provides a new deployment option for Amazon EMR that allows you to run Apache Spark on Amazon Elastic Kubernetes Service (Amazon EKS). If you already use Amazon EMR, you can now run Amazon EMR based applications with other types of applications on the same Amazon EKS cluster to improve resource utilization and simplify infrastructure management across multiple AWS Availability Zones.

If you already run big data frameworks on Amazon EKS, you can now use Amazon EMR to automate provisioning and management, and run Apache Spark up to 3x faster. Amazon EMR on Amazon EKS looks after building, configuring, and managing containers.

Availability US West (Oregon), US East (N Virginia), and Europe (Ireland)

Links [Announcement](#)
[Landing page](#)

Sessions re:Invent portal video:
What's new with Amazon EMR

Tags EMR, EKS, Apache Spark, Kubernetes

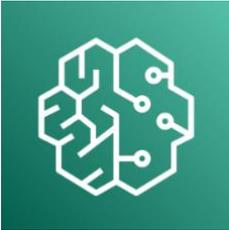
[Home](#)

Machine Learning Services

[Home](#)

Amazon SageMaker Data Wrangler

Amazon SageMaker



Using Amazon SageMaker Data Wrangler's data selection tool, you can choose the data you want from various data sources, including Amazon S3, Amazon Athena, Amazon Redshift, AWS Lake Formation, and Amazon SageMaker Feature Store, and import it with a single click. Amazon SageMaker Data Wrangler contains over 300 built-in data transformations so you can quickly normalize, transform, and combine features without having to write any code. With Amazon SageMaker Data Wrangler's visualization templates, you can quickly preview and inspect that these transformations are completed as you intended by viewing them in Amazon SageMaker Studio. Once your data is prepared, you can build fully automated ML workflows with Amazon SageMaker Pipelines and save them for reuse in the Amazon SageMaker Feature Store.

Availability Amazon SageMaker Data Wrangler is now generally available in all regions where Amazon SageMaker Studio is available.

Links [Announcement](#)
[Landing page](#)

Sessions [Introduction video](#)
re:Invent portal video:
[NEW LAUNCH!] Accelerate data preparation with Amazon SageMaker Data Wrangler

Tags Machine learning, SageMaker, Pre-trained model, One-click deployment

[Home](#)

Amazon SageMaker Feature Store

Amazon SageMaker



Amazon SageMaker Feature Store is a fully managed repository that helps maintain consistency between features used at the time of inference and model training, so you can confidently deploy models in production with more predictable behaviour allowing you to operate ML models at scale. Amazon SageMaker Feature Store enables metadata management and discovery of features with easy tagging and search, so data science teams can simply reuse an existing feature instead of having to rewrite and process features for each new model. For real time predictions, features can be served with low millisecond latency or extracted for model training or batch prediction use cases from the feature store. Amazon SageMaker Feature Store manages historical records of feature data so that features can easily be reproduced at a specific point in time. With Amazon SageMaker Feature Store, you can accelerate machine learning, increase productivity and scale across thousands of models.

Availability SageMaker Feature Store is available in all regions where SageMaker is available.

Links [Announcement](#)
[Landing page](#)

Sessions [Introduction video](#)
re:Invent portal video:
[NEW LAUNCH!] Amazon SageMaker Feature Store: Store, discover, and share features for ML apps

Tags Machine learning, SageMaker, Feature Store, Curated data, Meta data

[Home](#)

Amazon SageMaker Clarify

Amazon SageMaker Clarify



Amazon SageMaker Clarify detects potential bias during data preparation, after training, and in your deployed model by examining attributes you specify. For instance, you can check for bias related to age in your initial dataset or in your trained model and receive a detailed report that quantifies different types of possible bias. SageMaker Clarify also includes feature importance graphs that help you explain model predictions and produces reports that can be used to support internal presentations or to identify issues with your model that you can take steps to correct.

Availability Amazon SageMaker Clarify is now generally available in all regions where Amazon SageMaker is available.

Links [Announcement](#)
[Landing Page](#)

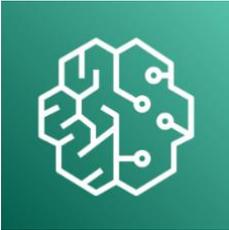
Sessions [Introduction Video - Bias detection](#)
[Introduction Video - Model explainability](#)
re:Invent portal video:
[\[NEW LAUNCH!\] Understand ML model predictions and biases with Amazon SageMaker Clarify](#)

Tags Machine learning, SageMaker, Feature importance, Fairness, Bias

[Home](#)

Amazon SageMaker JumpStart

Amazon SageMaker



Amazon SageMaker JumpStart provides a set of solutions for the most common use cases, such as fraud detection, predictive maintenance, and demand forecasting, that can be deployed readily with just a few clicks. The solutions are fully customizable and showcase the use of AWS CloudFormation templates and reference architectures so you can accelerate your ML journey. Amazon SageMaker JumpStart also provides one-click deployment and fine-tuning of more than 150 pre-trained models from popular model zoos, including PyTorch Hub and TensorFlow Hub. One-click deployment and fine-tuning features are available for natural language processing, object detection, and image classification models, so you can minimize the time to deploy open-source models for your own use case.

Availability Amazon SageMaker JumpStart is now generally available in all regions where Amazon SageMaker Studio is available.

Links [Announcement](#)
[Landing Page](#)

Sessions [Introduction Video - ML Solutions](#)
[Introduction Video - Model Zoo](#)
re:Invent portal video:
[\[New Launch!\] Get started with ML in minutes with Amazon SageMaker JumpStart](#)

Tags Machine learning, SageMaker, Pre-trained model, One-click

Amazon SageMaker Distributed training

Amazon SageMaker



Using partitioning algorithms, SageMaker distributed training automatically splits large deep learning models and training datasets across AWS GPU instances in a fraction of the time it takes to do manually. SageMaker achieves these efficiencies through two techniques: model parallelism and data parallelism. Model parallelism splits models too large to fit on a single GPU into smaller parts before distributing across multiple GPUs to train, and data parallelism splits large datasets to train concurrently in order to improve training speed.

With only a few lines of additional code, SageMaker will determine the best approach to split your model by using graph partitioning algorithms to balance the computation of each GPU while minimizing the communication between GPU instances. SageMaker also optimizes your distributed training jobs through algorithms that are designed to fully utilize AWS compute and network infrastructure in order to achieve near-linear scaling efficiency, which allows you to complete training faster than manual implementations.

Availability Amazon SageMaker Distributed training is now generally available in all regions where Amazon SageMaker is available.

Links [Landing page](#)

Sessions [Introduction video - Model Parallelism](#)

[Introduction video - Data Parallelism](#)

re:Invent portal video:

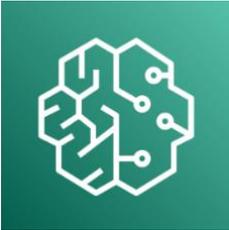
Fast training and near-linear scaling with DataParallel in Amazon SageMaker

Train billion-parameter models with model parallelism on Amazon SageMaker

[Home](#)

Amazon SageMaker Debugger and Profiling

Amazon SageMaker



Amazon SageMaker Debugger is a capability of Amazon SageMaker that makes it easy to train ML models faster by capturing real-time metrics such as learning gradients and weights, providing transparency into the training process, so you can correct anomalies such as losses, over-fitting, and over-training. SageMaker Debugger provides built-in techniques called rules to easily analyse emitted data including tensors that are critical for the success of training jobs such as identifying why your ML model is predicting a right traffic signal as left even though it trained at over 90% accuracy.

With new profiling capabilities, SageMaker Debugger now automatically monitors system resources such as CPU, GPU, network, I/O, and memory providing a complete resource utilization view of training jobs. You can also profile your entire training job, or portions thereof, to emit detailed framework metrics during different phases of the training job.

Availability Amazon SageMaker Debugger is now generally available in all AWS regions in the Americas and Europe, and some regions in Asia Pacific with additional regions coming soon.

Links [Announcement](#)
[Landing page](#)

Sessions [Deep profiling with Amazon SageMaker Debugger](#)

re:Invent portal video:

Train ML models faster with better insights using Amazon SageMaker Debugger

Tags Machine learning, SageMaker, Debugger, Monitor resources, Profiling

[Home](#)

Amazon SageMaker Model Monitor

Amazon SageMaker



Amazon SageMaker Model Monitor currently supports detecting data quality drift by tracking the difference between data that was used to train the models versus the data that is being presented to the model to score and alerting you of deviations to help you take timely actions such as auditing data or retraining models. Today, we are adding three new capabilities to SageMaker Model Monitor, enabling you to detect drift in model quality, model bias, and feature importance.

With model quality monitoring, you can monitor model characteristics (such as precision, accuracy, recall, and more) of your ML models in real time. SageMaker Model monitor reports how well a ML model is predicting outcomes by comparing model prediction to ground truth data. As the model is monitored, you can view exportable reports and graphs detailing model quality in Amazon S3, Amazon SageMaker Studio, and SageMaker Notebook instance. You can also configure Amazon CloudWatch to receive notifications if drift in model quality is observed.

Availability Amazon SageMaker Model Monitor is now generally available in all regions where Amazon SageMaker is available.

Links [Announcement](#)
[Landing page](#)

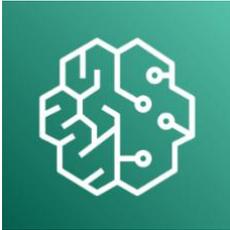
Sessions re:Invent portal video:
Detect machine learning (ML) model drift in production

Tags Machine learning, SageMaker, Model bias, Model quality, Detect drift

[Home](#)

Amazon SageMaker Edge Manager

Amazon SageMaker



Amazon SageMaker Edge Manager provides a software agent that runs on edge devices. The agent comes with a ML model optimized with SageMaker Neo automatically so you don't need to have Neo runtime installed on your devices in order to take advantage of the model optimizations. The agent also collects prediction data and sends a sample of the data to the cloud for monitoring, labeling, and retraining so you can keep models accurate over time.

All data can be viewed in the SageMaker Edge Manager dashboard which reports on the operation of deployed models. And, because SageMaker Edge Manager enables you to manage models separately from the rest of the application, you can update the model and the application independently reducing costly downtime and service disruptions. SageMaker Edge Manager also cryptographically signs your models so you can verify that it was not tampered with as it moves from the cloud to edge devices.

Availability US East (N. Virginia), US East (Ohio), US West (Oregon), EU (Ireland), EU (Frankfurt), and Asia Pacific (Tokyo)

Links [Announcement](#)
[Landing page](#)

Sessions [Introduction video](#)
re:Invent portal video:
[NEW LAUNCH!] MLOps for edge devices with Amazon SageMaker Edge Manager

Tags Machine learning, SageMaker, edge, IoT, industrial drift, neo

[Home](#)

Amazon SageMaker Pipelines

Amazon SageMaker



Amazon SageMaker Pipelines enables data science and engineering teams to collaborate seamlessly on ML projects and streamline building, automating, and scaling of end-to-end ML workflows. Amazon SageMaker SDK makes it easy to construct model building pipelines by defining the parameters and steps which can include Amazon SageMaker Data Wrangler,

Processing, Training, Batch Transform, conditional evaluation, and registering models to the central model registry.

Once the pipelines are built, Amazon SageMaker takes care of the execution of the pipelines and you can view the pipeline executions and the real-time metrics and logs for each step in Amazon SageMaker Studio. Models are registered to the new Amazon SageMaker model registry which automatically versions new models generated from pipelines and offers built-in approval workflows to select which models are deployed to production.

Availability Amazon SageMaker Pipelines is now generally available in all AWS Regions where Amazon SageMaker is available and the MLOps capabilities of Amazon SageMaker Pipelines are only available in the AWS Regions where AWS CodePipeline is also available.

Links [Announcement](#)
[Landing page](#)

Sessions [Introduction video](#)

re:Invent portal video:

[NEW LAUNCH!] How to create fully automated ML workflows with Amazon SageMaker Pipelines

Tags Machine learning, SageMaker, Pipeline, MLOps

[Home](#)

Amazon Neptune ML

Amazon Neptune



Amazon Neptune that uses Graph Neural Networks (GNNs), a machine learning (ML) technique purpose-built for graphs, to make easy, fast, and accurate predictions using graph data. With GNNs, you can improve the accuracy of most predictions for graphs by over 50% when compared to making predictions using non-graph methods based on published research from

Stanford University.

Using the Deep Graph Library (DGL), an open-source library to which AWS contributes that makes it easy to apply deep learning to graph data, Neptune ML automates the heavy lifting of selecting and training the best ML model for graph data, and lets users run ML on their graph directly using Neptune APIs and queries. As a result, you can now create, train, and apply ML on Neptune data in hours instead of weeks without the need to learn new tools and ML technologies.

Availability Neptune ML is generally available to customers using Neptune engine version 1.0.4.1 and in all regions Neptune is available.

Links [Announcement](#)
[Landing page](#)

Sessions re:Invent portal video:
Deep dive on Amazon Neptune
New capabilities to build graph apps quickly with Amazon Neptune

Tags Machine learning, GNN, Deep graph learning, ML on graph data

[Home](#)

Amazon Redshift ML

Amazon Redshift



With Amazon Redshift ML powered by Amazon SageMaker, you can use SQL statements to create and train machine learning models from your data in Amazon Redshift and then use these models for use cases such as churn prediction and fraud risk scoring directly in your queries and reports. Amazon Redshift ML automatically discovers and tunes the best model based on the training data using Amazon SageMaker Autopilot. SageMaker Autopilot chooses among the best regression, binary, or multi-class classification and linear models. Alternatively, you can choose a model type such as Xtreme Gradient Boosted tree (XGBoost), a problem type like regression or classification, and preprocessors or hyperparameters. Amazon Redshift ML uses your parameters to build, train, and deploy the model in the Amazon Redshift data warehouse. You can obtain predictions from these trained models using SQL queries as if you were invoking a user defined function (UDF) and leverage all benefits of Amazon Redshift, including massively parallel processing capabilities.

Availability The Redshift ML preview is available in the following regions: US East (Ohio), US East (N Virginia), US West (Oregon), US West (San Francisco), Canada (Central), Europe (Frankfurt), Europe (Ireland), Europe(London), Europe (Paris), Europe (Stockholm), Asia Pacific (Hong Kong) Asia Pacific (Tokyo), Asia Pacific (Singapore), Asia Pacific (Sydney), and South America (São Paulo).

Links [Announcement](#)
[Landing page](#)

Sessions [Demo video](#)

re:Invent portal video:

[NEW LAUNCH!] Introducing Amazon Redshift Machine Learning

Tags Machine learning, SageMaker, SQL, UDF

[Home](#)

Amazon Athena ML

Amazon Athena



Amazon Athena released a new feature that allows users to easily invoke machine learning models for inference directly from their SQL queries. The ability to use machine learning models in SQL queries makes complex tasks such as anomaly detection, customer cohort analysis, and sales predictions as simple as invoking a function in a SQL query.

Customers can use more than a dozen built-in machine learning algorithms provided by Amazon SageMaker, train their own models, or find and subscribe to model packages from AWS Marketplace and deploy on Amazon SageMaker Hosting Services. There is no additional setup required. Users can invoke these ML models in their SQL queries from the Athena console, Athena APIs, and through Athena's preview JDBC driver.

Availability Amazon Athena's ML functionality is available today in preview in the us-east-1 (N. Virginia) region. Begin your preview now by following these [steps](#).

Links [Announcement](#)
[User guide](#)

Tags Machine learning, SageMaker, SQL, Athena, JDBC

[Home](#)

Amazon Aurora ML

Amazon Aurora



Amazon Aurora machine learning enables you to add ML-based predictions to applications via the familiar SQL programming language, so you don't need to learn separate tools or have prior machine learning experience. It provides simple, optimized, and secure integration between Aurora and AWS ML services without having to build custom integrations or move data around.

When you run an ML query, Aurora calls [Amazon SageMaker](#) for a wide variety of ML algorithms or [Amazon Comprehend](#) for sentiment analysis, so your application doesn't need to call these services directly. This makes Aurora machine learning suitable for low-latency, real-time use cases such as fraud detection, ad targeting, and product recommendations. For example, you can build product recommendation systems by writing SQL queries in Aurora that pass customer profile, shopping history, and product catalog data to a SageMaker model, and get product recommendations returned as query results.

Availability Aurora machine learning integration is available for Amazon Aurora with MySQL 5.7 compatibility, PostgreSQL 10 compatibility, and PostgreSQL 11 compatibility.

Links [Landing page](#)
[Announcement](#)

Sessions re:Invent portal video:
What's new in Amazon Aurora

Tags Machine learning, SageMaker, SQL, aurora, RDBMS, RDS

[Home](#)

AWS Trainium

Amazon EC2, Amazon SageMaker



AWS Trainium is the second custom machine learning (ML) chip designed by AWS that provides the best price performance for training ML models in the cloud. In addition to delivering the most cost-effective ML training, Trainium offers the highest performance with the most teraflops (TFLOPS) of compute power for ML in the cloud and enables a broader set of ML applications. The Trainium chip is specifically optimized for deep learning training workloads for applications including image classification, semantic search, translation, voice recognition, natural language processing and recommendation engines.

Availability AWS Trainium will be available in 2021. To be notified about early access to AWS Trainium, [sign up here](#).

Links [Landing page](#)

Tags Machine learning, SageMaker, EC2, intel, GPU, TensorFlow, PyTorch, MXNet

[Home](#)

AI Services

[Home](#)

Amazon Monitron

Amazon Monitron



Amazon Monitron is an end-to-end system that monitors the condition of equipment in your facilities and automatically detects abnormal behaviour in industrial machinery, enabling you to take proactive action on potential failures and reduce unplanned downtime by implementing a predictive maintenance program.

Monitron sensors capture vibration and temperature data from rotating machines, such as bearings, gearboxes, motors, pumps, compressors, and fans. Sensors send vibration and temperature measurements hourly to a nearby Monitron gateway, using Bluetooth Low Energy (BLE) technology allowing the sensors to run for at least three years. The Monitron gateway is itself connected to your WiFi network, and sends sensor data to AWS, where it is stored and analysed using machine learning and ISO 20816 vibration standards.



Availability Is available today at US East (N. Virginia)

Links [Blog post](#)

[Landing page](#)

Sessions [Brief overview video](#)

[Introduction video](#)

re:Invent portal video:

[NEW LAUNCH!] Enable predictive maintenance for your industrial equipment

[Home](#)

Amazon Lookout for Equipment

Amazon Lookout



Amazon Lookout for Equipment uses the data from your sensors to detect abnormal equipment behaviour, so you can take action before machine failures occur and avoid unplanned downtime.

Customers that want to build ML models to monitor the health or efficiency of their equipment can directly upload their historical sensor data to Amazon Lookout for Equipment and automatically build a ML model that learns the normal behaviour patterns and alerts to abnormal behaviour. Customers can set up Amazon Lookout for Equipment to read real-time data from their equipment and detect the current behaviour of the asset.

Availability Is in preview today at US East (N. Virginia), Asia Pacific (Seoul), and Europe (Ireland)

Links [Blog post](#)
[Landing page](#)

Sessions re:Invent portal video:
[NEW LAUNCH!] Detect abnormal equipment behavior by analyzing sensor data

Tags AI, IoT, Hardware

[Home](#)

Amazon Lookout for Vision

Amazon Lookout



Amazon Lookout for Vision is a machine learning (ML) service that spots defects and anomalies in visual representations using computer vision (CV). With Amazon Lookout for Vision, companies can increase quality and reduce operational costs by quickly identifying differences in images at scale. For example, identifying missing components in products, damage to vehicles or structures, irregularities in production lines, miniscule defects in silicon wafers, and other visual inspection use cases.

Amazon Lookout for Vision uses ML to see and understand images from any camera as a person would, but with an even higher degree of accuracy and at a much larger scale, allowing customers to eliminate the need for costly and inconsistent manual inspection while improving quality control, defect and damage assessment, and compliance. In minutes, you can begin using Amazon Lookout for Vision to automate inspection of images and objects – with no machine learning expertise required.

Availability Is in preview today at US East (N. Virginia), US East (Ohio), US West (Oregon), Asia Pacific (Seoul), Asia Pacific (Tokyo), Europe (Frankfurt) and Europe (Ireland)

Links [Blog post](#)
[Landing page](#)

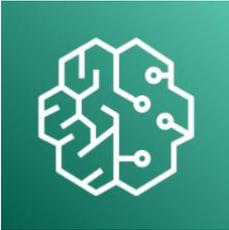
Sessions [Discussion video](#)
re:Invent portal video:
[NEW LAUNCH!] Amazon Lookout for Vision

Tags AI, IoT, Manufacturing

[Home](#)

Amazon Lookout for Metrics

Amazon Lookout



Amazon Lookout for Metrics uses machine learning (ML) to automatically detect and diagnose anomalies (i.e. outliers from the norm) in business and operational time series data, such as a sudden dip in sales revenue or customer acquisition rates. Start monitoring metrics that are important to your business by connecting to popular data stores like Amazon S3, Amazon Redshift, and Amazon Relational Database Service (RDS), as well as third-party SaaS applications, such as Salesforce, ServiceNow, Zendesk, and Marketo.

Amazon Lookout for Metrics automatically inspects and prepares the data from these sources and builds a custom ML model, informed by over 20 years of experience at Amazon, to detect anomalies with greater speed and accuracy than traditional methods used for anomaly detection. You can also provide feedback on detected anomalies to tune the results and improve accuracy over time. Amazon Lookout for Metrics makes it easy to diagnose detected anomalies by grouping together anomalies that are related to the same event and sending an alert that includes a summary of the potential root cause. It also ranks anomalies in order of severity so that you can prioritize your attention to what matters the most to your business.

Availability Now available to [request an invitation to the preview](#).

Links [Blog post](#)
[Landing page](#)

Sessions re:Invent portal video:
[NEW LAUNCH!] Proactively monitor the health of your business using Amazon Lookout for Metrics

Tags AI, anomaly detection, time series data

[Home](#)

AWS Panorama Appliance

AWS Panorama



AWS Panorama is a machine learning Appliance and accompanying Software Development Kit (SDK) that allows organizations to bring computer vision (CV) to on-premises cameras to make predictions locally with high accuracy and low latency. With AWS Panorama, you can automate tasks that have traditionally required human inspection to improve visibility into potential issues. For example, you can use AWS Panorama to evaluate manufacturing quality, identify bottlenecks in industrial processes, and monitor workplace safety and security - even in environments with limited or no internet connectivity.

The AWS Panorama Appliance is a hardware device that allows you to add CV to your internet protocol (IP) cameras that weren't built to accommodate computer vision. AWS Panorama Appliance turns your existing cameras into smart cameras that can run CV models on multiple concurrent video streams.

The AWS Panorama Device SDK enables third-party manufacturers to build new cameras that run more meaningful CV models at the edge for tasks like object detection or activity recognition. AWS Panorama-compatible cameras work out of the box with AWS machine learning services.

Availability Now available to [request an invitation to the preview](#).

Links [Blog post](#)
[Landing page](#)

Sessions re:Invent portal video:
[NEW LAUNCH!] Use computer vision at the edge to improve operations with AWS Panorama

Tags AI, ML, Computer Vision, Hardware

[Home](#)

Amazon DevOps Guru

Amazon DevOps Guru



Amazon DevOps Guru is a Machine Learning (ML) powered service that makes it easy to improve an application’s operational performance and availability. By detecting behaviours that deviate from normal operating patterns, you can identify operational issues long before they impact your customers.

With ready-to-use machine learning models, you can identify anomalous application behaviour (e.g., increased latency, error rates, resource constraints, etc.) and surface critical issues that could cause potential outages or service disruptions. When a critical issue is identified, an alert is sent automatically, providing a summary of related anomalies, the likely root cause, and context about when and where the issue occurred. When possible, recommendations on how to remediate the issue are provided.

Operational data is automatically ingested from your AWS applications and visualized on a unifying dashboard. No manual setup or machine learning expertise is required.

Availability Available for preview in US East (N. Virginia), US East (Ohio), US West (Oregon), Europe (Ireland), and Asia Pacific (Tokyo)

Links [Blog post](#)
[Landing page](#)

Sessions [Discussion video](#)

re:Invent portal video:
[NEW LAUNCH!] Improve application availability with ML-powered insights using Amazon DevOps Guru

Tags AI, ML, Application availability, DevOps

[Home](#)

Amazon Kendra Incremental Learning

Amazon Kendra



As unstructured data continues to grow within the enterprise, organizations need adaptive, intelligent, and nimble search solutions, particularly as search patterns and document trends change over time. Amazon Kendra's incremental learning capabilities enable it to learn how users interact with search results and adapt its models to re-rank information, so the most commonly sought content is presented first.

Incremental learning tunes future search results quickly in a way that's data-driven, and cost effective without the need for customers to train and deploy machine learning models.

Availability This is now generally available in all regions where Amazon Kendra is available.

Links [Announcement](#)
[Landing page](#)

Tags Machine learning, Search, Incremental learning, Improved search

[Home](#)

Amazon HealthLake

Amazon Health Lake



Amazon Health Lake is a HIPAA-eligible service that enables healthcare providers, health insurance companies, and pharmaceutical companies to store, transform, query, and analyse health data at petabyte scale. It removes the heavy lifting of organizing, indexing, and structuring patient information, to provide a complete view of the health of individual patients and entire patient populations in a secure, compliant, and auditable manner.

Availability In preview at US East (N. Virginia). [Sign up here](#)

Links [Blog post](#)
[Landing page](#)

Sessions [Introduction video](#)
re:Invent portal video:
[NEW LAUNCH!] Make sense of health data with Amazon HealthLake

Tags AI, Datalake, Health

[Home](#)

Amazon Forecast Weather Index

Amazon Forecast



Changes in local weather conditions can impact short term demand for products and services at particular locations for many customers in retail, hospitality, travel, entertainment, insurance and energy domains. While historical demand patterns show seasonal demand, advance planning for day-to-day variation is harder. In retail inventory management use cases, day-to-day weather variation impacts foot traffic and product mix. Typical demand forecasting systems do not take expected weather conditions into account.

The Amazon Forecast Weather Index combines multiple weather metrics from historical weather events and current forecasts at a given location to increase your demand forecast model accuracy. Amazon Forecast uses machine learning to generate more accurate demand forecasts, without requiring any prior ML experience. Forecast brings the same technology used at Amazon.com to developers as a fully managed service, removing the need for developers to manage resources or re-build their systems.

Availability This is now generally available in all regions where Amazon Forecast is available.

Links [Press release](#)
[Blog post](#)

Tags Machine learning, Forecast, Historical weather index, Weather forecast

[Home](#)

Sydney Region Availability

[Home](#)
Service
AP-SOUTHEAST-2
Data Analytics Services

AWS Glue DataBrew	✓
AWS Glue Elastic Views	✗ ¹
AWS Lake Formation New Features	✗ ¹
Amazon Redshift Data Sharing	✓
Amazon AQUA for Amazon Redshift	✗ ¹
Amazon Redshift federated querying MySQL	✓
Amazon Redshift console partner integration	✓
Amazon Redshift native JSON support	✓
Amazon Managed Workflows for Apache Airflow (MWAA)	✓
Amazon QuickSight Q	✗ ¹
Amazon EMR Studio	✗ ¹
Amazon EMR on Amazon EKS	✗

Machine Learning Services

Amazon SageMaker Data Wrangler	✓
Amazon SageMaker Feature Store	✓
Amazon SageMaker Clarify	✓
Amazon SageMaker JumpStart	✓
Amazon SageMaker Distributed training	✓
Amazon SageMaker Debugger and Profiling	✓
Amazon SageMaker Model Monitor	✓
Amazon SageMaker Edge Manager	✗
Amazon SageMaker Pipelines	✓
Amazon Neptune ML	✓
Amazon Redshift ML	✓
Amazon Athena ML	✗ ¹
Amazon Aurora ML	✓
AWS Trainium	✗

¹ In Preview

[Home](#)

Service

AP-SOUTHEAST-2

AI Services

Amazon Monitron	✘
Amazon Lookout for Equipment	✘ ¹
Amazon Lookout for Vision	✘ ¹
Amazon Lookout for Metrics	✘ ¹
AWS Panorama Appliance	✘ ¹
Amazon DevOps Guru	✘
Amazon Kendra Incremental Learning	✓
Amazon HealthLake	✘ ¹
Amazon Forecast Weather Index	✓