



Reference

AWS Glossary



AWS Glossary: Reference

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

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

Table of Contents

AWS Glossary 1

AWS Glossary

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

Numbers and symbols

100-continue A method that gives a client the ability to see whether a server can accept a request before actually sending it. For large PUT requests, this method can save both time and bandwidth charges.

A

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

AAD See [additional authenticated data](#).

access control list (ACL) A document that defines who can access a particular [bucket](#) or object. Each [bucket](#) and object in [Amazon S3](#) has an ACL. This document defines what each type of user can do, such as write and read permissions.

access identifiers See [credentials](#).

access key The combination of an [access key ID](#) (for example, AKIAIOSFODNN7EXAMPLE) and a [secret access key](#) (for example, wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY). You use access keys to sign API requests that you make to AWS.

access key ID A unique identifier that's associated with a [secret access key](#); the access key ID and secret access key are used together to sign programmatic AWS requests cryptographically.

access key rotation A method to increase security by changing the AWS access key ID. You can use this method to retire an old key at your discretion.

access policy language A language for writing documents (specifically, [policies](#)) that specify who can access a particular AWS [resource](#) and under what conditions.

account	<p>A formal relationship with AWS that's associated with all of the following:</p> <ul style="list-style-type: none">• The owner email address and password• The control of resources created under its umbrella• Payment for the AWS activity related to those resources <p>The AWS account has permission to do anything and everything with all the AWS account resources. This is in contrast to a user, which is an entity contained within the account.</p>
account activity	<p>A webpage showing your month-to-date AWS usage and costs. The account activity page is located at https://aws.amazon.com/account-activity/.</p>
AWS Account Management	<p>AWS Account Management is a tool that you can use to update the contact information for each of your AWS accounts.</p> <p>See Also https://aws.amazon.com/organizations.</p>
ACL	<p>See access control list (ACL).</p>
ACM	<p>AWS Certificate Manager is a web service for provisioning, managing, and deploying Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services.</p> <p>See Also https://aws.amazon.com/certificate-manager/.</p>
action	<p>An API function. Also called <i>operation</i> or <i>call</i>. The activity the principal has permission to perform. The action is B in the statement "A has permission to do B to C where D applies." For example, Jane sends a request to Amazon SQS with Action=ReceiveMessage.</p> <p>CloudWatch: The response initiated by the change in an alarm's state (for example, from OK to ALARM). The state change might be caused by a metric reaching the alarm threshold, or by a SetAlarmState request. Each alarm can have one or more actions assigned to each state. Actions are performed once each time the alarm changes to a state that has an action assigned. Example actions include an Amazon SNS notification, running an Amazon EC2 Auto Scaling policy, and an Amazon EC2 instance stop/terminate action.</p>

active trusted key groups	A list that shows each of the trusted key groups , and the IDs of the public keys in each key group, that are active for a distribution in Amazon CloudFront. CloudFront can use the public keys in these key groups to verify the signatures of CloudFront signed URLs and signed cookies .
active trusted signers	See active trusted key groups .
active-active	A class of high availability strategies in which a workload exists simultaneously in multiple Regions, uses multiple primary resources, and serves traffic from all of the Regions to which it's deployed. Sometimes referred to as active/active. See Also , , .
active-passive	A class of disaster recovery strategies that involve a primary Region and a standby Region in a back up and restore , hot standby , pilot light , or warm standby configuration. Sometimes referred to as active/passive.
additional authenticated data	Information that's checked for integrity but not encrypted, such as headers or other contextual metadata.
administrative suspension	Amazon EC2 Auto Scaling might suspend processes for Auto Scaling group that repeatedly fail to launch instances. Auto Scaling groups that most commonly experience administrative suspension have zero running instances, have been trying to launch instances for more than 24 hours, and have not succeeded in that time.
alarm	An item that watches a single metric over a specified time period and starts an Amazon SNS topic or an Amazon EC2 Auto Scaling policy . These actions are started if the value of the metric crosses a threshold value over a predetermined number of time periods.
allow	One of two possible outcomes (the other is deny) when an IAM access policy is evaluated. When a user makes a request to AWS, AWS evaluates the request based on all permissions that apply to the user and then returns either allow or deny.
Amazon Machine Image (AMI)	An Amazon Machine Image (AMI) is an encrypted machine image stored in Amazon EBS or Amazon S3 . AMIs function similarly to a template of a computer's root drive. They contain the operating system and can

	also include software and layers of your application, such as database servers, middleware, and web servers.
Amazon Web Services (AWS)	<p>An infrastructure web services platform in the cloud for companies of all sizes.</p> <p>See Also https://aws.amazon.com/what-is-cloud-computing/.</p>
AMI	See Amazon Machine Image (AMI) .
Amplify	<p>AWS Amplify is a complete solution that frontend web and mobile developers can use to build and deploy secure, scalable full-stack applications powered by AWS. Amplify provides two services: Amplify Hosting and Amplify Studio.</p> <p>See Also https://aws.amazon.com/amplify/ .</p>
Amplify Android	<p>Amplify Android is a collection of open-source client libraries that provides interfaces for specific use cases across many AWS services. Amplify Android is the recommended way to build native Android applications powered by AWS.</p> <p>See Also https://aws.amazon.com/amplify/ .</p>
Amplify Hosting	<p>AWS Amplify Hosting is a fully managed continuous integration and continuous delivery (CI/CD) and hosting service for fast, secure, and reliable static and server-side rendered apps. Amplify Hosting provides a Git-based workflow for hosting full-stack serverless web apps with continuous deployment.</p> <p>See Also https://aws.amazon.com/amplify/hosting/ .</p>
Amplify iOS	<p>Amplify iOS is a collection of open-source client libraries that provides interfaces for specific use cases across many AWS services. Amplify iOS is the recommended way to build native iOS applications powered by AWS.</p> <p>See Also https://aws.amazon.com/amplify/ .</p>
Amplify Studio	<p>AWS Amplify Studio is a visual development environment that web and mobile developers can use to build the frontend UI components and the backend environment for a full-stack application.</p> <p>See Also https://aws.amazon.com/amplify/studio/ .</p>
analysis rules	AWS Clean Rooms : The query restrictions that authorize a specific type of query.

analysis scheme	CloudSearch : Language-specific text analysis options that are applied to a text field to control stemming and configure stopwords and synonyms.
API Gateway	Amazon API Gateway is a fully managed service that developers can use to create, publish, maintain, monitor, and secure APIs at any scale. See Also https://aws.amazon.com/api-gateway .
AWS App2Container	AWS App2Container is a transformation tool that modernizes .NET and Java applications by migrating them into containerized applications. See Also https://aws.amazon.com/app2container .
AWS AppConfig	AWS AppConfig is a service used to update software at runtime without deploying new code. With AWS AppConfig, you can configure, validate, and deploy feature flags and application configurations. See Also https://aws.amazon.com/systems-manager/features/appconfig .
Amazon AppFlow	Amazon AppFlow is a fully managed integration service that you can use to transfer data securely between software as a service (SaaS) applications and AWS services. See Also https://aws.amazon.com/appflow .
application	Elastic Beanstalk : A logical collection of components, including environments, versions, and environment configurations. An application is conceptually similar to a folder. CodeDeploy : A name that uniquely identifies the application to be deployed. AWS CodeDeploy uses this name to ensure the correct combination of revision, deployment configuration, and deployment group are referenced during a deployment.
Application Auto Scaling	AWS Application Auto Scaling is a web service that you can use to configure automatic scaling for AWS resources beyond Amazon EC2, such as Amazon ECS services, Amazon EMR clusters, and DynamoDB tables. See Also https://aws.amazon.com/autoscaling/ .
Application Billing	The location where your customers manage the Amazon DevPay products they've purchased. The web address is http://www.amazon.com/dp-applications .

Application Composer	AWS Application Composer is a visual designer that you can use to build serverless applications from multiple AWS services. As you design an application, Application Composer automatically generates a YAML template with CloudFormation and AWS SAM template resources. See Also https://aws.amazon.com/application-composer/ .
Application Cost Profiler	AWS Application Cost Profiler is a solution to track the consumption of shared AWS resources used by software applications and report granular cost breakdown across tenant base. See Also https://aws.amazon.com/aws-cost-management/aws-application-cost-profiler/ .
Application Discovery Service	AWS Application Discovery Service is a web service that helps you plan to migrate to AWS by identifying IT assets in a data center—including servers, virtual machines, applications, application dependencies, and network infrastructure. See Also https://aws.amazon.com/application-discovery/ .
application revision	CodeDeploy : An archive file containing source content—such as source code, webpages, executable files, and deployment scripts—along with an application specification file . Revisions are stored in Amazon S3 buckets or GitHub repositories. For Amazon S3, a revision is uniquely identified by its Amazon S3 object key and its ETag, version, or both. For GitHub, a revision is uniquely identified by its commit ID.
application specification file	CodeDeploy : A YAML-formatted file used to map the source files in an application revision to destinations on the instance. The file is also used to specify custom permissions for deployed files and specify scripts to be run on each instance at various stages of the deployment process.
application version	Elastic Beanstalk : A specific, labeled iteration of an application that represents a functionally consistent set of deployable application code. A version points to an Amazon S3 object (a JAVA WAR file) that contains the application code.
AppSpec file	See application specification file .
AppStream 2.0	Amazon AppStream 2.0 is a fully managed, secure service for streaming desktop applications to users without rewriting those applications.

	See Also https://aws.amazon.com/appstream/ .
AWS AppSync	AWS AppSync is an enterprise-level, fully managed GraphQL service with real-time data synchronization and offline programming features. See Also https://aws.amazon.com/appsync/ .
ARN	See Amazon Resource Name (ARN) .
artifact	CodePipeline : A copy of the files or changes that are worked on by the pipeline.
asymmetric encryption	Encryption that uses both a public key and a private key.
asynchronous bounce	A type of bounce that occurs when a receiver initially accepts an email message for delivery and then subsequently fails to deliver it.
Athena	Amazon Athena is an interactive query service that you can use to analyze data in Amazon S3 using ANSI SQL. Athena is serverless, so there's no infrastructure to manage. Athena scales automatically and is simple to use, so you can start analyzing your datasets within seconds. See Also https://aws.amazon.com/athena/ .
atomic counter	DynamoDB: A method of incrementing or decrementing the value of an existing attribute without interfering with other write requests.
attribute	<p>A fundamental data element, something that doesn't need to be broken down any further. In DynamoDB, attributes are similar in many ways to fields or columns in other database systems.</p> <p>Amazon Machine Learning: A unique, named property within an observation in a dataset. In tabular data, such as spreadsheets or comma-separated values (.csv) files, the column headings represent the attributes, and the rows contain values for each attribute.</p>
AUC	Area Under a Curve. An industry-standard metric to evaluate the quality of a binary classification machine learning model. AUC measures the ability of the model to predict a higher score for positive examples, those that are "correct," than for negative examples, those that are "incorrect." The AUC metric returns a decimal value from 0 to 1. AUC values near 1 indicate an ML model that's highly accurate.

Aurora	Amazon Aurora is a fully managed MySQL-compatible relational database engine that combines the speed and availability of commercial databases with the simplicity and cost-effectiveness of open-source databases. See Also https://aws.amazon.com/rds/aurora/ .
authenticated encryption	Encryption that provides confidentiality, data integrity, and authenticity assurances of the encrypted data.
authentication	The process of proving your identity to a system.
AWS Auto Scaling	AWS Auto Scaling is a fully managed service that you can use to quickly discover the scalable AWS resources that are part of your application and to configure dynamic scaling. See Also https://aws.amazon.com/autoscaling/ .
Auto Scaling group	A representation of multiple EC2 instances that share similar characteristics, and that are treated as a logical grouping for the purposes of instance scaling and management.
Availability Zone	A distinct location within a Region that's insulated from failures in other Availability Zones, and provides inexpensive, low-latency network connectivity to other Availability Zones in the same Region.
AWS	See Amazon Web Services (AWS) .

B

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

back up and restore	A disaster recovery strategy in which backups of data in the primary Region are copied to a standby Region and can be restored from the standby Region. You must provision the infrastructure and other resources, such as compute, as part of a failover process. See Also , , , .
Backint Agent	AWS Backint Agent for SAP HANA is an SAP-certified backup and restore solution for SAP HANA workloads running on Amazon EC2 instances in the cloud.

	See Also https://aws.amazon.com/backint-agent .
AWS Backup	AWS Backup is a managed backup service that you can use to centralize and automate the backup of data across AWS services in the cloud and on premises. See Also https://aws.amazon.com/backup/ .
basic monitoring	Monitoring of AWS provided metrics derived at a 5-minute frequency.
batch	See document batch .
batch prediction	Amazon Machine Learning: An operation that processes multiple input data observations at one time (asynchronously). Unlike real-time predictions, batch predictions aren't available until all predictions have been processed. See Also real-time predictions .
BGP ASN	Border Gateway Protocol Autonomous System Number is a unique identifier for a network, for use in BGP routing. Amazon EC2 supports all 2-byte ASN numbers in the range of 1 – 65335, with the exception of 7224, which is reserved.
billing	See Billing and Cost Management .
Billing and Cost Management	AWS Billing and Cost Management is the AWS Cloud computing model where you pay for services on demand and use as much or as little as you need. While resources are active under your account, you pay for the cost of allocating those resources. You also pay for any incidental usage associated with those resources, such as data transfer or allocated storage. See Also https://aws.amazon.com/billing/new-user-faqs/ .
binary attribute	Amazon Machine Learning: An attribute for which one of two possible values is possible. Valid positive values are 1, y, yes, t, and true answers. Valid negative values are 0, n, no, f, and false. Amazon Machine Learning outputs 1 for positive values and 0 for negative values. See Also attribute .
binary classification model	Amazon Machine Learning: A machine learning model that predicts the answer to questions where the answer can be expressed as a binary

	<p>variable. For example, questions with answers of “1” or “0”, “yes” or “no”, “will click” or “will not click” are questions that have binary answers. The result for a binary classification model is always either a “1” (for a “true” or affirmative answers) or a “0” (for a “false” or negative answers).</p>
block	<p>A dataset. Amazon EMR breaks large amounts of data into subsets. Each subset is called a data block. Amazon EMR assigns an ID to each block and uses a hash table to keep track of block processing.</p>
block device	<p>A storage device that supports reading and (optionally) writing data in fixed-size blocks, sectors, or clusters.</p>
block device mapping	<p>A mapping structure for every AMI and instance that specifies the block devices attached to the instance.</p>
AWS Blockchain Templates	<p>See Managed Blockchain.</p>
blue/green deployment	<p>CodeDeploy: A deployment method where the instances in a deployment group (the original environment) are replaced by a different set of instances (the replacement environment).</p>
bootstrap action	<p>A user-specified default or custom action that runs a script or an application on all nodes of a job flow before Hadoop starts.</p>
Border Gateway Protocol Autonomous System Number	<p>See BGP ASN.</p>
bounce	<p>A failed email delivery attempt.</p>
Braket	<p>Amazon Braket is a fully managed quantum computing service that helps you run quantum algorithms to accelerate your research and discovery.</p> <p>See Also https://aws.amazon.com/braket.</p>
breach	<p>Amazon EC2 Auto Scaling: The condition where a user-set threshold (upper or lower boundary) is passed. If the duration of the breach is significant, as set by a breach duration parameter, it can possibly start a scaling activity.</p>

bucket	Amazon S3 : A container for stored objects. Every object is contained in a bucket. For example, if the object named photos/puppy.jpg is stored in the amzn-s3-demo-bucket bucket, then authorized users can access the object with the URL <code>https://amzn-s3-demo-bucket.s3.region-code.amazonaws.com/photos/puppy.jpg</code> .
bucket owner	The person or organization that owns a bucket in Amazon S3 . In the same way that Amazon is the only owner of the domain name Amazon.com, only one person or organization can own a bucket.
bundling	A commonly used term for creating an Amazon Machine Image (AMI) . It specifically refers to creating instance store-backed AMIs .

C

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

cache cluster	A logical cache distributed over multiple cache nodes . A cache cluster can be set up with a specific number of cache nodes.
cache cluster identifier	Customer-supplied identifier for the cache cluster that must be unique for that customer in an AWS Region .
cache engine version	The version of the Memcached service that's running on the cache node.
cache node	A fixed-size chunk of secure, network-attached RAM. Each cache node runs an instance of the Memcached service, and has its own DNS name and port. Multiple types of cache nodes are supported, each with varying amounts of associated memory.
cache node type	An EC2 instance type used to run the cache node.
cache parameter group	A container for cache engine parameter values that can be applied to one or more cache clusters.
cache security group	A group maintained by ElastiCache that combines inbound authorizations to cache nodes for hosts belonging to Amazon EC2 security groups that are specified through the console or the API or command line tools.

campaign	<p>Amazon Personalize: A deployed solution version (trained model) with provisioned dedicated transaction capacity for creating real-time recommendations for your application users. After you create a campaign, you use the <code>getRecommendations</code> or <code>getPersonalizedRanking</code> personalization operations to get recommendations.</p> <p>See Also recommendations, solution version.</p>	
canned access policy	A standard access control policy that you can apply to a bucket or object. Options include: private, public-read, public-read-write, and authenticated-read.	
canonicalization	The process of converting data into a standard format that a service such as Amazon S3 can recognize.	
capacity	The amount of available compute size at a given time. Each Auto Scaling group is defined with a minimum and maximum compute size. A scaling activity increases or decreases the capacity within the defined minimum and maximum values.	
Cartesian product	A mathematical operation that returns a product from multiple sets.	
Cartesian product processor	A processor that calculates a Cartesian product. Also known as a <i>Cartesian data processor</i> .	
AWS CDK	AWS Cloud Development Kit (AWS CDK) is an open-source software development framework for defining your cloud infrastructure in code and provisioning it through AWS CloudFormation.	See Also https://aws.amazon.com/cdk/ .
CDN	See content delivery network (CDN) .	
certificate	A credential that some AWS products use to authenticate AWS accounts and users. Also known as an X.509 certificate . The certificate is paired with a private key.	
chargeable resources	Features or services whose use incurs fees. Although some AWS products are free, others include charges. For example, in an CloudFormation stack , AWS resources that have been created incur charges. The amount charged depends on the usage load. Use the Amazon Web Services	

	Simple Monthly Calculator to estimate your cost prior to creating instances, stacks, or other resources.
Amazon Q Developer in chat applications	Amazon Q Developer in chat applications is an interactive agent that makes it easier to monitor, troubleshoot, and operate AWS resources in your Slack channels and Amazon Chime chat rooms. See Also https://aws.amazon.com/chatbot .
Amazon Chime	Amazon Chime is a secure, real-time, unified communications service that transforms meetings by making them more efficient and easier to conduct. See Also https://aws.amazon.com/chime/ .
CIDR block	Classless Inter-Domain Routing. An internet protocol address allocation and route aggregation methodology. See Also Classless Inter-Domain Routing on Wikipedia .
ciphertext	Information that has been encrypted , as opposed to plaintext , which is information that has not.
classification	In machine learning, a type of problem that seeks to place (classify) a data sample into a single category or "class." Often, classification problems are modeled to choose one category (class) out of two. These are binary classification problems. Problems with more than two available categories (classes) are called "multiclass classification" problems. See Also binary classification model , multiclass classification model .
AWS Clean Rooms	AWS Clean Rooms is an AWS service that helps multiple parties to join their data together in a secure collaboration workspace. See Also https://aws.amazon.com/clean-rooms/ .
Client VPN	AWS Client VPN is a client-based, managed VPN service that remote clients can use to securely access your AWS resources using an Open VPN-based software client. See Also https://aws.amazon.com/vpn/client-vpn .
AWS Cloud Control API	AWS Cloud Control API is a set of standardized application programming interfaces (APIs) that developers can use to create, read, update, delete, and list supported cloud infrastructure.

	See Also https://aws.amazon.com/cloudcontrolapi .
Cloud Directory	Amazon Cloud Directory is a service that provides a highly scalable directory store for your application's multihierarchical data. See Also https://aws.amazon.com/cloud-directory/ .
AWS Cloud Map	AWS Cloud Map is a service that you use to create and maintain a map of the backend services and resources that your applications depend on. With AWS Cloud Map, you can name and discover your AWS Cloud resources. See Also https://aws.amazon.com/cloud-map .
cloud service provider (CSP)	A cloud service provider is a company that provides subscribers with access to internet-hosted computing, storage, and software services.
AWS Cloud WAN	AWS Cloud WAN is a managed wide-area networking service used to build, manage, and monitor a unified global network. See Also https://aws.amazon.com/cloud-wan .
AWS Cloud9	AWS Cloud9 is a cloud-based integrated development environment (IDE) that you use to write, run, and debug code. See Also https://aws.amazon.com/cloud9/ .
CloudFormation	AWS CloudFormation is a service for writing or changing templates that create and delete related AWS resources together as a unit. See Also https://aws.amazon.com/cloudformation .
CloudFront	Amazon CloudFront is an AWS content delivery service that helps you improve the performance, reliability, and availability of your websites and applications. See Also https://aws.amazon.com/cloudfront .
CloudHSM	AWS CloudHSM is a web service that helps you meet corporate, contractual, and regulatory compliance requirements for data security by using dedicated hardware security module (HSM) appliances within the AWS Cloud. See Also https://aws.amazon.com/cloudhsm/ .
CloudSearch	Amazon CloudSearch is a fully managed service in the AWS Cloud that you can use to set up, manage, and scale a search solution for your website or application.

See Also <https://aws.amazon.com/cloudsearch/>.

CloudTrail

AWS CloudTrail is a web service that records AWS API calls for your account and delivers log files to you. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements that the AWS service returns.

See Also <https://aws.amazon.com/cloudtrail/>.

CloudWatch

Amazon CloudWatch is a web service that you can use to monitor and manage various metrics, and configure alarm actions based on data from those metrics.

See Also <https://aws.amazon.com/cloudwatch>.

CloudWatch Events

Amazon CloudWatch Events is a web service that you can use to deliver a timely stream of system events that describe changes in AWS [resources](#) to [Lambda](#) functions, streams in [Kinesis Data Streams](#), [Amazon SNS](#) topics, or built-in targets.

See Also <https://aws.amazon.com/cloudwatch>.

CloudWatch Logs

Amazon CloudWatch Logs is a web service for monitoring and troubleshooting your systems and applications from your existing system, application, and custom log files. You can send your existing log files to CloudWatch Logs and monitor these logs in near-real time.

See Also <https://aws.amazon.com/cloudwatch>.

cluster

A logical grouping of [container instances](#) that you can place [tasks](#) on.

[OpenSearch Service](#): A logical grouping of one or more data nodes, optional dedicated master nodes, and storage required to run Amazon OpenSearch Service (OpenSearch Service) and operate your OpenSearch Service domain.

See Also [data node](#), [dedicated master node](#), [node](#).

cluster compute instance

A type of [instance](#) that provides a great amount of CPU power coupled with increased networking performance, making it well suited for High Performance Compute (HPC) applications and other demanding network-bound applications.

cluster placement group	A logical cluster compute instance grouping to provide lower latency and high-bandwidth connectivity between the instances .
cluster status	OpenSearch Service : An indicator of the health of a cluster. A status can be green, yellow, or red. At the shard level, green means that all shards are allocated to nodes in a cluster, yellow means that the primary shard is allocated but the replica shards aren't, and red means that the primary and replica shards of at least one index aren't allocated. The shard status determines the index status, and the index status determines the cluster status.
CNAME	Canonical Name Record. A type of resource record in the Domain Name System (DNS) that specifies that the domain name is an alias of another, canonical domain name. Specifically, it's an entry in a DNS table that you can use to alias one fully qualified domain name to another.
Code Signing for AWS IoT	A service for signing code that you create for any IoT device that's supported by Amazon Web Services (AWS).
CodeBuild	AWS CodeBuild is a fully managed continuous integration service that compiles source code, runs tests, and produces software packages that are ready to deploy. See Also https://aws.amazon.com/codebuild .
CodeCommit	AWS CodeCommit is a fully managed source control service that companies can use to host secure and highly scalable private Git repositories. See Also https://aws.amazon.com/codecommit .
CodeDeploy	AWS CodeDeploy is a service that automates code deployments to any instance, including EC2 instances and instances running on-premises. See Also https://aws.amazon.com/codedeploy .
AWS CodeDeploy agent	AWS CodeDeploy agent is a software package that, when installed and configured on an instance, enables that instance to be used in CodeDeploy deployments.
CodeGuru	Amazon CodeGuru is a collection of developer tools that automate code reviews and provide intelligent recommendations to optimize application performance.

	See Also https://aws.amazon.com/codeguru .
CodePipeline	<p>AWS CodePipeline is a continuous delivery service for fast and reliable application updates.</p> <p>See Also https://aws.amazon.com/codepipeline.</p>
Amazon Cognito	<p>Amazon Cognito is a web service that you can use to save mobile user data in the AWS Cloud without writing any backend code or managing any infrastructure. Examples of mobile user data that you can save include app preferences and game states. Amazon Cognito offers mobile identity management and data synchronization across devices.</p> <p>See Also https://aws.amazon.com/cognito/.</p>
collaboration	<p>AWS Clean Rooms: A secure logical boundary in AWS Clean Rooms in which members can perform SQL queries on configured tables.</p>
AWS CLI	<p>AWS Command Line Interface is a unified downloadable and configurable tool for managing AWS services. Control multiple AWS services from the command line and automate them through scripts.</p> <p>See Also https://aws.amazon.com/cli/.</p>
complaint	<p>The event where a recipient who doesn't want to receive an email message chooses "Mark as Spam" within the email client, and the internet service provider (ISP) sends a notification to Amazon SES.</p>
compound query	<p>CloudSearch: A search request that specifies multiple search criteria using the Amazon CloudSearch structured search syntax.</p>
Amazon Comprehend	<p>Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text.</p> <p>See Also https://aws.amazon.com/comprehend/.</p>
Amazon Comprehend Medical	<p>Amazon Comprehend Medical is a HIPAA-eligible natural language processing (NLP) service that uses machine learning (ML), and has been pre-trained to understand and extract health data from medical text, such as prescriptions, procedures, or diagnoses.</p> <p>See Also https://aws.amazon.com/comprehend/medical.</p>
condition	<p>IAM: Any restriction or detail about a permission. The condition is <i>D</i> in the statement "A has permission to do B to C where D applies."</p>

	<p>AWS WAF: A set of attributes that AWS WAF searches for in web requests to AWS resources such as Amazon CloudFront distributions. Conditions can include values such as the IP addresses that web requests originate from or values in request headers. Based on the specified conditions, you can configure AWS WAF to allow or block web requests to AWS resources.</p>
conditional parameter	See mapping .
AWS Config	<p>AWS Config is a fully managed service that provides an AWS resource inventory, configuration history, and configuration change notifications for better security and governance. You can create rules that automatically check the configuration of AWS resources that AWS Config records.</p> <p>See Also https://aws.amazon.com/config/.</p>
configuration API	<p>CloudSearch: The API call that you use to create, configure, and manage search domains.</p>
configuration template	<p>A series of key–value pairs that define parameters for various AWS products so that Elastic Beanstalk can provision them for an environment.</p>
Amazon Connect	<p>Amazon Connect is a service solution that offers self-service configuration and provides dynamic, personal, and natural customer engagement at any scale.</p> <p>See Also https://aws.amazon.com/connect/.</p>
consistency model	<p>The method a service uses to achieve high availability. For example, it could involve replicating data across multiple servers in a data center.</p> <p>See Also eventual consistency.</p>
console	See AWS Management Console .
Console Mobile Application	<p>AWS Console Mobile Application lets AWS customers monitor and manage a select set of resources to stay informed and connected with their AWS resources while on the go.</p> <p>See Also https://aws.amazon.com/console/mobile.</p>

consolidated billing	A feature of the AWS Organizations service for consolidating payment for multiple AWS accounts. You create an organization that contains your AWS accounts, and you use the management account of your organization to pay for all member accounts. You can see a combined view of AWS costs that are incurred by all accounts in your organization, and you can get detailed cost reports for individual accounts.
container	A container is a standard unit of software that contains application code and all relevant dependencies.
container definition	A container definition specifies the details that are associated with running a container on Amazon ECS. More specifically, a container definition specifies details such as the container image to use and how much CPU and memory the container is allocated. The container definition is included as part of an Amazon ECS task definition .
container instance	A container instance is a self-managed EC2 instance or an on-premises server or virtual machine (VM) that's running the Amazon Elastic Container Service (Amazon ECS) container agent and has been registered into a cluster . A container instance serves as the infrastructure that your Amazon ECS workloads are run on.
container registry	A container registry is a collection of repositories that store container images. One example is Amazon Elastic Container Registry (Amazon ECR).
content delivery network (CDN)	A web service that speeds up distribution of your static and dynamic web content—such as .html, .css, .js, media files, and image files—to your users by using a worldwide network of data centers. When a user requests your content, the request is routed to the data center that provides the lowest latency (time delay). If the content is already in the location with the lowest latency, the CDN delivers it immediately. If not, the CDN retrieves it from an origin that you specify (for example, a web server or an Amazon S3 bucket). With some CDNs, you can help secure your content by configuring an HTTPS connection between users and data centers, and between data centers and your origin. Amazon CloudFront is an example of a CDN.

contextual metadata	<p>Amazon Personalize: Interactions data that you collect about a user's browsing context (such as device used or location) when an event (such as a click) occurs. Contextual metadata can improve recommendation relevance for new and existing users.</p> <p>See Also Interactions dataset, event.</p>
continuous delivery	<p>A software development practice where code changes are automatically built, tested, and prepared for a release to production.</p> <p>See Also https://aws.amazon.com/devops/continuous-delivery/.</p>
continuous integration	<p>A software development practice where developers regularly merge code changes into a central repository, after which automated builds and tests are run.</p> <p>See Also https://aws.amazon.com/devops/continuous-integration/.</p>
AWS Control Tower	<p>AWS Control Tower is a service used to set up and govern a secure, multi-account AWS environment.</p> <p>See Also https://aws.amazon.com/controltower.</p>
cooldown period	<p>Amount of time that Amazon EC2 Auto Scaling doesn't allow the desired size of the Auto Scaling group to be changed by any other notification from an CloudWatch alarm.</p>
core node	<p>An EC2 instance that runs Hadoop map and reduce tasks and stores data using the Hadoop Distributed File System (HDFS). Core nodes are managed by the master node, which assigns Hadoop tasks to nodes and monitors their status. The EC2 instances you assign as core nodes are capacity that must be allotted for the entire job flow run. Because core nodes store data, you can't remove them from a job flow. However, you can add more core nodes to a running job flow.</p> <p>Core nodes run both the DataNodes and TaskTracker Hadoop daemons.</p>
corpus	<p>CloudSearch: A collection of data that you want to search.</p>
Corretto	<p>Amazon Corretto is a no-cost, multiplatform, production-ready distribution of the Open Java Development Kit (OpenJDK).</p> <p>See Also https://aws.amazon.com/corretto/.</p>
coverage	<p>Amazon Personalize: An evaluation metric that tells you the proportion of unique items that Amazon Personalize might recommend using your</p>

	<p>model out of the total number of unique items in Interactions and Items datasets. To make sure Amazon Personalize recommends more of your items, use a model with a higher coverage score. Recipes that feature item exploration, such as user-personalization, have higher coverage than those that don't, such as popularity-count.</p> <p>See Also metrics, Items dataset, Interactions dataset, item exploration, user-personalization recipe, popularity-count recipe.</p>
credential helper	<p>CodeCommit: A program that stores credentials for repositories and supplies them to Git when making connections to those repositories. The AWS CLI includes a credential helper that you can use with Git when connecting to CodeCommit repositories.</p>
credentials	<p>Also called <i>access credentials</i> or <i>security credentials</i>. In authentication and authorization, a system uses credentials to identify who is making a call and whether to allow the requested access. In AWS, these credentials are typically the access key ID and the secret access key.</p>
cross-account access	<p>The process of permitting limited, controlled use of resources in one AWS account by a user in another AWS account. For example, in CodeCommit and CodeDeploy you can configure cross-account access so that a user in AWS account A can access an CodeCommit repository created by account B. Or a pipeline in CodePipeline created by account A can use CodeDeploy resources created by account B. In IAM you use a role to delegate temporary access to a user in one account to resources in another.</p>
cross-Region replication	<p>A solution for replicating data across different AWS Regions, in near-real time.</p>
Cryptographic Computing for Clean Rooms (C3R)	<p>AWS Clean Rooms: A capability in AWS Clean Rooms that organizations can use to bring sensitive data together to derive new insights from data analytics while cryptographically limiting what any party in the process can learn.</p>
customer gateway	<p>A router or software application on your side of a VPN tunnel that's managed by Amazon VPC. The internal interfaces of the customer gateway are attached to one or more devices in your home network.</p>

	The external interface is attached to the virtual private gateway (VGW) across the VPN tunnel.
customer managed policy	An IAM managed policy that you create and manage in your AWS account .
customer master key (CMK)	We no longer use customer master key or CMK. These terms are replaced by AWS KMS key (first mention) and KMS key (subsequent mention). For more information, see KMS key .

D

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

dashboard	See service health dashboard .
data consistency	A concept that describes when data is written or updated successfully and all copies of the data are updated in all AWS Regions . However, it takes time for the data to propagate to all storage locations. To support varied application requirements, DynamoDB supports both eventually consistent and strongly consistent reads. See Also eventual consistency , eventually consistent read , strongly consistent read .
AWS Data Exchange	AWS Data Exchange is a service that helps you find, subscribe to, and use third-party data in the cloud. See Also https://aws.amazon.com/data-exchange .
Amazon Data Lifecycle Manager	Amazon Data Lifecycle Manager is an Amazon service that automates and manages the lifecycle of Amazon EBS snapshots and Amazon EBS-backed AMIs .
data node	OpenSearch Service : An OpenSearch instance that holds data and responds to data upload requests. See Also dedicated master node , node .
Data Pipeline	AWS Data Pipeline is a web service for processing and moving data between different AWS compute and storage services, as well as on-premises data sources, at specified intervals.

See Also <https://aws.amazon.com/datapipeline>.

data schema

See [schema](#).

data source

The database, file, or repository that provides information required by an application or database. For example, in [OpsWorks](#), valid data sources include an [instance](#) for a stack's MySQL layer or a stack's [Amazon RDS](#) service layer. In [Amazon Redshift](#), valid data sources include text files in an [Amazon S3 bucket](#), in an [Amazon EMR](#) cluster, or on a remote host that a cluster can access through an SSH connection.

See Also [datasource](#).

database engine

The database software and version running on the [DB instance](#).

database name

The name of a database hosted in a [DB instance](#). A DB instance can host multiple databases, but databases hosted by the same DB instance must each have a unique name within that instance.

dataset

[Amazon Personalize](#): A container for the data used by Amazon Personalize. There are three types of Amazon Personalize datasets: Users, Items, and Interactions.

See Also [Interactions dataset](#), [Users dataset](#), [Items dataset](#).

dataset group

[Amazon Personalize](#): A container for Amazon Personalize components, including datasets, event trackers, solutions, filters, campaigns, and batch inference jobs. A dataset group organizes your resources into independent collections, so resources from one dataset group can't influence resources in any other dataset group.

See Also [dataset](#), [event tracker](#), [solution](#), [campaign](#).

datasource

[Amazon ML](#): An object that contains metadata about the input data. Amazon ML reads the input data, computes descriptive statistics on its attributes, and stores the statistics—along with a schema and other information—as part of the datasource object. Amazon ML uses datasources to train and evaluate a machine learning model and generate batch predictions.

See Also [data source](#).

DataSync	<p>AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between storage systems and services.</p> <p>See Also https://aws.amazon.com/datasync.</p>
DB compute class	<p>The size of the database compute platform used to run the instance.</p>
DB instance	<p>An isolated database environment running in the cloud. A DB instance can contain multiple user-created databases.</p>
DB instance identifier	<p>User-supplied identifier for the DB instance. The identifier must be unique for that user in an AWS Region.</p>
DB parameter group	<p>A container for database engine parameter values that apply to one or more DB instances.</p>
DB security group	<p>A method that controls access to the DB instance. By default, network access is turned off to DB instances. After inbound traffic is configured for a security group, the same rules apply to all DB instances associated with that group.</p>
DB snapshot	<p>A user-initiated point backup of a DB instance.</p>
Dedicated Host	<p>A physical server with EC2 instance capacity fully dedicated to a user.</p>
Dedicated Instance	<p>An instance that's physically isolated at the host hardware level and launched within a Amazon VPC.</p>
dedicated master node	<p>OpenSearch Service: An OpenSearch instance that performs cluster management tasks, but doesn't hold data or respond to data upload requests. Amazon OpenSearch Service (OpenSearch Service) uses dedicated master nodes to increase cluster stability.</p> <p>See Also data node, node.</p>
Dedicated Reserved Instance	<p>An option that you purchase to guarantee that sufficient capacity will be available to launch Dedicated Instances into a Amazon VPC.</p>
AWS DeepComposer	<p>AWS DeepComposer is a web service designed specifically to educate developers through tutorials, sample code, and training data.</p> <p>See Also https://aws.amazon.com/deepcomposer.</p>

AWS DeepLens	AWS DeepLens is a tool that provides AWS customers with a centralized place to search, discover, and connect with trusted APN Technology and Consulting Partners, based on customers' business needs. See Also https://aws.amazon.com/deeplens .
AWS DeepRacer	AWS DeepRacer is a cloud-based 3D racing simulator, global racing league, and fully autonomous 1/18th-scale race car driven by reinforcement learning. See Also https://aws.amazon.com/deepracer .
delegation	Within a single AWS account : Giving AWS users access to resources your AWS account. Between two AWS accounts: Setting up a trust between the account that owns the resource (the trusting account), and the account that contains the users that need to access the resource (the trusted account). See Also trust policy .
delete marker	An object with a key and version ID, but without content. Amazon S3 inserts delete markers automatically into versioned buckets when an object is deleted.
deliverability	The likelihood that an email message arrives at its intended destination.
deliveries	The number of email messages, sent through Amazon SES , that were accepted by an internet service provider (ISP) for delivery to recipients over a period of time.
deny	The result of a policy statement that includes deny as the effect, so that a specific action or actions are expressly forbidden for a user, group, or role. Explicit deny take precedence over explicit allow .
deployment configuration	CodeDeploy : A set of deployment rules and success and failure conditions used by the service during a deployment.
deployment group	CodeDeploy : A set of individually tagged instances or EC2 instances in Auto Scaling groups , or both.
Description property	A property added to parameters, resources , resource properties, mappings, and outputs to help you to document CloudFormation template elements.

detailed monitoring	Monitoring of AWS provided metrics derived at a 1-minute frequency.
Detective	Amazon Detective is a service that collects log data from your AWS resources to analyze and identify the root cause of security findings or suspicious activities. The Detective behavior graph provides visualizations to help you to determine the nature and extent of possible security issues and conduct an efficient investigation. See Also https://aws.amazon.com/detective/ .
Device Farm	AWS Device Farm is an app testing service that you can use to test Android, iOS, and web apps on real, physical phones and tablets that are hosted by AWS. See Also https://aws.amazon.com/device-farm/ .
Amazon DevOps Guru	Amazon DevOps Guru is a fully managed operations service powered by machine learning (ML), designed to improve an application's operational performance and availability. See Also https://aws.amazon.com/devops-guru/ .
dimension	A name–value pair (for example, InstanceType=m1.small, or EngineName=mysql), that contains additional information to identify a metric.
Direct Connect	AWS Direct Connect is a web service that simplifies establishing a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data center, office, or colocation environment. See Also https://aws.amazon.com/directconnect .
Directory Service	AWS Directory Service is a managed service for connecting your AWS resources to an existing on-premises Microsoft Active Directory or to set up and operate a new, standalone directory in the AWS Cloud. See Also https://aws.amazon.com/directoryservice .
discussion forums	A place where AWS users can post technical questions and feedback to help accelerate their development efforts and to engage with the AWS community. For more information, see the Amazon Web Services Discussion Forums .

distribution	A link between an origin server (such as an Amazon S3 bucket) and a domain name, which CloudFront automatically assigns. Through this link, CloudFront identifies the object you have stored in your origin server .
DKIM	DomainKeys Identified Mail is a standard that email senders use to sign their messages. ISPs use those signatures to verify that messages are legitimate. For more information, see https://tools.ietf.org/html/rfc6376 .
AWS DMS	AWS Database Migration Service is a web service that can help you migrate data to and from many widely used commercial and open-source databases. See Also https://aws.amazon.com/dms .
DNS	See Domain Name System .
Docker image	A layered file system template that's the basis of a Docker container . Docker images can comprise specific operating systems or applications.
document	CloudSearch : An item that can be returned as a search result. Each document has a collection of fields that contain the data that can be searched or returned. The value of a field can be either a string or a number. Each document must have a unique ID and at least one field.
document batch	CloudSearch : A collection of add and delete document operations. You use the document service API to submit batches to update the data in your search domain.
document service API	CloudSearch : The API call that you use to submit document batches to update the data in a search domain.
document service endpoint	CloudSearch : The URL that you connect to when sending document updates to an Amazon CloudSearch domain. Each search domain has a unique document service endpoint that remains the same for the life of the domain.
Amazon DocumentDB	Amazon DocumentDB (with MongoDB compatibility) is a managed database service that you can use to set up, operate, and scale MongoDB-compatible databases in the cloud.

See Also <https://aws.amazon.com/documentdb/>.

domain

[OpenSearch Service](#): The hardware, software, and data exposed by Amazon OpenSearch Service (OpenSearch Service) endpoints. An OpenSearch Service domain is a service wrapper around an OpenSearch cluster. An OpenSearch Service domain encapsulates the engine instances that process OpenSearch Service requests, the indexed data that you want to search, snapshots of the domain, access policies, and metadata.

See Also [cluster](#), [Elasticsearch](#).

Domain Name System

Domain Name System is a service that routes internet traffic to websites by translating human-readable domain names (for example, `www.example.com`) into the numeric IP addresses, such as `192.0.2.1`, which computers use to connect to each other.

Donation button

An HTML-coded button to provide a simple and secure way for US-based, IRS-certified 501(c)(3) nonprofit organizations to solicit donations.

DynamoDB

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. See Also <https://aws.amazon.com/dynamodb/>.

Amazon DynamoDB Encryption Client

Amazon DynamoDB Encryption Client is a software library that helps you protect your table data before you send it to [DynamoDB](#).

Amazon DynamoDB Storage Backend for Titan

Amazon DynamoDB Storage Backend for Titan is a graph database implemented on top of Amazon DynamoDB. Titan is a scalable graph database optimized for storing and querying graphs. See Also <https://aws.amazon.com/dynamodb/>.

DynamoDB Streams

Amazon DynamoDB Streams is an AWS service that captures a time-ordered sequence of item-level modifications in any Amazon DynamoDB table. This service also stores this information in a log for up to 24 hours. Applications can access this log and view the data items as they appeared before and after they were modified, in near-real time. See Also <https://aws.amazon.com/dynamodb/>.

E

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

Amazon EBS	Amazon Elastic Block Store is a service that provides block level storage volumes or use with EC2 instances . See Also https://aws.amazon.com/ebs .
Amazon EBS-backed AMI	An Amazon EBS-backed AMI is a type of Amazon Machine Image (AMI) whose instances use an Amazon EBS volume as their root device. Compare this with instances launched from instance store-backed AMIs , which use the instance store as the root device.
Amazon EC2	Amazon Elastic Compute Cloud is a web service for launching and managing Linux/UNIX and Windows Server instances in Amazon data centers. See Also https://aws.amazon.com/ec2 .
Amazon EC2 Auto Scaling	Amazon EC2 Auto Scaling is a web service that launches or terminates instances automatically based on user-defined policies , schedules, and health checks . See Also https://aws.amazon.com/ec2/autoscaling .
EC2 instance	A compute instance in the Amazon EC2 service. Other AWS services use the term <i>EC2 instance</i> to distinguish these instances from other types of instances they support.
Amazon ECR	Amazon Elastic Container Registry (Amazon ECR) is a fully managed Docker container registry that you can use to store, manage, and deploy Docker container images. Amazon ECR is integrated with Amazon ECS and IAM . See Also https://aws.amazon.com/ecr .
Amazon ECS	Amazon Elastic Container Service (Amazon ECS) is a highly scalable, fast, container management service that you can use to run, stop, and manage Docker containers on a cluster of EC2 instances. See Also https://aws.amazon.com/ecs .

edge location	edge location is a data center that an AWS service uses to perform service-specific operations. For example, CloudFront uses edge locations to cache copies of your content, so the content is closer to your users and can be delivered faster regardless of their location. Route 53 uses edge locations to speed up the response to public DNS queries.
Amazon EFS	Amazon Elastic File System is a file storage service for EC2 instances . Amazon EFS provides an interface that you can use to create and configure file systems. Amazon EFS storage capacity grows and shrinks automatically as you add and remove files. See Also https://aws.amazon.com/efs/ .
Amazon EKS	Amazon Elastic Kubernetes Service is a managed service that you can use to run Kubernetes on AWS without needing to stand up or maintain your own Kubernetes control plane. See Also https://aws.amazon.com/eks/ .
Elastic	A company that provides open-source solutions—including OpenSearch, Logstash, Kibana, and Beats—that take data from any source and search, analyze, and visualize it in real time. Amazon OpenSearch Service (OpenSearch Service) is an AWS managed service for deploying, operating, and scaling OpenSearch in the AWS Cloud. See Also OpenSearch Service , Elasticsearch .
Elastic Beanstalk	AWS Elastic Beanstalk is a web service for deploying and managing applications in the AWS Cloud without worrying about the infrastructure that runs those applications. See Also https://aws.amazon.com/elasticbeanstalk .
Elastic Block Store	See Amazon EBS .
Elastic Inference	Amazon Elastic Inference is a resource that customers can use to attach low-cost GPU-powered acceleration to Amazon EC2 and SageMaker AI instances, or Amazon ECS tasks, to reduce the cost of running deep learning inference by up to 75%. See Also https://aws.amazon.com/machine-learning/elastic-inference .

Elastic IP address	<p>A fixed (static) IP address that you have allocated in Amazon EC2 or Amazon VPC and then attached to an instance. Elastic IP addresses are associated with your account, not a specific instance. They are <i>elastic</i> because you can easily allocate, attach, detach, and free them as your needs change. Unlike traditional static IP addresses, Elastic IP addresses allow you to mask instance or Availability Zone failures by rapidly remapping your public IP addresses to another instance.</p>
ELB	<p>Elastic Load Balancing is a web service that improves an application's availability by distributing incoming traffic between two or more EC2 instances.</p> <p>See Also https://aws.amazon.com/elasticloadbalancing.</p>
elastic network interface	<p>An additional network interface that can be attached to an instance. Elastic network interfaces include a primary private IP address, one or more secondary private IP addresses, an Elastic IP Address (optional), a MAC address, membership in specified security groups, a description, and a source/destination check flag. You can create an elastic network interface, attach it to an instance, detach it from an instance, and attach it to another instance.</p>
Elastic Transcoder	<p>Amazon Elastic Transcoder is a cloud-based media transcoding service. Elastic Transcoder is a highly scalable tool for converting (or <i>transcoding</i>) media files from their source format into versions that play on devices such as smartphones, tablets, and PCs.</p> <p>See Also https://aws.amazon.com/elastictranscoder/.</p>
ElastiCache	<p>Amazon ElastiCache is a web service that simplifies deploying, operating, and scaling an in-memory cache in the cloud. The service improves the performance of web applications by providing information retrieval from fast, managed, in-memory caches, instead of relying entirely on slower disk-based databases.</p> <p>See Also https://aws.amazon.com/elasticache/.</p>
Elasticsearch	<p>An open-source, real-time distributed search and analytics engine used for full-text search, structured search, and analytics. OpenSearch was developed by the Elastic company.</p>

	<p>Amazon OpenSearch Service (OpenSearch Service) is an AWS managed service for deploying, operating, and scaling OpenSearch in the AWS Cloud.</p> <p>See Also OpenSearch Service, Elastic.</p>
AWS Elemental MediaConnect	<p>AWS Elemental MediaConnect is a fully-managed live video distribution service that reliably and securely ingests video into the AWS Cloud and transports it to multiple destinations within the AWS network and the internet.</p> <p>See Also https://aws.amazon.com/mediacconnect.</p>
AWS Elemental MediaConvert	<p>AWS Elemental MediaConvert is a file-based media conversion service that transforms content into formats for traditional broadcast and internet streaming.</p> <p>See Also https://aws.amazon.com/mediaconvert.</p>
AWS Elemental MediaLive	<p>AWS Elemental MediaLive is a cloud-based live video encoding service that creates high-quality streams for delivery to broadcasts and internet-connected devices.</p> <p>See Also https://aws.amazon.com/medialive.</p>
AWS Elemental MediaPackage	<p>AWS Elemental MediaPackage is a highly-scalable video origination and packaging service that delivers video securely and reliably.</p> <p>See Also https://aws.amazon.com/mediapackage.</p>
AWS Elemental MediaStore	<p>AWS Elemental MediaStore is a storage service optimized for media that provides the performance, consistency, and low latency required to deliver live and on-demand video content at scale.</p> <p>See Also https://aws.amazon.com/mediastore.</p>
AWS Elemental MediaTailor	<p>AWS Elemental MediaTailor is a channel assembly and personalized ad-insertion service for over-the-top (OTT) video and audio applications.</p> <p>See Also https://aws.amazon.com/mediatailor.</p>
EMP	<p>The AWS End-of-Support Migration Program for Windows Server provides the technology and guidance to migrate your applications running on Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 to the latest, supported versions of Windows Server running on Amazon Web Services (AWS).</p>

Amazon EMR	<p>Amazon Elastic Map Reduce is a web service that you can use to process large amounts of data efficiently. Amazon EMR uses Hadoop processing combined with several AWS products to do such tasks as web indexing, data mining, log file analysis, machine learning, scientific simulation, and data warehousing.</p> <p>See Also https://aws.amazon.com/elasticmapreduce.</p>
encrypt	<p>To use a mathematical algorithm to make data unintelligible to unauthorized users. Encryption also gives authorized users a method (such as a key or password) to convert the altered data back to its original state.</p>
encryption context	<p>A set of key–value pairs that contains additional information associated with AWS KMS–encrypted information.</p>
AWS Encryption SDK	<p>AWS Encryption SDK is a client-side encryption library that you can use to encrypt and decrypt data using industry standards and best practices. See Also https://aws.amazon.com/blogs/security/tag/aws-encryption-sdk/.</p>
endpoint	<p>A URL that identifies a host and port as the entry point for a web service. Every web service request contains an endpoint. Most AWS products provide endpoints for a Region to enable faster connectivity.</p> <p>ElastiCache: The DNS name of a cache node.</p> <p>Amazon RDS: The DNS name of a DB instance.</p> <p>CloudFormation: The DNS name or IP address of the server that receives an HTTP request.</p>
endpoint port	<p>ElastiCache: The port number used by a cache node.</p> <p>Amazon RDS: The port number used by a DB instance.</p>
envelope encryption	<p>The use of a master key and a data key to algorithmically protect data. The master key is used to encrypt and decrypt the data key and the data key is used to encrypt and decrypt the data itself.</p>
environment	<p>Elastic Beanstalk: A specific running instance of an application. The application has a CNAME and includes an application version and</p>

	<p>a customizable configuration (which is inherited from the default container type).</p> <p>CodeDeploy: Instances in a deployment group in a blue/green deployment. At the start of a blue/green deployment, the deployment group is made up of instances in the original environment. At the end of the deployment, the deployment group is made up of instances in the replacement environment.</p>
environment configuration	A collection of parameters and settings that define how an environment and its associated resources behave.
ephemeral store	See instance store .
epoch	The date from which time is measured. For most Unix environments, the epoch is January 1, 1970.
ETL	See extract, transform, and load (ETL) .
evaluation	<p>Amazon Machine Learning: The process of measuring the predictive performance of a machine learning (ML) model.</p> <p>Also a machine learning object that stores the details and result of an ML model evaluation.</p>
evaluation datasource	The data that Amazon Machine Learning uses to evaluate the predictive accuracy of a machine learning model.
event	<p>Amazon Personalize: A user activity—such as a click, a purchase, or a video viewing—that you record and upload to an Amazon Personalize Interactions dataset. You record events individually in real time or record and upload events in bulk.</p> <p>See Also dataset, Interactions dataset.</p>
event tracker	<p>Amazon Personalize: Specifies a destination dataset group for event data that you record in real time. When you record events in real time, you provide the ID of the event tracker so that Amazon Personalize knows where to add the data.</p> <p>See Also dataset group, event.</p>

EventBridge	<p>Amazon EventBridge is a serverless event bus service that you can use to connect your applications with data from a variety of sources and routes that data to targets such as AWS Lambda. You can set up routing rules to determine where to send your data to build application architectures that react in real time to all of your data sources.</p> <p>See Also https://aws.amazon.com/eventbridge/.</p>
eventual consistency	<p>The method that AWS services use to achieve high availability. This involves replicating data across multiple servers in Amazon data centers. When data is written or updated and Success is returned, all copies of the data are updated. However, it takes time for the data to propagate to all storage locations. The data will eventually be consistent, but an immediate read might not show the change. Consistency is usually reached within seconds.</p> <p>See Also data consistency, eventually consistent read, strongly consistent read.</p>
eventually consistent read	<p>A read process that returns data from only one Region and might not show the most recent write information. However, if you repeat your read request after a short time, the response should eventually return the latest data.</p> <p>See Also data consistency, eventual consistency, strongly consistent read.</p>
eviction	<p>The deletion by CloudFront of an object from an edge location before its expiration time. If an object in an edge location isn't frequently requested, CloudFront might evict the object (remove the object before its expiration date) to make room for objects that are more popular.</p>
exbibyte (EiB)	<p>A contraction of exa binary byte. An exbibyte (EiB) is 2^{60} or 1,152,921,504,606,846,976 bytes. An exabyte (EB) is 10^{18} or 1,000,000,000,000,000,000 bytes. 1,024 EiB is a zebibyte (ZiB).</p>
expiration	<p>For CloudFront caching, the time when CloudFront stops responding to user requests with an object. If you don't use headers or CloudFront distribution settings to specify how long you want objects to stay in an edge location, the objects expire after 24 hours. The next time a user requests an object that has expired, CloudFront forwards the request to the origin.</p>

explicit impressions	<p>Amazon Personalize: A list of items that you manually add to an Amazon Personalize Interactions dataset to influence future recommendations. Unlike <i>implicit impressions</i>, where Amazon Personalize automatically derives the impressions data, you choose what to include in explicit impressions.</p> <p>See Also recommendations, Interactions dataset, impressions data, implicit impressions.</p>
explicit launch permission	<p>An Amazon Machine Image (AMI) launch permission granted to a specific AWS account.</p>
exponential backoff	<p>A strategy that incrementally increases the wait between retry attempts in order to reduce the load on the system and increase the likelihood that repeated requests will succeed. For example, client applications might wait up to 400 milliseconds before attempting the first retry, up to 1600 milliseconds before the second, and up to 6400 milliseconds (6.4 seconds) before the third.</p>
expression	<p>CloudSearch: A numeric expression that you can use to control how search hits are sorted. You can construct Amazon CloudSearch expressions using numeric fields, other rank expressions, a document's default relevance score, and standard numeric operators and functions. When you use the <code>sort</code> option to specify an expression in a search request, the expression is evaluated for each search hit and the hits are listed according to their expression values.</p>
extract, transform, and load (ETL)	<p>A process that's used to integrate data from multiple sources. Data is collected from sources (extract), converted to an appropriate format (transform), and written to a target data store (load) for purposes of analysis and querying.</p> <p>ETL tools combine these three functions to consolidate and move data from one environment to another. AWS Glue is a fully managed ETL service for discovering and organizing data, transforming it, and making it available for search and analytics.</p>

F

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

facet	CloudSearch : An index field that represents a category that you want to use to refine and filter search results.
facet enabled	CloudSearch : An index field option that enables facet information to be calculated for the field.
AWS Fargate	AWS Fargate is a serverless, pay-as-you-go compute engine that you can use to build applications on AWS. You can use Amazon Elastic Container Service (Amazon ECS) or Amazon Elastic Kubernetes Service (Amazon EKS) to maintain container applications using AWS Fargate. See Also https://aws.amazon.com/fargate/ .
Fault Injection Simulator (AWS FIS)	AWS Fault Injection Service is a managed service that you can use to perform fault injection experiments on your AWS workloads. See Also https://aws.amazon.com/fis .
FBL	See feedback loop (FBL) .
feature transformation	Amazon Machine Learning: The machine learning process of constructing more predictive input representations or “features” from the raw input variables to optimize a machine learning model’s ability to learn and generalize. Also known as <i>data transformation</i> or <i>feature engineering</i> .
federated identity management (FIM)	Allows individuals to sign in to different networks or services, using the same group or personal credentials to access data across all networks. With identity federation in AWS, external identities (federated users) are granted secure access to resources in an AWS account without having to create IAM users . These external identities can come from a corporate identity store (such as LDAP or Windows Active Directory) or from a third party (such as Login with Amazon, Facebook, or Google). AWS federation also supports SAML 2.0.
federated user	See federated identity management (FIM) .

federation	See federated identity management (FIM) .
feedback loop (FBL)	The mechanism by which a mailbox provider (for example, an internet service provider (ISP)) forwards a recipient's complaint back to the sender .
field weight	The relative importance of a text field in a search index. Field weights control how much matches in particular text fields affect a document's relevance score.
filter	A criterion that you specify to limit the results when you list or describe your Amazon EC2 resources .
filter query	A way to filter search results without affecting how the results are scored and sorted. Specified with the CloudSearch fq parameter.
FIM	See federated identity management (FIM) .
FinSpace	Amazon FinSpace is a data management and analytics service purpose-built for the financial services industry (FSI). See Also https://aws.amazon.com/finspace .
Firehose	See Firehose .
Firewall Manager	AWS Firewall Manager is a service that you use with AWS WAF to simplify your AWS WAF administration and maintenance tasks across multiple accounts and resources. With AWS Firewall Manager, you set up your firewall rules only once. The service automatically applies your rules across your accounts and resources, even as you add new resources. See Also https://aws.amazon.com/firewall-manager .
Forecast	Amazon Forecast is a fully managed service that uses statistical and machine learning algorithms to produce highly accurate time-series forecasts. See Also https://aws.amazon.com/forecast/ .
format version	See template format version .
forums	See discussion forums .
function	See intrinsic function .

fuzzy search A simple search query that uses approximate string matching (fuzzy matching) to correct for typographical errors and misspellings.

G

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

GameKit AWS GameKit is an open-source SDK and game engine plugin that empowers game developers to build and deploy cloud-based features with AWS from their game engine.

See Also <https://aws.amazon.com/gamekit/>.

Amazon GameLift Servers Amazon GameLift Servers is a managed service for deploying, operating, and scaling session-based multiplayer games.

See Also <https://aws.amazon.com/gamelift/>.

GameSparks Amazon GameSparks is a fully managed AWS service that provides a multi-service backend for game developers.

See Also <https://aws.amazon.com/gamesparks/>.

geospatial search A search query that uses locations specified as a latitude and longitude to determine matches and sort the results.

gibibyte (GiB) A contraction of giga binary byte, a gibibyte is 2^{30} or 1,073,741,824 bytes. A gigabyte (GB) is 10^9 or 1,000,000,000 bytes. 1,024 GiB is a [tebibyte \(TiB\)](#).

GitHub A web-based repository that uses Git for version control.

Global Accelerator AWS Global Accelerator is a network layer service that you use to create accelerators that direct traffic to optimal endpoints over the AWS global network. This improves the availability and performance of your internet applications that are used by a global audience.

See Also <https://aws.amazon.com/global-accelerator>.

global consistency An [active-active](#) strategy in which all reads and writes for a workload are handled in the Region where the request originates and are replicated synchronously to all other Regions in the architecture.

See Also , .

global secondary index	<p>An index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered global because queries on the index can span all of the data in a table, across all partitions.</p> <p>See Also local secondary index.</p>
AWS Glue	<p>AWS Glue is a fully managed extract, transform, and load (ETL) service that you can use to catalog data and load it for analytics. With AWS Glue, you can discover your data, develop scripts to transform sources into targets, and schedule and run ETL jobs in a serverless environment.</p> <p>See Also https://aws.amazon.com/glue.</p>
AWS GovCloud (US)	<p>AWS GovCloud (US) is an isolated AWS Region that hosts sensitive workloads in the cloud, ensuring that this work meets the US government's regulatory and compliance requirements. The AWS GovCloud (US) Region adheres to United States International Traffic in Arms Regulations (ITAR), Federal Risk and Authorization Management Program (FedRAMP) requirements, Department of Defense (DOD) Cloud Security Requirements Guide (SRG) Levels 2 and 4, and Criminal Justice Information Services (CJIS) Security Policy requirements.</p> <p>See Also https://aws.amazon.com/govcloud-us/.</p>
grant	<p>AWS KMS: A mechanism for giving AWS principals long-term permissions to use KMS keys.</p>
grant token	<p>A type of identifier that allows the permissions in a grant to take effect immediately.</p>
ground truth	<p>The observations used in the machine learning (ML) model training process that include the correct value for the target attribute. To train an ML model to predict house sales prices, the input observations would typically include prices of previous house sales in the area. The sale prices of these houses constitute the ground truth.</p>
group	<p>A collection of IAM users. You can use IAM groups to simplify specifying and managing permissions for multiple users.</p>
GuardDuty	<p>Amazon GuardDuty is a continuous security monitoring service. Amazon GuardDuty can help to identify unexpected and potentially unauthorized or malicious activity in your AWS environment.</p>

See Also <https://aws.amazon.com/guardduty/>.

H

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

Hadoop	Software that enables distributed processing for big data by using clusters and simple programming models. For more information, see http://hadoop.apache.org .
hard bounce	A persistent email delivery failure such as "mailbox does not exist."
hardware VPN	A hardware-based IPsec VPN connection over the internet.
AWS Health	AWS Health is a service that provides ongoing visibility into AWS customers' accounts and the availability of their AWS services and resources. See Also https://aws.amazon.com/premiumsupport/technology/aws-health-dashboard .
health check	A system call to check on the health status of each instance in an Amazon EC2 Auto Scaling group.
HealthLake	AWS HealthLake is a HIPAA-eligible service that helps customers store, query, and generate artificial intelligence (AI) and machine learning (ML) insights from healthcare data and enables healthcare data interoperability. See Also https://aws.amazon.com/healthlake .
highlight enabled	CloudSearch : An index field option that enables matches within the field to be highlighted.
highlights	CloudSearch : Excerpts returned with search results that show where the search terms appear within the text of the matching documents.
high-quality email	Email that recipients find valuable and want to receive. Value means different things to different recipients and can come in such forms as offers, order confirmations, receipts, or newsletters.

hit	A document that matches the criteria specified in a search request. Also referred to as a <i>search result</i> .
HMAC	Hash-based Message Authentication Code is a specific construction for calculating a message authentication code (MAC) involving a cryptographic hash function in combination with a secret key. You can use it to verify both the data integrity and the authenticity of a message at the same time. AWS calculates the HMAC using a standard, cryptographic hash algorithm, such as SHA-256.
hosted zone	A collection of resource record sets that Route 53 hosts. Similar to a traditional DNS zone file, a hosted zone represents a collection of records that are managed together under a single domain name.
hot standby	An active-passive disaster recovery strategy in which a workload is fully scaled up in both the primary and standby Regions, but serves traffic from only the primary Region. See Also , , .
HRNN	Amazon Personalize : A hierarchical recurrent neural network machine learning algorithm that models changes in user behavior and predicts the items that a user might interact with in personal recommendation applications.
HTTP-Query	See Query .
HVM virtualization	Hardware Virtual Machine virtualization. Allows the guest VM to run as though it's on a native hardware platform, except that it still uses paravirtual (PV) network and storage drivers for improved performance. See Also PV virtualization .

I

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

IAM	AWS Identity and Access Management is a web service that Amazon Web Services (AWS) customers can use to manage users and user permissions within AWS.
-----	---

	See Also https://aws.amazon.com/iam .
IAM Access Analyzer	Access Management Access Analyzer is a feature of IAM that you can use to identify the resources in your organization and accounts that are shared with an external entity. Example resources include Amazon S3 buckets or IAM roles. See Also https://aws.amazon.com/about-aws/whats-new/2019/12/introducing-aws-identity-and-access-management-access-analyzer/ .
IAM group	See group .
IAM Identity Center	AWS IAM Identity Center is a cloud-based service that brings together administration of users and their access to AWS accounts and cloud applications. You can control single sign-on access and user permissions across all your AWS accounts in AWS Organizations. See Also https://aws.amazon.com/single-sign-on/ .
IAM policy simulator	See policy simulator .
IAM role	See role .
IAM user	See user .
Identity and Access Management	See IAM .
identity provider (IdP)	An IAM entity that holds metadata about external identity providers.
IdP	See identity provider (IdP) .
image	See Amazon Machine Image (AMI) .
Image Builder	EC2 Image Builder is a service that facilitates building, maintaining, and distributing customized server images that launch EC2 instances, or that run in Docker containers. See Also https://aws.amazon.com/image-builder .
implicit impressions	Amazon Personalize : The recommendations that your application shows a user. Unlike <i>explicit impressions</i> , where you manually record each impression, Amazon Personalize automatically derives implicit impressions from your recommendation data.

	See Also recommendations , impressions data , explicit impressions .
import log	A report that contains details about how Import/Export processed your data.
Import/Export	AWS Import/Export is a service for transferring large amounts of data between AWS and portable storage devices. See Also https://aws.amazon.com/importexport .
import/export station	A machine that uploads or downloads your data to or from Amazon S3 .
impressions data	Amazon Personalize : The list of items that you presented to a user when they interacted with a particular item such as by clicking it, watching it, or purchasing it. Amazon Personalize uses impressions data to calculate the relevance of new items for a user based on how frequently users have selected or ignored the same item. See Also explicit impressions , implicit impressions .
index	See search index .
index field	A name–value pair that's included in an CloudSearch domain's index. An index field can contain text or numeric data, dates, or a location.
indexing options	Configuration settings that define an CloudSearch domain's index fields, how document data is mapped to those index fields, and how the index fields can be used.
inline policy	An IAM policy that's embedded in a single IAM user , group , or role .
in-place deployment	CodeDeploy: A deployment method where the application on each instance in the deployment group is stopped, the latest application revision is installed, and the new version of the application is started and validated. You can choose to use a load balancer so each instance is deregistered during its deployment and then restored to service after the deployment is complete.
input data	Amazon Machine Learning: The observations that you provide to Amazon Machine Learning to train and evaluate a machine learning model and generate predictions.

Amazon Inspector	Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed report with prioritized steps for remediation. See Also https://aws.amazon.com/inspector .
instance	A copy of an Amazon Machine Image (AMI) running as a virtual server in the AWS Cloud.
instance family	A general instance type grouping using either storage or CPU capacity.
instance group	A Hadoop cluster contains one master instance group that contains one master node , a core instance group that contains one or more core node and an optional task node instance group, which can contain any number of task nodes.
instance profile	A container that passes IAM role information to an EC2 instance at launch.
instance store	Disk storage that's physically attached to the host computer for an EC2 instance , and therefore has the same lifespan as the instance. When the instance is terminated, you lose any data in the instance store.
instance store-backed AMI	A type of Amazon Machine Image (AMI) whose instances use an instance store volume as the root device. Compare this with instances launched from Amazon EBS-backed AMIs , which use an Amazon EBS volume as the root device.
instance type	A specification that defines the memory, CPU, storage capacity, and usage cost for an instance . Some instance types are for standard applications, whereas others are for CPU-intensive, memory-intensive applications.
Interactions dataset	Amazon Personalize : A container for historical and real-time data collected from interactions between users and items (called events). Interactions data can include impressions data and contextual metadata. See Also dataset , event , impressions data , contextual metadata .

internet gateway	Connects a network to the internet. You can route traffic for IP addresses outside your Amazon VPC to the internet gateway.
internet service provider (ISP)	A company that provides subscribers with access to the internet. Many ISPs are also mailbox providers . Mailbox providers are sometimes referred to as ISPs, even if they only provide mailbox services.
intrinsic function	A special action in a CloudFormation template that assigns values to properties not available until runtime. These functions follow the format <i>Fn::Attribute</i> , such as <code>Fn::GetAtt</code> . Arguments for intrinsic functions can be parameters, pseudo parameters, or the output of other intrinsic functions.
AWS IoT 1-Click	AWS IoT 1-Click is a service that simple devices can use to launch AWS Lambda functions. See Also https://aws.amazon.com/iot-1-click .
AWS IoT Analytics	AWS IoT Analytics is a fully managed service used to run sophisticated analytics on massive volumes of IoT data. See Also https://aws.amazon.com/iot-analytics .
AWS IoT Core	AWS IoT Core is a managed cloud platform that lets connected devices easily and securely interact with cloud applications and other devices. See Also https://aws.amazon.com/iot .
AWS IoT Device Defender	AWS IoT Device Defender is an AWS IoT security service that you can use to audit the configuration of your devices, monitor your connected devices to detect abnormal behavior, and to mitigate security risks. See Also https://aws.amazon.com/iot-device-defender .
AWS IoT Device Management	AWS IoT Device Management is a service used to securely onboard, organize, monitor, and remotely manage IoT devices at scale. See Also https://aws.amazon.com/iot-device-management .
AWS IoT Events	AWS IoT Events is a fully managed AWS IoT service that you can use to detect and respond to events from IoT sensors and applications. See Also https://aws.amazon.com/iot-events .
AWS IoT FleetWise	AWS IoT FleetWise is a service that you can use to collect, transform, and transfer vehicle data to the cloud at scale. See Also https://aws.amazon.com/iot-fleetwise .

AWS IoT Greengrass	AWS IoT Greengrass is a software that you can use to run local compute, messaging, data caching, sync, and ML inference capabilities for connected devices in a secure way. See Also https://aws.amazon.com/greengrass .
AWS IoT RoboRunner	AWS IoT RoboRunner is a solution that provides infrastructure for integrating robots with work management systems and building robotics fleet management applications. See Also https://aws.amazon.com/roborunner .
AWS IoT SiteWise	AWS IoT SiteWise is a managed service that you can use to collect, organize, and analyze data from industrial equipment at scale. See Also https://aws.amazon.com/iot-sitewise .
AWS IoT Things Graph	AWS IoT Things Graph is a service that you can use to visually connect different devices and web services to build IoT applications. See Also https://aws.amazon.com/iot-things-graph .
IP address	A numerical address (for example, 192.0.2.44) that networked devices use to communicate with one another using the Internet Protocol (IP). Each EC2 instance is assigned two IP addresses at launch, which are directly mapped to each other through network address translation (NAT): a private IP address (following RFC 1918) and a public IP address. Instances launched in a VPC are assigned only a private IP address. Instances launched in your default VPC are assigned both a private IP address and a public IP address.
IP match condition	AWS WAF : An attribute that specifies the IP addresses or IP address ranges that web requests originate from. Based on the specified IP addresses, you can configure AWS WAF to allow or block web requests to AWS resources such as Amazon CloudFront distributions.
AWS IQ	AWS IQ is a cloud service that AWS customers can use to find, engage, and pay AWS Certified third-party experts for on-demand project work. See Also ???TITLE??? .
ISP	See internet service provider (ISP) .
issuer	The person who writes a policy to grant permissions to a resource . The issuer (by definition) is always the resource owner. AWS doesn't permit

	<p>Amazon SQS users to create policies for resources they don't own. If John is the resource owner, AWS authenticates John's identity when he submits the policy he's written to grant permissions for that resource.</p>
item	<p>A group of attributes that's uniquely identifiable among all of the other items. Items in DynamoDB are similar in many ways to rows, records, or tuples in other database systems.</p>
item exploration	<p>Amazon Personalize: The process that Amazon Personalize uses to test different item recommendations, including recommendations of new items with no or little interaction data, and learn how users respond. You configure item exploration at the campaign level for solution versions created with the user-personalization recipe.</p> <p>See Also recommendations, campaign, solution version, user-personalization recipe.</p>
Items dataset	<p>Amazon Personalize: A container for metadata about items, such as price, genre, or availability.</p> <p>See Also dataset.</p>
item-to-item similarities (SIMS) recipe	<p>Amazon Personalize: A RELATED_ITEMS recipe that uses the data from an Interactions dataset to make recommendations for items that are similar to a specified item. The SIMS recipe calculates similarity based on the way users interact with items instead of matching item metadata, such as price or age.</p> <p>See Also recipe, RELATED_ITEMS recipes, Interactions dataset.</p>

J

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

job flow	<p>Amazon EMR: One or more steps that specify all of the functions to be performed on the data.</p>
job ID	<p>A five-character, alphanumeric string that uniquely identifies an Import/Export storage device in your shipment. AWS issues the job ID in response to a CREATE JOB email command.</p>

job prefix	An optional string that you can add to the beginning of an Import/Export log file name to prevent collisions with objects of the same name. See Also key prefix .
JSON	JavaScript Object Notation. A lightweight data interchange format. For information about JSON, see http://www.json.org/ .
junk folder	The location where email messages that various filters determine to be of lesser value are collected so that they don't arrive in the recipient's inbox but are still accessible to the recipient. This is also referred to as a spam or bulk folder.

K

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

Amazon Kendra	Amazon Kendra is a search service powered by machine learning (ML) that developers can use to add search capabilities to their applications so their end users can discover information stored within the vast amount of content spread across their company. See Also https://aws.amazon.com/kendra/ .
key	A credential that identifies an AWS account or user to AWS (such as the AWS secret access key). Amazon S3 , Amazon EMR : The unique identifier for an object in a bucket . Every object in a bucket has exactly one key. Because a bucket and key together uniquely identify each object, you can think of Amazon S3 as a basic data map between the <i>bucket + key</i> , and the object itself. You can uniquely address every object in Amazon S3 through the combination of the web service endpoint, bucket name, and key, as in this example: <code>http://doc.s3.amazonaws.com/2006-03-01/AmazonS3.wsd1</code> , where <code>doc</code> is the name of the bucket, and <code>2006-03-01/AmazonS3.wsd1</code> is the key. Import/Export : The name of an object in Amazon S3. It's a sequence of Unicode characters whose UTF-8 encoding can't exceed 1024 bytes. If

a key (for example, logPrefix + import-log-JOBID) is longer than 1024 bytes, [Elastic Beanstalk](#) returns an InvalidManifestField error.

[IAM](#): In a [policy](#), a specific characteristic that's the basis for restricting access (such as the current time or the IP address of the requester).

Tagging resources: A general [tag](#) label that acts like a category for more specific tag values. For example, you might have [EC2 instance](#) with the tag key of *Owner* and the tag value of *Jan*. You can tag an AWS [resource](#) with up to 10 key–value pairs. Not all AWS resources can be tagged.

key pair

A set of security credentials that you use to prove your identity electronically. A key pair consists of a private key and a public key.

key prefix

A string of characters that is a subset of an object key name, starting with the first character. The prefix can be any length, up to the maximum length of the object key name (1,024 bytes).

Amazon Keyspaces

Amazon Keyspaces (for Apache Cassandra) is a scalable, highly available, and managed Apache Cassandra-compatible database service. See Also <https://aws.amazon.com/keyspaces/>.

kibibyte (KiB)

A contraction of kilo binary byte, a kibibyte is 2^{10} or 1,024 bytes. A kilobyte (KB) is 10^3 or 1,000 bytes. 1,024 KiB is a [mebibyte \(MiB\)](#).

Kinesis

Amazon Kinesis is a platform for streaming data on AWS. Kinesis offers services that simplify the loading and analysis of streaming data. See Also <https://aws.amazon.com/kinesis/>.

Firehose

Amazon Data Firehose is a fully managed service for loading streaming data into AWS. Firehose can capture and automatically load streaming data into [Amazon S3](#) and [Amazon Redshift](#), enabling near real-time analytics with existing business intelligence tools and dashboards. Firehose automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, and encrypt the data before loading it. See Also <https://aws.amazon.com/kinesis/firehose/>.

Kinesis Data Streams

Amazon Kinesis Data Streams is a web service for building custom applications that process or analyze streaming data for specialized

needs. Amazon Kinesis Data Streams can continuously capture and store terabytes of data per hour from hundreds of thousands of sources. See Also <https://aws.amazon.com/kinesis/streams/>.

AWS KMS

AWS Key Management Service is a managed service that simplifies the creation and control of encryption keys that are used to encrypt data. See Also <https://aws.amazon.com/kms>.

KMS key

The primary resource in AWS Key Management Service. In general, KMS keys are created, used, and deleted entirely within KMS. KMS supports symmetric and asymmetric KMS keys for encryption and signing. KMS keys can be either customer managed, AWS managed, or AWS owned. For more information, see AWS KMS keys in the [AWS Key Management Service Developer Guide](#).

L

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

labeled data

In machine learning, data for which you already know the target or “correct” answer.

Lake Formation

AWS Lake Formation is a managed service that makes it easy to set up, secure, and manage your data lakes. Lake Formation helps you discover your data sources and then catalog, cleanse, and transform the data. See Also <https://aws.amazon.com/lake-formation>.

Lambda

AWS Lambda is a web service that you can use to run code without provisioning or managing servers. You can run code for virtually any type of application or backend service with zero administration. You can set up your code to automatically start from other AWS services or call it directly from any web or mobile app. See Also <https://aws.amazon.com/lambda/>.

launch configuration

A set of descriptive parameters used to create new [EC2 instances](#) in an [Amazon EC2 Auto Scaling](#) activity.

A template that an [Auto Scaling group](#) uses to launch new EC2 instances. The launch configuration contains information such as the

	<p>Amazon Machine Image (AMI) ID, the instance type, key pairs, security groups, and block device mappings, among other configuration settings.</p>
launch permission	An Amazon Machine Image (AMI) attribute that allows users to launch an AMI.
Launch Wizard	<p>AWS Launch Wizard is a cloud solution that offers a guided way of sizing, configuring, and deploying AWS resources for third-party applications, such as Microsoft SQL Server Always On and HANA based SAP systems, without the need to manually identify and provision individual AWS resources.</p> <p>See Also https://aws.amazon.com/launchwizard.</p>
Amazon Lex	<p>Amazon Lex is a fully managed artificial intelligence (AI) service with advanced natural language models to design, build, test, and deploy conversational interfaces in applications.</p> <p>See Also https://aws.amazon.com/lex/.</p>
lifecycle	The lifecycle state of the EC2 instance contained in an Auto Scaling group . EC2 instances progress through several states over their lifespan; these include <i>Pending</i> , <i>InService</i> , <i>Terminating</i> and <i>Terminated</i> .
lifecycle action	An action that can be paused by Auto Scaling, such as launching or terminating an EC2 instance.
lifecycle hook	A feature for pausing Auto Scaling after it launches or terminates an EC2 instance so that you can perform a custom action while the instance isn't in service.
Lightsail	<p>Amazon Lightsail is a service used to launch and manage a virtual private server with AWS. Lightsail offers bundled plans that include everything you need to deploy a virtual private server, for a low monthly rate.</p> <p>See Also https://aws.amazon.com/lightsail/.</p>
load balancer	A DNS name combined with a set of ports, which together provide a destination for all requests intended for your application. A load balancer can distribute traffic to multiple application instances across every Availability Zone within a Region . Load balancers can span multiple Availability Zones within an AWS Region into which an Amazon

	<p>EC2 instance was launched. But load balancers can't span multiple Regions.</p>
local secondary index	<p>An index that has the same partition key as the table, but a different sort key. A local secondary index is local in the sense that every partition of a local secondary index is scoped to a table partition that has the same partition key value.</p> <p>See Also local secondary index.</p>
Amazon Location	<p>Amazon Location Service is a fully managed service that makes it easy for a developer to add location functionality, such as maps, points of interest, geocoding, routing, tracking, and geofencing, to their applications, without sacrificing data security, user privacy, data quality, or cost.</p> <p>See Also https://aws.amazon.com/location/.</p>
logical name	<p>A case-sensitive unique string within an CloudFormation template that identifies a resource, mapping, parameter, or output. In an AWS CloudFormation template, each parameter, resource, property, mapping, and output must be declared with a unique logical name. You use the logical name when dereferencing these items using the Ref function.</p>
Lookout for Equipment	<p>Amazon Lookout for Equipment is a machine learning service that uses data from sensors mounted on factory equipment to detect abnormal behavior so you can take action before machine failures occur.</p> <p>See Also https://aws.amazon.com/lookout-for-equipment/.</p>
Lookout for Metrics	<p>Amazon Lookout for Metrics is a machine learning (ML) service that automatically detects and diagnoses anomalies in business and operational data, such as a sudden dip in sales revenue or customer acquisition rates.</p> <p>See Also https://aws.amazon.com/lookout-for-metrics.</p>
Lookout for Vision	<p>Amazon Lookout for Vision is a machine learning service that uses computer vision (CV) to find defects in industrial products. Amazon Lookout for Vision can identify missing components in an industrial product, damage to vehicles or structures, irregularities in production lines, and even minuscule defects in silicon wafers—or any other physical item where quality is important.</p>

See Also <https://aws.amazon.com/lookout-for-vision/>.

Lumberyard See [O3DE](#).

M

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

Macie Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. See Also <http://aws.amazon.com/macie/>.

Mail Transfer Agent (MTA) Software that transports email messages from one computer to another by using a client-server architecture.

mailbox provider An organization that provides email mailbox hosting services. Mailbox providers are sometimes referred to as [internet service providers \(ISPs\)](#), even if they only provide mailbox services.

mailbox simulator A set of email addresses that you can use to test an [Amazon SES](#)-based email-sending application without sending messages to actual recipients. Each email address represents a specific scenario (such as a bounce or complaint) and generates a typical response that's specific to the scenario.

main route table The default [route table](#) that any new [Amazon VPC subnet](#) uses for routing. You can associate a subnet with a different route table of your choice. You can also change which route table is the main route table.

AWS Mainframe Modernization AWS Mainframe Modernization service is a cloud native platform for migration, modernization, execution, and operation of mainframe applications. See Also <https://aws.amazon.com/mainframe-modernization>.

Managed Blockchain Amazon Managed Blockchain is a fully managed service for creating and managing scalable blockchain networks using popular open source frameworks. See Also <http://aws.amazon.com/managed-blockchain/>.

Amazon Managed Grafana	<p>Amazon Managed Grafana is a fully managed and secure data visualization service that you can use to instantly query, correlate, and visualize operational metrics, logs, and traces from multiple data sources.</p> <p>See Also https://aws.amazon.com/grafana/.</p>
AWS managed key	<p>One type of KMS key in AWS KMS.</p>
managed policy	<p>A standalone IAM policy that you can attach to multiple users, groups, and roles in your IAM account. Managed policies can either be AWS managed policies (which are created and managed by AWS) or customer managed policies (which you create and manage in your AWS account).</p>
AWS managed policy	<p>An IAM managed policy that's created and managed by AWS.</p>
Amazon Managed Service for Prometheus	<p>Amazon Managed Service for Prometheus is a service that provides highly available, secure, and managed monitoring for your containers.</p> <p>See Also https://aws.amazon.com/prometheus/.</p>
AWS Management Console	<p>AWS Management Console is a graphical interface to manage compute, storage, and other cloud resources.</p> <p>See Also https://aws.amazon.com/console.</p>
management portal	<p>AWS Management Portal for vCenter is a web service for managing your AWS resources using VMware vCenter. You install the portal as a vCenter plugin within your existing vCenter environment. After it's installed, you can migrate VMware VMs to Amazon EC2 and manage AWS resources from within vCenter.</p> <p>See Also https://aws.amazon.com/ec2/vcenter-portal/.</p>
manifest	<p>When sending a <i>create job</i> request for an import or export operation, you describe your job in a text file called a manifest. The manifest file is a YAML-formatted file that specifies how to transfer data between your storage device and the AWS Cloud.</p>
manifest file	<p>Amazon Machine Learning: The file used for describing batch predictions. The manifest file relates each input data file with its associated batch prediction results. It's stored in the Amazon S3 output location.</p>

mapping	A way to add conditional parameter values to an CloudFormation template. You specify mappings in the template's optional Mappings section and retrieve the desired value using the FN: :FindInMap function.
marker	See pagination token .
AWS Marketplace	AWS Marketplace is a web portal where qualified partners market and sell their software to AWS customers. AWS Marketplace is an online software store that helps customers find, buy, and immediately start using the software and services that run on AWS. See Also https://aws.amazon.com/partners/aws-marketplace/ .
master node	A process running on an Amazon Machine Image (AMI) that keeps track of the work its core and task nodes complete.
maximum price	The maximum price you pay to launch one or more Spot Instances . If your maximum price exceeds the current Spot price and your restrictions are met, Amazon EC2 launches instances on your behalf.
maximum send rate	The maximum number of email messages that you can send per second using Amazon SES .
mean reciprocal rank at 25	Amazon Personalize : An evaluation metric that assesses the relevance of a model's highest ranked recommendation. Amazon Personalize calculates this metric using the average accuracy of the model when ranking the most relevant recommendation out of the top 25 recommendations over all requests for recommendations. See Also metrics , recommendations .
mebibyte (MiB)	A contraction of mega binary byte. A mebibyte (MiB) is 2^{20} or 1,048,576 bytes. A megabyte (MB) is 10^6 or 1,000,000 bytes. 1,024 MiB is a gibibyte (GiB) .
member resources	See resource .
MemoryDB	Amazon MemoryDB is a Redis-compatible, durable, in-memory database service that's purpose-built for modern applications with microservices architectures. See Also https://aws.amazon.com/memorydb .

message ID	<p>Amazon SES: A unique identifier that's assigned to every email message that's sent.</p> <p>Amazon SQS: The identifier returned when you send a message to a queue.</p>
metadata	<p>Information about other data or objects. In Amazon S3 and Amazon EMR metadata takes the form of name–value pairs that describe the object. These include default metadata such as the date last modified and standard HTTP metadata (for example, Content-Type). Users can also specify custom metadata at the time they store an object. In Amazon EC2 metadata includes data about an EC2 instance that the instance can retrieve to determine things about itself, such as the instance type or the IP address.</p>
metric	<p>An element of time-series data defined by a unique combination of exactly one namespace, exactly one metric name, and between zero and ten dimensions. Metrics and the statistics derived from them are the basis of CloudWatch.</p>
metric name	<p>The primary identifier of a metric, used with a namespace and optional dimensions.</p>
metrics	<p>Amazon Personalize: Evaluation data that Amazon Personalize generates when you train a model. You use metrics to evaluate the performance of the model, view the effects of modifying a solution's configuration, and compare results between solutions that use the same training data but were created with different recipes.</p> <p>See Also solution, recipe.</p>
MFA	<p>See multi-factor authentication (MFA).</p>
micro instance	<p>A type of EC2 instance that's more economical to use if you have occasional bursts of high CPU activity.</p>
AWS Microservice Extractor for .NET	<p>AWS Microservice Extractor for .NET is an assistive modernization tool that helps to reduce the time and effort required to break down large, monolithic applications running on the AWS Cloud or on premises into smaller, independent services. These services can be operated and managed independently.</p>

Migration Hub	AWS Migration Hub is a service that provides a single location to track migration tasks across multiple AWS tools and partner solutions. See Also https://aws.amazon.com/migration-hub/ .
MIME	See Multipurpose Internet Mail Extensions (MIME) .
Amazon ML	Amazon Machine Learning is a cloud-based service that creates machine learning (ML) models by finding patterns in your data, and uses these models to process new data and generate predictions. See Also http://aws.amazon.com/machine-learning/ .
ML model	In machine learning (ML), a mathematical model that generates predictions by finding patterns in data. Amazon Machine Learning supports three types of ML models: binary classification, multiclass classification, and regression. Also known as a <i>predictive model</i> . See Also binary classification model , multiclass classification model , regression model .
Mobile Analytics	Amazon Mobile Analytics is a service for collecting, visualizing, understanding, and extracting mobile app usage data at scale. See Also https://aws.amazon.com/mobileanalytics .
Mobile Hub	See Amplify .
AWS Mobile SDK	See Amplify .
Mobile SDK for Android	See Amplify Android .
Mobile SDK for iOS	See Amplify iOS .
Mobile SDK for Unity	The AWS Mobile SDK for Unity is included in the AWS SDK for .NET .
Mobile SDK for Xamarin	The AWS Mobile SDK for Xamarin is included in the AWS SDK for .NET .
Amazon Monitron	Amazon Monitron is an end-to-end system that uses machine learning (ML) to detect abnormal behavior in industrial machinery. Use Amazon Monitron to implement predictive maintenance and reduce unplanned downtime. See Also https://aws.amazon.com/monitron/ .

Amazon MQ	Amazon MQ is a managed message broker service for Apache ActiveMQ that you can use to set up and operate message brokers in the cloud. See Also https://aws.amazon.com/amazon-mq/ .
MTA	See Mail Transfer Agent (MTA) .
Multi-AZ deployment	A primary DB instance that has a synchronous standby replica in a different Availability Zone . The primary DB instance is synchronously replicated across Availability Zones to the standby replica.
multiclass classification model	A machine learning model that predicts values that belong to a limited, pre-defined set of permissible values. For example, "Is this product a book, movie, or clothing?"
multi-factor authentication (MFA)	An optional AWS account security feature. After you enable AWS MFA, you must provide a six-digit, single-use code in addition to your sign-in credentials whenever you access secure AWS webpages or the AWS Management Console . You get this single-use code from an authentication device that you keep in your physical possession. See Also https://aws.amazon.com/mfa/ .
multipart upload	A feature that you can use to upload a single object as a set of parts.
Multipurpose Internet Mail Extensions (MIME)	An internet standard that extends the email protocol to include non-ASCII text and nontext elements, such as attachments.
Multitool	A cascading application that provides a simple command-line interface for managing large datasets.
multi-valued attribute	An attribute with more than one value.
Amazon MWAA	Amazon Managed Workflows for Apache Airflow is a managed orchestration service for Apache Airflow to assist in setting up and operating end-to-end data pipelines in the cloud at scale. See Also https://aws.amazon.com/managed-workflows-for-apache-airflow .

N

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

namespace	An abstract container that provides context for the items (names, or technical terms, or words) it holds, and allows disambiguation of homonym items residing in different namespaces.
NAT	Network address translation. A strategy of mapping one or more IP addresses to another while data packets are in transit across a traffic routing device. This is commonly used to restrict internet communication to private instances while allowing outgoing traffic. See Also Network Address Translation and Protocol Translation , NAT gateway , NAT instance .
NAT gateway	A NAT device, managed by AWS, that performs network address translation in a private subnet , to secure inbound internet traffic. A NAT gateway uses both NAT and port address translation. See Also NAT instance .
NAT instance	A NAT device, configured by a user, that performs network address translation in a Amazon VPC public subnet to secure inbound internet traffic. See Also NAT gateway .
Neptune	Amazon Neptune is a managed graph database service that you can use to build and run applications that work with highly connected datasets. Neptune supports the popular graph query languages Apache TinkerPop Gremlin and W3C's SPARQL, enabling you to build queries that efficiently navigate highly connected datasets. See Also https://aws.amazon.com/neptune/ .
network ACL	An optional layer of security that acts as a firewall for controlling traffic in and out of a subnet . You can associate multiple subnets with a single network ACL , but a subnet can be associated with only one network ACL at a time.
Network Address Translation and Protocol Translation	(NAT-PT) An internet protocol standard defined in RFC 2766. See Also NAT instance , NAT gateway .
Network Firewall	AWS Network Firewall is a managed service that deploys essential network protections for all Amazon Virtual Private Clouds (Amazon VPCs).

	See Also https://aws.amazon.com/network-firewall .
n-gram processor	A processor that performs n-gram transformations. See Also n-gram transformation .
n-gram transformation	Amazon Machine Learning: A transformation that aids in text string analysis. An n-gram transformation takes a text variable as input and outputs strings by sliding a window of size n words, where n is specified by the user, over the text, and outputting every string of words of size n and all smaller sizes. For example, specifying the n-gram transformation with window size =2 returns all the two-word combinations and all of the single words.
NICE Desktop Cloud Visualization	A remote visualization technology for securely connecting users to graphic-intensive 3D applications hosted on a remote, high-performance server.
Nimble Studio	Amazon Nimble Studio is a managed AWS cloud service for creative studios to produce visual effects, animation, and interactive content—from storyboard to final deliverable. See Also https://aws.amazon.com/nimble-studio/ .
node	OpenSearch Service : An OpenSearch instance. A node can be either a data instance or a dedicated master instance. See Also dedicated master node .
NoEcho	A property of CloudFormation parameters that prevent the otherwise default reporting of names and values of a template parameter. Declaring the NoEcho property causes the parameter value to be masked with asterisks in the report by the <code>cfn-describe-stacks</code> command.
normalized discounted cumulative gain (NCDG) at K (5/10/25)	Amazon Personalize : An evaluation metric that tells you about the relevance of your model's highly ranked recommendations, where K is a sample size of 5, 10, or 25 recommendations. Amazon Personalize calculates this by assigning weight to recommendations based on their position in a ranked list, where each recommendation is discounted (given a lower weight) by a factor dependent on its position. The normalized discounted cumulative gain at K assumes

that recommendations that are lower on a list are less relevant than recommendations higher on the list.

See Also [metrics](#), [recommendations](#).

NoSQL

Nonrelational database systems that are highly available, scalable, and optimized for high performance. Instead of the relational model, NoSQL databases (for example, [DynamoDB](#)) use alternate models for data management, such as key–value pairs or document storage.

null object

A null object is one whose version ID is null. [Amazon S3](#) adds a null object to a [bucket](#) when [versioning](#) for that bucket is suspended. It's possible to have only one null object for each key in a bucket.

number of passes

The number of times that you allow Amazon Machine Learning to use the same data records to train a machine learning model.

O

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

O3DE

Open 3D Engine (successor to Amazon Lumberyard) is an open-source 3D development engine for creating games and simulations. O3DE is licensed under Apache 2.0 and maintained by a community of contributors, including Amazon.

See Also <https://www.o3de.org/>, <https://aws.amazon.com/lumberyard/>, <https://docs.aws.amazon.com/lumberyard/>.

object

[Amazon S3](#): The fundamental entity type stored in Amazon S3. Objects consist of object data and metadata. The data portion is opaque to Amazon S3.

[CloudFront](#): Any entity that can be served either over HTTP or a version of RTMP.

observation

Amazon Machine Learning: A single instance of data that Amazon Machine Learning (Amazon ML) uses to either train a machine learning model how to predict or to generate a prediction. Each row in an Amazon ML input data file is an observation.

On-Demand Instance	An Amazon EC2 pricing option that charges you for compute capacity by the hour or second (minimum of 60 seconds) with no long-term commitment.
Open 3D Engine	See O3DE .
OpenSearch Service	Amazon OpenSearch Service is an AWS managed service for deploying, operating, and scaling OpenSearch, an open-source search and analytics engine, in the AWS Cloud. Amazon OpenSearch Service (OpenSearch Service) also offers security options, high availability, data durability, and direct access to the OpenSearch API. See Also https://aws.amazon.com/elasticsearch-service .
operation	An API function. Also called an <i>action</i> .
OpsWorks	AWS OpsWorks is a configuration management service that helps you use Chef to configure and operate groups of instances and applications. You can define the application's architecture and the specification of each component including package installation, software configuration, and resources such as storage. You can automate tasks based on time, load, or lifecycle events. See Also https://aws.amazon.com/opsworks/ .
optimistic locking	A strategy to ensure that an item that you want to update has not been modified by others before you perform the update. For DynamoDB , optimistic locking support is provided by the AWS SDKs.
opt-in Region	An AWS Region that is disabled by default. To use an opt-in Region, you must enable it. Regions introduced after March 20, 2019 are opt-in Regions. For a list of opt-in Regions, see Considerations before enabling and disabling Regions in the <i>AWS Account Management Guide</i> . See Also Region that is enabled by default .
organization	Organizations : An entity that you create to consolidate and manage your AWS accounts. An organization has one management account along with zero or more member accounts.
organizational unit	Organizations : A container for accounts within a root of an organization. An organizational unit (OU) can contain other OUs.

Organizations	<p>AWS Organizations is an account management service that you can use to consolidate multiple AWS accounts into an organization that you create and centrally manage.</p> <p>See Also https://aws.amazon.com/organizations/.</p>
origin access identity	<p>Also called OAI. When using Amazon CloudFront to serve content with an Amazon S3 bucket as the origin, a virtual identity that you use to require users to access your content through CloudFront URLs instead of Amazon S3 URLs. Usually used with CloudFront private content.</p>
origin server	<p>The Amazon S3 bucket or custom origin containing the definitive original version of the content you deliver through CloudFront.</p>
original environment	<p>The instances in a deployment group at the start of an CodeDeploy blue/green deployment.</p>
OSB transformation	<p>Orthogonal sparse bigram transformation. In machine learning, a transformation that aids in text string analysis and that's an alternative to the n-gram transformation. OSB transformations are generated by sliding the window of size n words over the text, and outputting every pair of words that includes the first word in the window.</p> <p>See Also n-gram transformation.</p>
OU	<p>See organizational unit.</p>
Outposts	<p>AWS Outposts is a fully managed service by AWS that extends AWS infrastructure, services, APIs, and tools to on-premises data centers and edge locations. Use AWS Outposts for workloads and devices requiring low latency access to on-premises systems, local data processing, data residency, and application migration with local system interdependencies.</p> <p>See Also https://aws.amazon.com/outposts.</p>
output location	<p>Amazon Machine Learning: An Amazon S3 location where the results of a batch prediction are stored.</p>

P

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

pagination	<p>The process of responding to an API request by returning a large list of records in small separate parts. Pagination can occur in the following situations:</p> <ul style="list-style-type: none">• The client sets the maximum number of returned records to a value below the total number of records.• The service has a default maximum number of returned records that's lower than the total number of records. <p>When an API response is paginated, the service sends a subset of the large list of records and a pagination token that indicates that more records are available. The client includes this pagination token in a subsequent API request, and the service responds with the next subset of records. This continues until the service responds with a subset of records and no pagination token, indicating that all records have been sent.</p>
pagination token	<p>A marker that indicates that an API response contains a subset of a larger list of records. The client can return this marker in a subsequent API request to retrieve the next subset of records until the service responds with a subset of records and no pagination token, indicating that all records have been sent.</p> <p>See Also pagination.</p>
paid AMI	<p>An Amazon Machine Image (AMI) that you sell to other Amazon EC2 users on AWS Marketplace.</p>
AWS Panorama	<p>AWS Panorama is a machine learning (ML) Appliance and Software Development Kit (SDK) that organizations can use to bring computer vision (CV) to on-premises cameras to make predictions locally.</p> <p>See Also https://aws.amazon.com/panorama.</p>
AWS ParallelCluster	<p>AWS ParallelCluster is an AWS supported open source cluster management tool that helps you to deploy and manage high performance computing (HPC) clusters in the AWS Cloud.</p>
paravirtual virtualization	<p>See PV virtualization.</p>

part	A contiguous portion of the object's data in a multipart upload request.
partition	<p>A group of AWS Regions. Each Region is in only one partition, and each partition contains one or more Regions. Partitions have independent instances of the AWS Identity and Access Management (IAM) infrastructure. In other words, a partition is comprised of Regions that share the same authentication, account, and resource stack. Each AWS account is scoped to one partition. You can't use IAM credentials from one partition to interact with resources in a different partition.</p> <p>Some AWS services are designed to provide cross-Region functionality. Such cross-Region functionality is supported only between Regions in the same partition. AWS commercial Regions are in the AWS partition, China Regions are in the AWS-cn partition, and AWS GovCloud (US) Regions are in the AWS-us-gov partition.</p>
partition key	<p>A simple primary key, composed of one attribute (also known as a <i>hash attribute</i>).</p> <p>See Also primary key, sort key.</p>
PAT	Port address translation.
pebibyte (PiB)	A contraction of peta binary byte, a pebibyte is 2^{50} or 1,125,899,906,842,624 bytes. A petabyte (PB) is 10^{15} or 1,000,000,000,000,000 bytes. 1,024 PiB is an exbibyte (EiB) .
period	See sampling period .
permission	A statement within a policy that allows or denies access to a particular resource . You can state any permission in the following way: "A has permission to do B to C." For example, Jane (A) has permission to read messages (B) from John's Amazon SQS queue (C). Whenever Jane sends a request to Amazon SQS to use John's queue, the service checks to see if she has permission. It further checks to see if the request satisfies the conditions John set forth in the permission.
persistent storage	A data storage solution where the data remains intact until it's deleted. Options within AWS include: Amazon S3 , Amazon RDS , DynamoDB , and other services.

Amazon Personalize	Amazon Personalize is an artificial intelligence service for creating individualized product and content recommendations. See Also https://aws.amazon.com/personalize/ .
PERSONALIZED_RANKING recipes	Amazon Personalize : Recipes that provide item recommendations in ranked order based on the predicted interest for a user. See Also recipe , recommendations , personalized-ranking recipe , popularity-count recipe .
personalized-ranking recipe	Amazon Personalize : A PERSONALIZED_RANKING recipe that ranks a collection of items that you provide based on the predicted interest level for a specific user. Use the personalized-ranking recipe to create curated lists of items or ordered search results that are personalized for a specific user. See Also recipe , PERSONALIZED_RANKING recipes .
physical name	A unique label that CloudFormation assigns to each resource when creating a stack . Some AWS CloudFormation commands accept the physical name as a value with the <code>--physical-name</code> parameter.
pilot light	An active-passive disaster recovery strategy in which you replicate data from the primary Region as standby, then provision a replica that contains only the core workload infrastructure. To make this infrastructure functional and serve requests, you must provision the remaining resources, such as compute. See Also , , .
Amazon Pinpoint	Amazon Pinpoint is a multichannel communications service that helps organizations send timely, targeted content through SMS, email, mobile push notifications, voice messages, and in-application channels. See Also https://aws.amazon.com/pinpoint .
pipeline	CodePipeline : A workflow construct that defines the way software changes go through a release process.
plaintext	Information that has not been encrypted , as opposed to ciphertext .
policy	IAM : A document defining permissions that apply to a user, group, or role; the permissions in turn determine what users can do in AWS. A

	<p>policy typically allows access to specific actions, and can optionally grant that the actions are allowed for specific resources, such as EC2 instances or Amazon S3 buckets. Policies can also explicitly deny access.</p> <p>Amazon EC2 Auto Scaling: An object that stores the information that's needed to launch or terminate instances for an Auto Scaling group. Running the policy causes instances to be launched or terminated. You can configure an alarm to invoke an Auto Scaling policy.</p>
policy generator	A tool in the IAM AWS Management Console that helps you build a policy by selecting elements from lists of available options.
policy simulator	A tool in the IAM AWS Management Console that helps you test and troubleshoot policies so you can see their effects in real-world scenarios.
policy validator	A tool in the IAM AWS Management Console that examines your existing IAM access control policies to ensure that they comply with the IAM policy grammar.
Amazon Polly	Amazon Polly is a text-to-speech (TTS) service that turns text into natural-sounding human speech. Amazon Polly provides dozens of lifelike voices across a broad set of languages so that you can build speech-enabled applications that work in many different countries. See Also https://aws.amazon.com/polly/ .
popularity-count recipe	Amazon Personalize : A USER_PERSONALIZATION recipe that recommends the items that have had the most interactions with unique users. See Also recipe , USER_PERSONALIZATION recipes .
Porting Assistant for .NET	Porting Assistant for .NET is a compatibility analyzer that reduces the manual effort required to port Microsoft .NET Framework applications to open source .NET Core.
precision at K (5/10/25)	Amazon Personalize : An evaluation metric that tells you how relevant your model's recommendations are based on a sample size of K (5, 10, or 25) recommendations. Amazon Personalize calculates this metric based on the number of relevant recommendations out of the top K recommendations, divided by K, where K is 5, 10, or 25.

	See Also metrics , recommendations .
prefix	See job prefix .
Premium Support	A one-on-one, fast-response support channel that AWS customers can subscribe to for support for AWS infrastructure services. See Also https://aws.amazon.com/premiumsupport/ .
presigned URL	A web address that uses query string authentication .
primary key	One or two attributes that uniquely identify each item in a DynamoDB table, so that no two items can have the same key. See Also partition key , sort key .
primary shard	See shard .
principal	The user , service, or account that receives permissions that are defined in a policy . The principal is A in the statement "A has permission to do B to C."
AWS Private CA	AWS Private Certificate Authority is a hosted private certificate authority service for issuing and revoking private digital certificates . See Also https://aws.amazon.com/certificate-manager/private-certificate-authority/ .
private content	When using Amazon CloudFront to serve content with an Amazon S3 bucket as the origin, a method of controlling access to your content by requiring users to use signed URLs. Signed URLs can restrict user access based on the current date and time, the IP addresses that the requests originate from, or both.
private IP address	A private numerical address (for example, 192.0.2.44) that networked devices use to communicate with one another using the Internet Protocol (IP). Each EC2 instance is assigned two IP addresses at launch, which are directly mapped to each other through network address translation (NAT): a private address (following RFC 1918) and a public address. <i>Exception:</i> Instances launched in Amazon VPC are assigned only a private IP address.
private subnet	A Amazon VPC subnet whose instances can't be reached from the internet.

product code	An identifier provided by AWS when you submit a product to AWS Marketplace .
properties	See resource property .
property rule	A JSON -compliant markup standard for declaring properties, mappings, and output values in an CloudFormation template.
Provisioned IOPS	A storage option that delivers fast, predictable, and consistent I/O performance. When you specify an IOPS rate while creating a DB instance, Amazon RDS provisions that IOPS rate for the lifetime of the DB instance.
pseudo parameter	A predefined setting (for example, <code>AWS:StackName</code>) that can be used in CloudFormation templates without having to declare them. You can use pseudo parameters anywhere you can use a regular parameter.
public AMI	An Amazon Machine Image (AMI) that all AWS accounts have permission to launch.
public dataset	<p>A large collection of public information that can be seamlessly integrated into applications that are based in the AWS Cloud. Amazon stores public datasets at no charge to the community and, similar to other AWS services, users pay only for the compute and storage they use for their own applications. These datasets currently include data from the Human Genome Project, the US Census, Wikipedia, and other sources.</p> <p>See Also https://aws.amazon.com/publicdatasets.</p>
public IP address	A public numerical address (for example, 192.0.2.44) that networked devices use to communicate with one another using the Internet Protocol (IP). Each EC2 instance is assigned two IP addresses at launch, which are directly mapped to each other through Network Address Translation (NAT): a private address (following RFC 1918) and a public address. <i>Exception:</i> Instances launched in Amazon VPC are assigned only a private IP address.
public subnet	A subnet whose instances can be reached from the internet.
PV virtualization	Paravirtual virtualization allows guest VMs to run on host systems that don't have special support extensions for full hardware and CPU

virtualization. Because PV guests run a modified operating system that doesn't use hardware emulation, they can't provide hardware-related features, such as enhanced networking or GPU support.

See Also [HVM virtualization](#).

Q

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

Amazon QLDB	Amazon Quantum Ledger Database (Amazon QLDB) is a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log owned by a central trusted authority. See Also https://aws.amazon.com/qldb .
quartile binning transformation	Amazon Machine Learning: A process that takes two inputs, a numerical variable and a parameter called a bin number, and outputs a categorical variable. Quartile binning transformations discover non-linearity in a variable's distribution by enabling the machine learning model to learn separate importance values for parts of the numeric variable's distribution.
Query	A type of web service that generally uses only the GET or POST HTTP method and a query string with parameters in the URL. See Also REST .
query string authentication	An AWS feature that you can use to place the authentication information in the HTTP request query string instead of in the Authorization header, which provides URL-based access to objects in a bucket .
queue	A sequence of messages or jobs that are held in temporary storage awaiting transmission or processing.
queue URL	A web address that uniquely identifies a queue.
QuickSight	Amazon QuickSight is a fast, cloud-powered business analytics service that you can use to build visualizations, perform analysis, and quickly get business insights from your data. See Also https://aws.amazon.com/quicksight/ .

quota The maximum value for your resources, actions, and items in your AWS account

R

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

AWS RAM AWS Resource Access Manager is a web service that AWS customers can use to securely share AWS resources with any AWS account or within your organization.
See Also <https://aws.amazon.com/ram>.

range GET A request that specifies a byte range of data to get for a download. If an object is large, you can break up a download into smaller units by sending multiple range GET requests that each specify a different byte range to GET.

raw email A type of *sendmail* request with which you can specify the email headers and MIME types.

Amazon RDS Amazon Relational Database Service is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks.
See Also <https://aws.amazon.com/rds>.

read local/write global An [active-active](#) strategy in which all writes for a workload are sent to one primary Region and all read traffic is served from the Region where the request originates. Typically architected with an asynchronous data store. Sometimes referred to as *read local-write global*.
See Also , .

read local/write local An [active-active](#) strategy in which all writes for a workload are sent to one primary Region and all read traffic is served from the Region where the request originates. Typically architected with an asynchronous data store. Sometimes referred to as *read local-write global*.
See Also , .

read replica	Amazon RDS : An active copy of another DB instance. Any updates to the data on the source DB instance are replicated to the read replica DB instance using the built-in replication feature of MySQL 5.1.
real-time predictions	Amazon Machine Learning: Synchronously generated predictions for individual data observations. See Also batch prediction .
receipt handle	Amazon SQS : An identifier that you get when you receive a message from the queue. This identifier is required to delete a message from the queue or when changing a message's visibility timeout.
receiver	The entity that consists of the network systems, software, and policies that manage email delivery for a recipient .
recipe	Amazon Personalize : An Amazon Personalize algorithm that's preconfigured to predict the items that a user interacts with (for <code>USER_PERSONALIZATION</code> recipes), or calculate items that are similar to specific items that a user has shown interest in (for <code>RELATED_ITEMS</code> recipes), or rank a collection of items that you provide based on the predicted interest for a specific user (for <code>PERSONALIZED_RANKING</code> recipes). See Also USER_PERSONALIZATION recipes , RELATED_ITEMS recipes , PERSONALIZED_RANKING recipes .
recipient	Amazon SES : The person or entity receiving an email message. For example, a person named in the "To" field of a message.
recommendations	Amazon Personalize : A list of items that Amazon Personalize predicts that a user interacts with. Depending on the Amazon Personalize recipe used, recommendations can be either a list of items (with <code>USER_PERSONALIZATION</code> recipes and <code>RELATED_ITEMS</code> recipes), or a ranking of a collection of items you provided (with <code>PERSONALIZED_RANKING</code> recipes). See Also recipe , campaign , solution version , USER_PERSONALIZATION recipes , RELATED_ITEMS recipes , PERSONALIZED_RANKING recipes .
Redis	A fast, open-source, in-memory key-value data structure store. Redis comes with a set of versatile in-memory data structures with which you can easily create a variety of custom applications.

Amazon Redshift	Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. With Amazon Redshift, you can analyze your data using your existing business intelligence tools. See Also https://aws.amazon.com/redshift/ .
reference	A means of inserting a property from one AWS resource into another. For example, you could insert an Amazon EC2 security group property into an Amazon RDS resource.
Region	A named set of AWS resources that's in the same geographical area. A Region is comprised of at least three Availability Zones . AWS Regions are divided into partitions . AWS commercial Regions are in the AWS partition, China Regions are in the AWS-cn partition, and AWS GovCloud (US) Regions are in the AWS-us-gov partition.
Region that is enabled by default	An AWS Region that is enabled by default. Regions that were introduced before March 20, 2019 are enabled by default and can't be disabled. For a list of Regions that aren't enabled by default (opt-in Region), see Considerations before enabling and disabling Regions in the <i>AWS Account Management Guide</i> .
regression model	Amazon Machine Learning: Preformatted instructions for common data transformations that fine-tune machine learning model performance.
regression model	A type of machine learning model that predicts a numeric value, such as the exact purchase price of a house.
regularization	A machine learning (ML) parameter that you can tune to obtain higher-quality ML models. Regularization helps prevent ML models from memorizing training data examples instead of learning how to generalize the patterns it sees (called overfitting). When training data is overfitted, the ML model performs well on the training data, but doesn't perform well on the evaluation data or on new data.
Amazon Rekognition	Amazon Rekognition is a machine learning service that identifies objects, people, text, scenes, and activities, including inappropriate content, in either image or video files. With Amazon Rekognition Custom Labels, you can create a customized ML model that detects objects and scenes specific to your business in images.

	See Also https://aws.amazon.com/rekognition/ .
RELATED_ITEMS recipes	Amazon Personalize Recipes that recommend items that are similar to a specified item, such as the item-to-item (SIMS) recipe. See Also recipe , item-to-item similarities (SIMS) recipe .
replacement environment	The instances in a deployment group after the CodeDeploy blue/green deployment.
replica shard	See shard .
reply path	The email address that an email reply is sent to. This is different from the return path .
representational state transfer	See REST .
reputation	<ol style="list-style-type: none">1. An Amazon SES metric, based on factors that might include bounces, complaints, and other metrics, regarding whether a customer is sending high-quality email.2. A measure of confidence, as judged by an internet service provider (ISP) or other entity that an IP address that they are receiving email from isn't the source of spam.
requester	The person (or application) that sends a request to AWS to perform a specific action. When AWS receives a request, it first evaluates the requester's permissions to determine whether the requester is allowed to perform the request action (if applicable, for the requested resource).
Requester Pays	An Amazon S3 feature that allows a bucket owner to specify that anyone who requests access to objects in a particular bucket must pay the data transfer and request costs.
reservation	A collection of EC2 instances started as part of the same launch request. This is not to be confused with a Reserved Instance .
Reserved Instance	A pricing option for EC2 instances that discounts the on-demand usage charge for instances that meet the specified parameters. Customers pay for the entire term of the instance, regardless of how they use it.

Reserved Instance Marketplace	An online exchange that matches sellers who have reserved capacity that they no longer need with buyers who are looking to purchase additional capacity. reserved instances that you purchase from third-party sellers have less than a full standard term remaining and can be sold at different upfront prices. The usage or reoccurring fees remain the same as the fees set when the Reserved Instances were originally purchased. Full standard terms for Reserved Instances available from AWS run for one year or three years.
Resilience Hub	AWS Resilience Hub gives you a central place to define, validate, and track the resiliency of your AWS application. It helps you to protect your applications from disruptions, and reduce recovery costs to optimize business continuity to help meet compliance and regulatory requirements. See Also https://aws.amazon.com/resilience-hub .
resource	An entity that users can work with in AWS, such as an EC2 instance , an DynamoDB table, an Amazon S3 bucket , an IAM user, or an OpsWorks stack .
Resource Groups	AWS Resource Groups is a web service that AWS customers can use to manage and automate tasks on large numbers of resources at one time. See Also AWS Resource Groups .
Amazon Resource Name (ARN)	Amazon Resource Name is a standardized way to refer to an AWS resource (for example, <code>arn:aws:iam::123456789012:user/division_abc/subdivision_xyz/Bob</code>).
resource property	A value required when including an AWS resource in an CloudFormation stack . Each resource can have one or more properties associated with it. For example, an <code>AWS::EC2::Instance</code> resource might have a <code>UserData</code> property. In an AWS CloudFormation template, resources must declare a properties section, even if the resource has no properties.
resource record	Also called <i>resource record set</i> . The fundamental information elements in the Domain Name System (DNS). See Also Domain Name System on Wikipedia.
REST	Representational state transfer. A simple stateless architecture that generally runs over HTTPS/TLS. REST emphasizes that resources have

unique and hierarchical identifiers (URIs), are represented by common media types (such as HTML, XML, or [JSON](#)), and that operations on the resources are either predefined or discoverable within the media type. In practice, this generally results in a limited number of operations. See Also [Query](#), [WSDL](#), [SOAP](#).

RESTful web service	Also known as RESTful API. A web service that follows REST architectural constraints. The API operations must use HTTP methods explicitly, expose hierarchical URIs, and transfer either XML, JSON , or both.
return enabled	CloudSearch : An index field option that enables the field's values to be returned in the search results.
return path	The email address that bounced email is returned to. The return path is specified in the header of the original email. This is different from the reply path .
revision	CodePipeline : A change that's made to a source that's configured in a source action, such as a pushed commit to a GitHub repository or an update to a file in a versioned Amazon S3 bucket .
AWS RoboMaker	AWS RoboMaker is a cloud-based simulation service that robotics developers use to run, scale, and automate simulation without managing any infrastructure. See Also https://aws.amazon.com/robomaker .
role	A tool for giving temporary access to AWS resources in your AWS account .
rollback	A return to a previous state that follows the failure to create an object, such as CloudFormation stack . All resources that are associated with the failure are deleted during the rollback. For AWS CloudFormation, you can override this behavior using the <code>--disable-rollback</code> option on the command line.
root	Organizations : A parent container for the accounts in your organization. If you apply a service control policy to the root, it applies to every organizational unit and account in the organization.
root credentials	Authentication information associated with the AWS account owner.

root device volume	A volume that contains the image used to boot the instance (also known as a <i>root device</i>). If you launched the instance from an AMI backed by instance store , this is an instance store volume created from a template stored in Amazon S3 . If you launched the instance from an AMI backed by Amazon EBS , this is an Amazon EBS volume created from an Amazon EBS snapshot.
route table	A set of routing rules that controls the traffic leaving any subnet that's associated with the route table. You can associate multiple subnets with a single route table, but a subnet can be associated with only one route table at a time.
Route 53	Amazon Route 53 is a web service that you can use to create a new DNS service or to migrate your existing DNS service to the cloud. See Also https://aws.amazon.com/route53 .
row identifier	Amazon Machine Learning: An attribute in the input data that you can include in the evaluation or prediction output to make it easier to associate a prediction with an observation.
rule	AWS WAF : A set of conditions that AWS WAF searches for in web requests to AWS resources such as Amazon CloudFront distributions. You add rules to a web ACL , and then specify whether you want to allow or block web requests based on each rule.

S

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

Amazon S3	Amazon S3 is storage for the internet. You can use it to store and retrieve any amount of data at any time, from anywhere on the web. See Also https://aws.amazon.com/s3 .
Amazon S3 Glacier	Amazon S3 Glacier is a secure, durable, and low-cost storage service for data archiving and long-term backup. You can reliably store large or small amounts of data for significantly less than on-premises solutions. S3 Glacier is optimized for infrequently accessed data, where a retrieval time of several hours is suitable.

	See Also https://aws.amazon.com/glacier/ .
Amazon S3-Backed AMI	See instance store-backed AMI .
SageMaker AI	Amazon SageMaker AI is a fully managed cloud service that builds, trains, and deploys machine learning (ML) models by using AWS infrastructure, tools, and workflows. See Also https://aws.amazon.com/sagemaker .
AWS SAM	AWS Serverless Application Model is an open-source framework for building and running serverless applications. AWS SAM provides a command line interface tool and a shorthand syntax template specification that you can use to quickly iterate through your serverless application lifecycle. See Also https://aws.amazon.com/serverless/sam/ .
sampling period	A defined duration of time, such as one minute, which CloudWatch computes a statistic over.
sandbox	A testing location where you can test the functionality of your application without affecting production, incurring charges, or purchasing products. Amazon SES : An environment that developers can use to test and evaluate the service. In the sandbox, you have full access to the Amazon SES API, but you can only send messages to verified email addresses and the mailbox simulator. To get out of the sandbox, you must apply for production access. Accounts in the sandbox also have lower sending limits than production accounts.
scale in	To remove EC2 instances from an Auto Scaling group .
scale out	To add EC2 instances to an Auto Scaling group .
scaling activity	A process that changes the size, configuration, or makeup of an Auto Scaling group by launching or terminating instances.
scaling policy	A description of how Auto Scaling automatically scales an Auto Scaling group in response to changing demand.

	See Also scale in , scale out .
scheduler	The method used for placing tasks on container instances .
schema	Amazon Machine Learning: The information that's needed to interpret the input data for a machine learning model, including attribute names and their assigned data types, and the names of special attributes.
score cut-off value	Amazon Machine Learning: A binary classification model outputs a score that ranges from 0 to 1. To decide whether an observation is classified as 1 or 0, you pick a classification threshold, or cut-off, and Amazon ML compares the score against it. Observations with scores higher than the cut-off are predicted as target equals 1, and scores lower than the cut-off are predicted as target equals 0.
SCP	See service control policy .
AWS SCT	AWS Schema Conversion Tool is a desktop application that automates heterogeneous database migrations. You can use AWS SCT to convert database schemas and code objects, SQL code in your applications, and ETL scripts to a format compatible with the target database. Then, you can use AWS SCT data extraction agents to migrate data to your target database. See Also https://aws.amazon.com/dms/schema-conversion-tool .
AWS SDK for .NET	AWS SDK for .NET is a software development kit that provides .NET API operations for AWS services including Amazon S3 , Amazon EC2 , IAM , and more. You can download the SDK as multiple service-specific packages on NuGet. See Also https://aws.amazon.com/sdk-for-net/ .
SDK for C++	AWS SDK for C++ is a software development kit that provides C++ APIs for many AWS services including Amazon S3 , Amazon EC2 , DynamoDB , and more. The single, downloadable package includes the AWS C++ library, code examples, and documentation. See Also https://aws.amazon.com/sdk-for-cpp/ .
SDK for Go	AWS SDK for Go is a software development kit for integrating your Go application with the full suite of AWS services. See Also https://aws.amazon.com/sdk-for-go/ .

SDK for Java	<p>AWS SDK for Java is a software development kit that provides Java API operations for many AWS services including Amazon S3, Amazon EC2, DynamoDB, and more. The single, downloadable package includes the AWS Java library, code examples, and documentation.</p> <p>See Also https://aws.amazon.com/sdk-for-java/.</p>
SDK for JavaScript in Node.js	<p>AWS SDK for JavaScript in Node.js is a software development kit for accessing AWS services from JavaScript in Node.js. The SDK provides JavaScript objects for AWS services, including Amazon S3, Amazon EC2, DynamoDB, and Amazon SWF. The single, downloadable package includes the AWS JavaScript library and documentation.</p> <p>See Also https://docs.aws.amazon.com/sdk-for-javascript/v2/developer-guide/.</p>
SDK for JavaScript in the Browser	<p>AWS SDK for JavaScript in the Browser is a software development kit for accessing AWS services from JavaScript code running in the browser. Authenticate users through Facebook, Google, or Login with Amazon using web identity federation. Store application data in DynamoDB, and save user files to Amazon S3.</p> <p>See Also https://docs.aws.amazon.com/sdk-for-javascript/v2/developer-guide/.</p>
SDK for PHP	<p>AWS SDK for PHP is a software development kit and open-source PHP library for integrating your PHP application with AWS services such as Amazon S3, Amazon S3 Glacier, and DynamoDB.</p> <p>See Also https://aws.amazon.com/sdk-for-php/.</p>
SDK for Python (Boto3)	<p>AWS SDK for Python (Boto3) is a software development kit for using Python to access AWS services such as Amazon EC2, Amazon EMR, Amazon EC2 Auto Scaling, Kinesis, or Lambda.</p> <p>See Also http://boto.readthedocs.org/en/latest/.</p>
SDK for Ruby	<p>AWS SDK for Ruby is a software development kit for accessing AWS services from Ruby. The SDK provides Ruby classes for many AWS services including Amazon S3, Amazon EC2, DynamoDB and more. The single, downloadable package includes the AWS Ruby Library and documentation.</p> <p>See Also https://aws.amazon.com/sdk-for-ruby/.</p>

SDK for Rust	AWS SDK for Rust is a software development kit that provides APIs and utilities for developers. It enables Rust applications to integrate with AWS services such as Amazon S3 and Amazon EC2.
SDK for Swift	AWS SDK for Swift is a software development kit that provides support for accessing AWS infrastructure and services using the Swift language.
search API	CloudSearch : The API that you use to submit search requests to a search domain .
search domain	CloudSearch : Encapsulates your searchable data and the search instances that handle your search requests. You typically set up a separate Amazon CloudSearch domain for each different collection of data that you want to search.
search domain configuration	CloudSearch : A domain's indexing options, analysis schemes , expressions , suggesters , access policies, and scaling and availability options.
search enabled	CloudSearch : An index field option that enables the field data to be searched.
search endpoint	CloudSearch : The URL that you connect to when sending search requests to a search domain. Each Amazon CloudSearch domain has a unique search endpoint that remains the same for the life of the domain.
search index	CloudSearch : A representation of your searchable data that facilitates fast and accurate data retrieval.
search instance	CloudSearch : A compute resource that indexes your data and processes search requests. An Amazon CloudSearch domain has one or more search instances, each with a finite amount of RAM and CPU resources. As your data volume grows, more search instances or larger search instances are deployed to contain your indexed data. When necessary, your index is automatically partitioned across multiple search instances. As your request volume or complexity increases, each search partition is automatically replicated to provide additional processing capacity.
search request	CloudSearch : A request that's sent to an Amazon CloudSearch domain's search endpoint to retrieve documents from the index that match particular search criteria.

search result	CloudSearch : A document that matches a search request. Also referred to as a <i>search hit</i> .
secret access key	A key that's used with the access key ID to cryptographically sign programmatic AWS requests. Signing a request identifies the sender and prevents the request from being altered. You can generate secret access keys for your AWS account , individual IAM users and temporary sessions.
Secrets Manager	AWS Secrets Manager is a service for securely encrypting, storing, and rotating credentials for databases and other services. See Also https://aws.amazon.com/secrets-manager/ .
security group	A named set of allowed inbound network connections for an instance. (Security groups in Amazon VPC also include support for outbound connections.) Each security group consists of a list of protocols, ports, and IP address ranges. A security group can apply to multiple instances, and multiple groups can regulate a single instance.
Security Hub	AWS Security Hub is a service that provides a comprehensive view of the security state of your AWS resources. Security Hub collects security data from AWS accounts and services and helps you analyze your security trends to identify and prioritize the security issues across your AWS environment. See Also https://aws.amazon.com/security-hub/ .
sender	The person or entity sending an email message.
Sender ID	A Microsoft controlled version of SPF . An email authentication and anti-spoofing system. For more information about Sender ID, see Sender ID in Wikipedia.
sending limits	The sending quota and maximum send rate that are associated with every Amazon SES account.
sending quota	The maximum number of email messages that you can send using Amazon SES in a 24-hour period.
AWS Serverless Application Repository	AWS Serverless Application Repository is a managed repository that teams, organizations, and individual developers can use to store and

	<p>share reusable applications, and assemble and deploy serverless architectures in powerful new ways.</p> <p>See Also https://aws.amazon.com/serverless/serverlessrepo/.</p>
server-side encryption (SSE)	<p>The encrypting of data at the server level. Amazon S3 supports three modes of server-side encryption: SSE-S3, where Amazon S3 manages the keys; SSE-C, where the customer manages the keys; and SSE-KMS, where AWS KMS manages keys.</p>
Service Catalog	<p>AWS Service Catalog is a web service that helps organizations create and manage catalogs of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multitier application architectures.</p> <p>See Also https://aws.amazon.com/servicecatalog/.</p>
service control policy	<p>Organizations: A policy-based control that specifies the services and actions that users and roles can use in the accounts that the service control policy (SCP) affects.</p>
service endpoint	<p>See endpoint.</p>
service health dashboard	<p>A webpage showing up-to-the-minute information about AWS service availability. The dashboard is located at http://status.aws.amazon.com/.</p>
AWS Service Management Connector	<p>AWS Service Management Connector enables customers to provision, manage, and operate AWS resources and capabilities in familiar IT Service Management (ITSM) tooling.</p> <p>See Also https://aws.amazon.com/service-management-connector.</p>
Service Quotas	<p>A service for viewing and managing your quotas easily and at scale as your AWS workloads grow. Quotas, also referred to as limits, are the maximum number of resources that you can create in an AWS account.</p>
service role	<p>An IAM role that grants permissions to an AWS service so it can access AWS resources. The policies that you attach to the service role determine which AWS resources the service can access and what it can do with those resources.</p>
Amazon SES	<p>Amazon Simple Email Service is a simple and cost-effective email solution for applications.</p>

	See Also https://aws.amazon.com/ses .
session	The period when the temporary security credentials that are provided by AWS STS allow access to your AWS account.
SHA	Secure Hash Algorithm. SHA1 is an earlier version of the algorithm, which AWS has replaced with SHA256.
shard	OpenSearch Service : A partition of data in an index. You can split an index into multiple shards, which can include primary shards (original shards) and replica shards (copies of the primary shards). Replica shards provide failover. This means that, if a cluster node that contains a primary shard fails, a replica shard is promoted to a primary shard. Replica shards also can handle requests.
shared AMI	An Amazon Machine Image (AMI) that a developer builds and makes available for others to use.
Shield	AWS Shield is a service that helps to protect your resources—such as Amazon EC2 instances, Elastic Load Balancing load balancers, Amazon CloudFront distributions, and Route 53 hosted zones—against DDoS attacks. AWS Shield is automatically included at no extra cost beyond what you already pay for AWS WAF and your other AWS services. For added protection against DDoS attacks, AWS offers AWS Shield Advanced. See Also https://aws.amazon.com/shield .
shutdown action	Amazon EMR : A predefined bootstrap action that launches a script that runs a series of commands in parallel before terminating the job flow.
signature	Refers to a <i>digital signature</i> , which is a mathematical way to confirm the authenticity of a digital message. AWS uses signatures to authenticate the requests you send to our web services. For more information, to https://aws.amazon.com/security .
SIGNATURE file	Import/Export : A file that you copy to the root directory of your storage device. The file contains a job ID, manifest file, and a signature.
Signature Version 4	Protocol for authenticating inbound API requests to AWS services in all AWS Regions.

Signer	AWS Signer is a fully managed code-signing service used to ensure the authenticity and integrity of an AWS customer's code.
Silk	Amazon Silk is a next-generation web browser that's available only on Fire OS tablets and phones. Built on a split architecture that divides processing between the client and the AWS Cloud, Amazon Silk creates a faster, more responsive mobile browsing experience.
Simple Mail Transfer Protocol	See SMTP .
Simple Object Access Protocol	See SOAP .
SimSpace Weaver	AWS SimSpace Weaver is a managed service that you can use to build and run large-scale spatial simulations in the AWS Cloud. See Also https://aws.amazon.com/simspaceweaver/ .
SIMS recipe	See item-to-item similarities (SIMS) recipe .
single sign-on	An authentication scheme that allows users to sign in one time to access multiple applications and websites. The service name AWS Single Sign-On is now AWS IAM Identity Center. See Also IAM Identity Center .
Single-AZ DB instance	A standard (non-Multi-AZ) DB instance that's deployed in one Availability Zone , without a standby replica in another Availability Zone. See Also Multi-AZ deployment .
Site-to-Site VPN	AWS Site-to-Site VPN is a fully managed service that you can use to establish Internet Protocol security (IPsec) VPN connections between your AWS networks and your on-premises networks. See Also https://aws.amazon.com/vpn/site-to-site-vpn .
sloppy phrase search	A search for a phrase that specifies how close the terms must be to one another to be considered a match.
AWS SMS	AWS Server Migration Service is a service that combines data collection tools with automated server replication to speed the migration of on-premises servers to AWS.

	See Also https://aws.amazon.com/server-migration-service .
SMTP	Simple Mail Transfer Protocol. The standard that's used to exchange email messages between internet hosts for the purpose of routing and delivery.
snapshot	Amazon EBS : A backup of your volumes that's stored in Amazon S3 . You can use these snapshots as the starting point for new Amazon EBS volumes or to protect your data for long-term durability. See Also DB snapshot .
Snowball	AWS Snowball is a petabyte-scale data transport solution that uses devices that are secure to transfer large amounts of data into and out of the AWS Cloud. See Also https://aws.amazon.com/snowball .
Amazon SNS	Amazon Simple Notification Service is a web service that applications, users, and devices can use to instantly send and receive notifications from the cloud. See Also https://aws.amazon.com/sns .
SOAP	Simple Object Access Protocol. An XML-based protocol that you can use to exchange information over a particular protocol (for example, HTTP or SMTP) between applications. See Also REST , WSDL .
soft bounce	A temporary email delivery failure such as one resulting from a full mailbox.
software VPN	A software appliance-based VPN connection over the internet.
solution	Amazon Personalize : The recipe, customized parameters, and trained models (solution versions) that can be used to generate recommendations. See Also recipe , solution version , recommendations .
solution version	Amazon Personalize : A trained model that you create as part of a solution in Amazon Personalize. You deploy a solution version in a campaign to generate recommendations. See Also solution , campaign , recommendations .

sort enabled	CloudSearch : An index field option that enables a field to be used to sort the search results.
sort key	An attribute used to sort the order of partition keys in a composite primary key (also known as a <i>range attribute</i>). See Also partition key , primary key .
source/destination checking	A security measure to verify that an EC2 instance is the origin of all traffic that it sends and the ultimate destination of all traffic that it receives. In other words, this measure verifies that the instance isn't relaying traffic. By default, source/destination checking is turned on. For instances that function as gateways, such as Amazon VPC NAT instances, source/destination checking must be disabled.
spam	Unsolicited bulk emails.
spamtrap	An email address that's set up by an anti- spam entity. This email address isn't for correspondence but rather for monitoring unsolicited emails. This is also called a <i>honeypot</i> .
SPF	Sender Policy Framework. A standard for authenticating email.
SPICE	A robust in-memory engine that is part of Amazon QuickSight . Engineered for the cloud, SPICE (Super-fast, Parallel, In-memory Calculation Engine) uses a combination of storage and in-memory technologies. It uses these to get faster results from interactive queries and advanced calculations on large datasets. SPICE automatically replicates data for high availability. SPICE makes it possible for Amazon QuickSight to support hundreds of thousands of simultaneous analyses across a variety of data sources.
Spot Instance	A type of EC2 instance that you can bid on to use unused Amazon EC2 capacity.
Spot price	The price for a Spot Instance at any given time. If your maximum price exceeds the current price and your restrictions are met, Amazon EC2 launches instances on your behalf.
SQL injection match condition	AWS WAF : An attribute that specifies the part of web requests (such as a header or a query string) that AWS WAF inspects for malicious SQL

	<p>code. Based on the specified conditions, you can configure AWS WAF to allow or block web requests to an AWS resource, such as an Amazon CloudFront distribution.</p>
Amazon SQS	<p>Amazon Simple Queue Service is a reliable and scalable hosted queues for storing messages as they travel between computers.</p> <p>See Also https://aws.amazon.com/sqs.</p>
Amazon SWF	<p>Amazon Simple Workflow Service is a fully managed service that helps developers build, run, and scale background jobs that have parallel or sequential steps. Amazon SWF functions similar to a state tracker and task coordinator in the AWS Cloud.</p> <p>See Also https://aws.amazon.com/swf/.</p>
SSE	<p>See server-side encryption (SSE).</p>
SSL	<p>Secure Sockets Layer</p> <p>See Also Transport Layer Security (TLS).</p>
stack	<p>CloudFormation: A collection of AWS resources that you create and delete as a single unit.</p> <p>OpsWorks: A set of instances that you manage collectively, typically because they have a common purpose such as serving PHP applications. A stack serves as a container and handles tasks that apply to the group of instances as a whole, such as managing applications and cookbooks.</p>
station	<p>CodePipeline: A portion of a pipeline workflow where one or more actions are performed.</p>
station	<p>A place at an AWS facility where your AWS Import/Export data is transferred on to, or off of, your storage device.</p>
statistic	<p>One of five functions of the values submitted for a given sampling period. These functions are Maximum, Minimum, Sum, Average, and SampleCount.</p>
stem	<p>The common root or substring shared by a set of related words.</p>
stemming	<p>The process of mapping related words to a common stem. This enables matching on variants of a word. For example, a search for "horse" could</p>

	<p>return matches for horses, horseback, and horsing, as well as horse. CloudSearch supports both dictionary based and algorithmic stemming.</p>
step	<p>Amazon EMR: A single function applied to the data in a job flow. The sum of all steps comprises a job flow.</p>
Step Functions	<p>AWS Step Functions is a web service that coordinates the components of distributed applications as a series of steps in a visual workflow. See Also https://aws.amazon.com/step-functions/.</p>
step type	<p>Amazon EMR: The type of work done in a step. There are a limited number of step types, such as moving data from Amazon S3 to Amazon EC2 or from Amazon EC2 to Amazon S3.</p>
sticky session	<p>A feature of the ELB load balancer that binds a user's session to a specific application instance. This is so that all requests that are coming from the user during the session are sent to the same application instance. By contrast, a load balancer defaults to route each request independently to the application instance with the smallest load.</p>
stopping	<p>The process of filtering stop words from an index or search request.</p>
stopword	<p>A word that isn't indexed and is automatically filtered out of search requests because it's either insignificant or so common that including it results in too many matches to be useful. Stopwords are language specific.</p>
Storage Gateway	<p>AWS Storage Gateway is a hybrid cloud storage service that provides on-premises access to virtually unlimited cloud storage. See Also AWS Storage Gateway.</p>
streaming	<p>Amazon EMR: A utility that comes with Hadoop that you can use to develop MapReduce executables in languages other than Java.</p> <p>CloudFront: The ability to use a media file in real time—as it's transmitted in a steady stream from a server.</p>
streaming distribution	<p>A special kind of distribution that serves streamed media files using a Real Time Messaging Protocol (RTMP) connection.</p>
Streams	<p>See Kinesis Data Streams.</p>

string match condition	AWS WAF : An attribute that specifies the strings that AWS WAF searches for in a web request, such as a value in a header or a query string. Based on the specified strings, you can configure AWS WAF to allow or block web requests to an AWS resource , such as a CloudFront distribution.
string-to-sign	Before you calculate an HMAC signature, you first assemble the required components in a canonical order. The preencrypted string is the string-to-sign.
strongly consistent read	A read process that returns a response with the most up-to-date data. This data reflects the updates from all previous write operations that were successful—regardless of the Region. See Also data consistency , eventual consistency , eventually consistent read .
structured query	Search criteria that are specified using the CloudSearch structured query language. You use the structured query language to construct compound queries that use advanced search options and combine multiple search criteria using Boolean operators.
AWS STS	AWS Security Token Service is a web service for requesting temporary, limited-privilege credentials for IAM users or for users that you authenticate (federated users). See Also https://aws.amazon.com/iam/ .
subnet	A segment of the IP address range of a Amazon VPC that an EC2 instance can be attached to. You can create subnets to group instances according to security and operational needs.
Subscription button	An HTML-coded button that provides a simple way to charge customers a recurring fee.
suggester	CloudSearch : Specifies an index field for getting autocomplete suggestions and options that can enable fuzzy matches and control how suggestions are sorted.
suggestions	Documents that contain a match for the partial search string in the field that's designated by the suggester . CloudSearch suggestions include the document IDs and field values for each matching document. To be a

match, the string must match the contents of the field starting from the beginning of the field.

Sumerian Amazon Sumerian is a set of tools for creating and running high-quality 3D, augmented reality (AR), and virtual reality (VR) applications on the web.

See Also <https://aws.amazon.com/sumerian/>.

supported AMI An [Amazon Machine Image \(AMI\)](#) similar to a [paid AMI](#), except that the owner charges for additional software or a service that customers use with their own AMIs.

SWF See [Amazon SWF](#).

symmetric encryption [Encryption](#) that uses a private key only.
See Also [asymmetric encryption](#).

synchronous bounce A type of [bounce](#) that occurs while the email servers of the [sender](#) and [receiver](#) are actively communicating.

synonym A word that's the same or nearly the same as an indexed word and that likely produces the same results when specified in a search request. For example, a search for "Rocky Four" or "Rocky 4" likely returns the fourth *Rocky* movie. You can do this by designating that four and 4 are synonyms for IV. Synonyms are language specific.

Systems Manager AWS Systems Manager is the operations hub for AWS and hybrid cloud environments that can help achieve secure operations at scale. It provides a unified user interface for users to view operations data from multiple AWS services and automate tasks across their AWS resources.
See Also <https://aws.amazon.com/systems-manager>.

T

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

table A collection of data. Similar to other database systems, DynamoDB stores data in tables.

tag	Metadata that you can define and assign to AWS resources , such as an EC2 instance . Not all AWS resources can be tagged.
tagging	Tagging resources: Applying a tag to an AWS resource . Amazon SES : Also called <i>labeling</i> . A way to format return path email addresses so that you can specify a different return path for each recipient of a message. You can use tagging to support VERP . For example, if Andrew manages a mailing list, he can use the return paths <code>andrew+recipient1@example.net</code> and <code>andrew+recipient2@example.net</code> so that he can determine which email bounced.
target attribute	Amazon Machine Learning (Amazon ML): The attribute in the input data that contains the “correct” answers. Amazon ML uses the target attribute to learn how to make predictions on new data. For example, if you were building a model for predicting the sale price of a house, the target attribute would be “target sale price in USD.”
target revision	CodeDeploy : The most recent version of the application revision that has been uploaded to the repository and will be deployed to the instances in a deployment group. In other words, the application revision currently targeted for deployment. This is also the revision that will be pulled for automatic deployments.
task	An instantiation of a task definition that's running on a container instance .
task definition	The blueprint for your task. Specifies the name of the task , revisions, container definitions , and volume information.
task node	An EC2 instance that runs Hadoop map and reduce tasks, but doesn't store data. Task nodes are managed by the master node , which assigns Hadoop tasks to nodes and monitors their status. While a job flow is running, you can increase and decrease the number of task nodes. Because they don't store data and can be added and removed from a job flow, you can use task nodes to manage the EC2 instance capacity your job flow uses, increasing capacity to handle peak loads and decreasing it later.

	Task nodes only run a TaskTracker Hadoop daemon.
tebibyte (TiB)	A contraction of tera binary byte. A tebibyte (TiB) is 2^{40} or 1,099,511,627,776 bytes. A terabyte (TB) is 10^{12} or 1,000,000,000,000 bytes. 1,024 TiB is a pebibyte (PiB) .
template format version	The version of an CloudFormation template design that determines the available features. If you omit the <code>AWSTemplateFormatVersion</code> section from your template, AWS CloudFormation assumes the most recent format version.
template validation	The process of confirming the use of JSON code in an CloudFormation template. You can validate any AWS CloudFormation template using the <code>cfn-validate-template</code> command.
temporary security credentials	Authentication information that's provided by AWS STS when you call an STS API action. Includes an access key ID , a secret access key , a session token , and an expiration time.
Amazon Textract	Amazon Textract is a service that automatically extracts text and data from scanned documents. Amazon Textract goes beyond simple optical character recognition (OCR) to also identify the contents of fields in forms and information stored in tables. See Also https://aws.amazon.com/textract/ .
throttling	The automatic restricting or slowing down of a process based on one or more limits. For example, Kinesis Data Streams throttles operations if an application (or group of applications operating on the same stream) attempts to get data from a shard at a rate faster than the shard limit. API Gateway uses throttling to limit the steady-state request rates for a single account. Amazon SES uses throttling to reject attempts to send email that exceeds the sending limits .
time-series data	Data that's provided as part of a metric. The time value is assumed to be when the value occurred. A metric is the fundamental concept for CloudWatch and represents a time-ordered set of data points. You publish metric data points into CloudWatch and later retrieve statistics about those data points as a time-series ordered dataset.

timestamp	A date/time string in the ISO 8601 format (more specifically, in the YYYY-MM-DD format).
Timestream	Amazon Timestream is a scalable and serverless time series database service for real-time analytics, DevOps, and IoT applications that you can use to store and analyze trillions of events per day. See Also https://aws.amazon.com/timestream .
TLS	See Transport Layer Security (TLS) .
tokenization	The process of splitting a stream of text into separate tokens on detectable boundaries such as white space and hyphens.
AWS Toolkit for Eclipse	AWS Toolkit for Eclipse is an open-source plugin for the Eclipse Java integrated development environment (IDE) that makes it easier to develop, debug, and deploy Java applications using Amazon Web Services. See Also https://aws.amazon.com/eclipse/ .
AWS Toolkit for JetBrains	AWS Toolkit for JetBrains is an open-source plugin for the integrated development environments (IDEs) from JetBrains that makes it easier to develop, debug, and deploy serverless applications using Amazon Web Services. See Also https://aws.amazon.com/intellij/ , https://aws.amazon.com/pycharm/ .
AWS Toolkit for Microsoft Azure DevOps	AWS Toolkit for Microsoft Azure DevOps provides tasks you can use in build and release definitions in VSTS to interact with AWS services. See Also https://aws.amazon.com/vsts/ .
AWS Toolkit for Visual Studio	AWS Toolkit for Visual Studio is an extension for Visual Studio that helps in developing, debugging, and deploying .NET applications using Amazon Web Services. See Also https://aws.amazon.com/visualstudio/ .
AWS Toolkit for Visual Studio Code	AWS Toolkit for Visual Studio Code is an open-source plugin for the Visual Studio Code (VS Code) editor that makes it easier to develop, debug, and deploy applications using Amazon Web Services. See Also https://aws.amazon.com/visualstudiocode/ .

AWS Tools for PowerShell	AWS Tools for PowerShell is a set of PowerShell cmdlets to help developers and administrators manage their AWS services from the PowerShell scripting environment. See Also https://aws.amazon.com/powershell/ .
topic	A communication channel to send messages and subscribe to notifications. It provides an access point for publishers and subscribers to communicate with each other.
Traffic Mirroring	An Amazon VPC feature that you can use to copy network traffic from an elastic network interface of Amazon EC2 instances. You can then send this network traffic to out-of-band security and monitoring appliances for content inspection, threat monitoring, and troubleshooting. See Also https://aws.amazon.com/vpc/ .
training datasource	A datasource that contains the data that Amazon Machine Learning uses to train the machine learning model to make predictions.
Amazon Transcribe	Amazon Transcribe is a machine learning service that uses automatic speech recognition (ASR) to quickly and accurately convert speech to text. See Also https://aws.amazon.com/transcribe/ .
Amazon Transcribe Medical	Amazon Transcribe Medical is an automatic speech recognition (ASR) service for adding medical speech-to-text capabilities to voice-enabled clinical documentation applications. See Also https://aws.amazon.com/transcribe/medical/ .
Transfer Family	AWS Transfer Family offers fully managed support for transferring files over SFTP, FTPS, and FTP into and out of Amazon S3 or Amazon EFS, as well as support for the Applicability Statement 2 (AS2) protocol for business-to-business (B2B) transfers. See Also https://aws.amazon.com/aws-transfer-family .
transition	CodePipeline : The act of a revision in a pipeline continuing from one stage to the next in a workflow.
Amazon Translate	Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. See Also https://aws.amazon.com/translate/ .

Transport Layer Security (TLS)	A cryptographic protocol that provides security for communication over the internet. Its predecessor is Secure Sockets Layer (SSL).
trust policy	An IAM policy that's an inherent part of an IAM role . The trust policy specifies which principals are allowed to use the role.
Trusted Advisor	AWS Trusted Advisor is a web service that inspects your AWS environment and makes recommendations for saving money, improving system availability and performance, and helping to close security gaps. See Also https://aws.amazon.com/premiumsupport/trustedadvisor/ .
trusted key groups	Amazon CloudFront key groups whose public keys CloudFront can use to verify the signatures of CloudFront signed URLs and signed cookies .
trusted signers	See trusted key groups .
tuning	Selecting the number and type of AMIs to run a Hadoop job flow most efficiently.
tunnel	A route for transmission of private network traffic that uses the internet to connect nodes in the private network. The tunnel uses encryption and secure protocols such as PPTP to prevent the traffic from being intercepted as it passes through public routing nodes.

U

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

unbounded	The number of potential occurrences isn't limited by a set number. This value is often used when defining a data type that's a list (for example, <code>maxOccurs="unbounded"</code>), in WSDL .
unit	Standard measurement for the values submitted to CloudWatch as metric data. Units include seconds, percent, bytes, bits, count, bytes/second, bits/second, count/second, and none.
usage report	An AWS record that details your usage of a particular AWS service. You can generate and download usage reports from https://aws.amazon.com/usage-reports/ .

user	A person or application under an account that makes API calls to AWS products. Each user has a unique name within the AWS account, and a set of security credentials that aren't shared with other users. These credentials are separate from the security credentials for the AWS account. Each user is associated with one and only one AWS account.
USER_PERSONALIZATION recipes	Amazon Personalize : Recipes that are used to build a recommendation system that predicts the items that a user interacts with based on data provided in Interactions, Items, and Users datasets. See Also recipe , user-personalization recipe , popularity-count recipe , HRNN .
user-personalization recipe	Amazon Personalize : An HRNN-based USER_PERSONALIZATION recipe that predicts the items that a user interacts with. The user-personalization recipe can use item exploration and impressions data to generate recommendations for new items. See Also HRNN , recipe , USER_PERSONALIZATION recipes , item exploration , impressions data , recommendations .
Users dataset	Amazon Personalize : A container for metadata about your users, such as age, gender, or loyalty membership. See Also dataset .

V

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

validation	See template validation .
value	Instances of attributes for an item, such as cells in a spreadsheet. An attribute might have multiple values. Tagging resources: A specific tag label that acts as a descriptor within a tag category (key). For example, you might have EC2 instance with the tag key of <i>Owner</i> and the tag value of <i>Jan</i> . You can tag an AWS resource with up to 10 key–value pairs. Not all AWS resources can be tagged.
Variable Envelope Return Path	See VERP .

verification	The process of confirming that you own an email address or a domain so that you can send email from or to it.
VERP	Variable Envelope Return Path. A way that email-sending applications can match bounced email with the undeliverable address that caused the bounce by using a different return path for each recipient. VERP is typically used for mailing lists. With VERP, the recipient's email address is embedded in the address of the return path, which is where bounced email is returned. This makes it possible to automate the processing of bounced email without having to open the bounce messages, which might vary in content.
versioning	Every object in Amazon S3 has a key and a version ID. Objects with the same key, but different version IDs can be stored in the same bucket . Versioning is enabled at the bucket layer using PUT Bucket versioning.
VGW	See virtual private gateway (VGW) .
virtual private gateway (VGW)	The Amazon side of a VPN connection that maintains connectivity. The internal interfaces of the virtual private gateway connect to your Amazon VPC through the VPN attachment. The external interfaces connect to the VPN connection, which leads to the customer gateway .
virtualization	Allows multiple guest virtual machines (VM) to run on a host operating system. Guest VMs can run on one or more levels above the host hardware, depending on the type of virtualization. See Also PV virtualization , HVM virtualization .
visibility timeout	The period of time that a message is invisible to the rest of your application after an application component gets it from the queue. During the visibility timeout, the component that received the message usually processes it, and then deletes it from the queue. This prevents multiple components from processing the same message.
VM Import/Export	VM Import/Export is a service for importing virtual machine (VM) images from your existing virtualization environment to Amazon EC2 and then exporting them back. See Also https://aws.amazon.com/ec2/vm-import .

volume	A fixed amount of storage on an instance . You can share volume data between more than one container and persist the data on the container instance when the containers are no longer running.
Amazon VPC	Amazon Virtual Private Cloud is a web service for provisioning a logically isolated section of the AWS Cloud virtual network that you define. You control your virtual networking environment by selecting your own IP address range, creating subnets and configuring route tables and network gateways. See Also https://aws.amazon.com/vpc .
VPC endpoint	A feature that you can use to create a private connection between your Amazon VPC and another AWS service without requiring access over the internet, through a NAT instance, a VPN connection , or Direct Connect .
VPNG	See virtual private gateway (VGW) .
AWS VPN	AWS Virtual Private Network provides functionality that establishes encrypted connections between your network or device, and AWS. AWS VPN is comprised of two services: AWS Client VPN and AWS Site-to-Site VPN . See Also https://aws.amazon.com/vpn .
AWS VPN CloudHub	AWS VPN CloudHub is a feature that enables secure communication between branch offices using a simple hub-and-spoke model, with or without a VPN.
VPN connection	Amazon Web Services (AWS) : The IPsec connection that's between a Amazon VPC and some other network, such as a corporate data center, home network, or colocation facility.

W

[Numbers and symbols](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#), [Y](#), [Z](#)

AWS WAF	AWS WAF is a web application firewall service that controls access to content by allowing or blocking web requests based on criteria that you specify. For example, you can filter access based on the header values or the IP addresses that the requests originate from. AWS WAF helps
---------	--

	<p>protect web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources.</p> <p>See Also https://aws.amazon.com/waf/.</p>
Amazon WAM	<p>Amazon WorkSpaces Application Manager (Amazon WAM) is a web service for deploying and managing applications for WorkSpaces. Amazon WAM accelerates software deployment, upgrades, patching, and retirement by packaging Windows desktop applications into virtualized application containers.</p> <p>See Also https://aws.amazon.com/workspaces/applicationmanager.</p>
warm standby	<p>An active-passive disaster recovery strategy in which a workload is scaled down in the passive standby Region, but is otherwise fully functional. This is not an Amazon EC2 Auto Scaling term, but an industry-standard resilience term.</p> <p>See Also , , .</p>
AWS Wavelength	<p>AWS Wavelength is a service by AWS that embeds AWS compute and storage services within 5G networks to provide mobile edge computing infrastructure. Use AWS Wavelength to develop, deploy, and scale ultra-low-latency applications to mobile devices and end users.</p> <p>See Also https://aws.amazon.com/wavelength.</p>
web access control list (web ACL)	<p>AWS WAF: A set of rules that defines the conditions that AWS WAF searches for in web requests to an AWS resource, such as a Amazon CloudFront distribution. A web access control list (web ACL) specifies if to allow, block, or count the requests.</p>
Web Services Description Language	<p>See WSDL.</p>
WorkDocs	<p>Amazon WorkDocs is a managed, secure enterprise document storage and sharing service with administrative controls and feedback capabilities.</p> <p>See Also https://aws.amazon.com/workdocs/.</p>
Amazon WorkLink	<p>Amazon WorkLink is a cloud-based service that provides secure access to internal websites and web apps from mobile devices.</p>

See Also <https://aws.amazon.com/worklink/>.

WorkMail

Amazon WorkMail is a managed, secure business email and calendar service with support for existing desktop and mobile email clients. See Also <https://aws.amazon.com/workmail/>.

WorkSpaces

Amazon WorkSpaces is a managed, secure desktop computing service for provisioning cloud-based desktops and providing users access to documents, applications, and [resources](#) from supported devices. See Also <https://aws.amazon.com/workspaces/>.

WSDL

Web Services Description Language. A language that's used to describe the actions that a web service can perform, along with the syntax of action requests and responses. See Also [REST](#), [SOAP](#).

X, Y, Z

X.509 certificate

A digital document that uses the X.509 public key infrastructure (PKI) standard to verify that a public key belongs to the entity that's described in the [certificate](#).

X-Ray

AWS X-Ray is a web service that collects data about requests that your application serves. X-Ray provides tools that you can use to view, filter, and gain insights into that data to identify issues and opportunities for optimization. See Also <https://aws.amazon.com/xray/>.

yobibyte (YiB)

A contraction of yotta binary byte. A yobibyte (YiB) is 2^{80} or 1,208,925,819,614,629,174,706,176 bytes. A yottabyte (YB) is 10^{24} or 1,000,000,000,000,000,000,000,000 bytes.

zebibyte (ZiB)

A contraction of zetta binary byte. A zebibyte (ZiB) is 2^{70} or 1,180,591,620,717,411,303,424 bytes. A zettabyte (ZB) is 10^{21} or 1,000,000,000,000,000,000,000 bytes. 1,024 ZiB is a [yobibyte \(YiB\)](#).

zone awareness

[OpenSearch Service](#): A configuration that distributes nodes in a cluster across two [Availability Zones](#) in the same Region. Zone awareness helps to prevent data loss and minimizes downtime if a node and data center

fails. If you enable zone awareness, you must have an even number of data instances in the instance count, and you also must use the Amazon OpenSearch Service Configuration API to replicate your data for your OpenSearch cluster.