



System and Organization Controls 3 (SOC 3) Report
Report on the Amazon Web Services System Relevant to
Security, Availability, and Confidentiality
For the Period April 1, 2019 - September 30, 2019



Report of Independent Accountants

To the Management of Amazon Web Services, Inc.

Scope:

We have examined management's assertion, contained within the accompanying "Report on the Amazon Web Services System Relevant to Security, Availability, and Confidentiality" (Assertion), that Amazon Web Services, Inc.'s (AWS) controls over the Amazon Web Services System (System) were effective throughout the period April 1, 2019 to September 30, 2019, to provide reasonable assurance that its principal service commitments and system requirements were achieved based on the criteria relevant to security, availability, and confidentiality (applicable trust services criteria) set forth in the AICPA's TSP section 100, *2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy*.

Management's Responsibilities

AWS' management is responsible for its assertion, selecting the trust services categories and associated criteria on which its assertion is based, and having a reasonable basis for its assertion. It is also responsible for:

- Identifying the Amazon Web Services System and describing the boundaries of the System
- Identifying our principal service commitments and system requirements and the risks that would threaten the achievement of its principal service commitments and service requirements that are the objectives of our system
- identifying, designing, implementing, operating, and monitoring effective controls over the Amazon Web Services System to mitigate risks that threaten the achievement of the principal service commitments and system requirement

Our Responsibilities

Our responsibility is to express an opinion on the Assertion, based on our examination. Our examination was conducted in accordance with attestation standards established by the American Institute of Certified Public Accountants. Those standards require that we plan and perform our examination to obtain reasonable assurance about whether management's assertion is fairly stated, in all material respects. An examination involves performing procedures to obtain evidence about management's assertion, which includes (1) obtaining an understanding of AWS' relevant security, availability, and confidentiality policies, processes and controls, (2) testing and evaluating the operating effectiveness of the controls and (3) performing such other procedures as we considered necessary in the circumstances. The nature, timing, and extent of the procedures selected depend on our judgment, including an assessment of the risk of material misstatement, whether due to fraud or error. We believe that the evidence obtained during our examination is sufficient to provide a reasonable basis for our opinion.

Our examination was not conducted for the purpose of evaluating AWS' cybersecurity risk management program. Accordingly, we do not express an opinion or any other form of assurance on its cybersecurity risk management program.

Inherent limitations:

Because of their nature and inherent limitations, controls may not prevent, or detect and correct all misstatements that may be considered relevant. Furthermore, the projection of any evaluations of effectiveness to future periods, or conclusions about the suitability of the design of the controls to achieve AWS' principal service commitments and system requirements, is subject to the risk that controls may become inadequate because of changes in conditions, that the degree of compliance with such controls may deteriorate, or that changes made to the system or controls, or the failure to make needed changes to the system or controls, may alter the validity of such evaluations. Examples of inherent limitations of internal controls related to security include (a) vulnerabilities in information technology components as a result of design by their manufacturer or developer; (b) breakdown of internal control at a vendor or business partner; and (c) persistent attackers with the resources to use advanced technical means and sophisticated social engineering techniques specifically targeting the entity.

Opinion:

In our opinion, AWS' controls over the system were effective throughout the period April 1, 2019 to September 30, 2019, to provide reasonable assurance that its principal service commitments and system requirements were achieved based on the applicable trust services criteria.

Ernst & Young LLP

October 28, 2019



Amazon Web Services
410 Terry Avenue North
Seattle, WA 98109-5210

**Management's Report of its Assertions on the Effectiveness of Its Controls
Over the Amazon Web Services System
Based on the Trust Services Criteria for Security, Availability, and Confidentiality**

We, as management of, Amazon Web Services, Inc. are responsible for:

- Identifying the AWS Web Services System (System) and describing the boundaries of the System, which are presented in Attachment A
- Identifying our principal service commitments and system requirements
- Identifying the risks that would threaten the achievement of its principal service commitments and service requirements that are the objectives of our system, which are presented in Attachment B
- identifying, designing, implementing, operating, and monitoring effective controls over the System to mitigate risks that threaten the achievement of the principal service commitments and system requirement
- Selecting the trust services categories that are the basis of our assertion

We assert that the controls over the system were effective throughout the period April 1, 2019 to September 30, 2019 to provide reasonable assurance that the principle service commitments and system requirements were achieved based on the criteria relevant to security, availability and confidentiality set forth in the AICPA's TSP section 100, *2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality and Privacy*.

Very truly yours,

Amazon Web Services Management



Attachment A – Amazon Web Services System

AWS Background

Since 2006, Amazon Web Services (AWS) has provided flexible, scalable and secure IT infrastructure to businesses of all sizes around the world. With AWS, customers can deploy solutions on a cloud computing environment that provides compute power, storage, and other application services over the Internet as their business needs demand. AWS affords businesses the flexibility to employ the operating systems, application programs, and databases of their choice.

The scope covered in this report consists of the following services (the service name is followed by the services' namespace¹ in parenthesis):

- AWS Amplify Console (`amplify`)
- API Gateway (`apigateway`)
- Amazon AppStream 2.0 (`appstream`)
- AWS AppSync (`appsync`)
- Amazon Athena (`athena`)
- AWS Auto Scaling (`autoscaling`)
- AWS Backup (`backup`)
- AWS Batch (`batch`)
- AWS Certificate Manager (`acm`)
- Amazon Cloud Directory (`clouddirectory`)
- AWS CloudFormation (`cloudformation`)
- Amazon CloudFront (`cloudfront`)
- AWS CloudHSM (`cloudhsm`)
- AWS CloudTrail (`cloudtrail`)
- Amazon CloudWatch (`cloudwatch`, `events`, `logs`)
- CloudWatch SDK Metrics for Enterprise Support (`sdkmetrics`)
- AWS CodeBuild (`codebuild`)
- AWS CodeCommit (`codecommit`)
- AWS CodeDeploy (`codedeploy`)
- AWS CodePipeline (`codepipeline`)
- Amazon Cognito (`cognito-idp`, `cognito-identity`, `cognito-sync`)
- Amazon Comprehend (`comprehend`)
- Amazon Comprehend Medical (`comprehendmedical`)
- AWS Config (`config`)
- AWS IoT Events (`iotevents`)
- AWS IoT Greengrass (`greengrass`)
- AWS Key Management Service (`kms`)
- Amazon Kinesis Data Analytics (`kinesisanalytics`)
- Amazon Kinesis Data Firehose (`firehose`)
- Amazon Kinesis Data Streams (`kinesis`)
- Amazon Kinesis Video Streams (`kinesisvideo`)
- AWS Lambda (`lambda`)
- Amazon Lex (`runtime.lex`, `models.lex`)
- AWS License Manager (`license-manager`)
- Amazon Macie (`macie`)
- AWS Managed Services
- Amazon Managed Streaming for Kafka (Amazon MSK) (`kafka`)
- Amazon MQ (`mq`)
- Amazon Neptune (`neptune-db`)
- AWS OpsWorks (`opsworks`)
- AWS OpsWorks for Chef Automate or AWS OpsWorks for Puppet Enterprise (`opsworks-cm`)
- AWS Organizations (`organizations`)
- Amazon Personalize (`personalize`)
- Amazon Pinpoint (`mobiletargeting`)
- Amazon Polly (`polly`)

¹ When customers create IAM policies or work with Amazon Resource Names (ARNs), customers identify an AWS service using a *namespace*. For example, the namespace for Amazon S3 is `s3`, and the namespace for Amazon EC2 is `ec2`. Customers use namespaces when identifying actions and resources across AWS.



- Amazon Connect (`connect`)
- AWS Control Tower (`controltower`)
- AWS Database Migration Service (`dms`)
- AWS DataSync (`datasync`)
- AWS Direct Connect (`directconnect`)
- AWS Directory Service (`ds`) – [Excludes Simple Active Directory]
- Amazon DocumentDB (with MongoDB compatibility)
- Amazon DynamoDB (`dynamodb`)
- AWS Elastic Beanstalk (`elasticbeanstalk`)
- Amazon Elastic Block Store (`ec2`)
- Amazon Elastic Compute Cloud (`ec2`)
- Amazon Elastic Container Registry (`ecr`)
- Amazon Elastic Container Service (`ecs`) – [both Fargate and EC2 launch types]
- Amazon Elastic Container Service for Kubernetes (`eks`)
- Amazon Elastic File System (`elasticfilesystem`)
- Amazon Elasticsearch Service (`es`)
- Elastic Load Balancing (`elasticloadbalancing`)
- Amazon ElastiCache (`elasticache`)
- AWS Elemental MediaConnect (`mediaconnect`)
- Amazon EMR (`elasticmapreduce`)
- AWS Firewall Manager (`fms`)
- Amazon Forecast (`amazonforecast`)
- Amazon FreeRTOS (`freertos`)
- Amazon FSx (`fsx`)
- Amazon Glacier (`glacier`)
- AWS Global Accelerator (`globalaccelerator`)
- AWS Glue (`glue`)
- AWS GuardDuty (`guardduty`)
- AWS Identity and Access Management (`iam`)
- VM Import/Export
- Amazon Inspector (`inspector`)
- AWS IoT Core (`iot`)
- AWS IoT Device Management (`iot`)
- Amazon QuickSight (`quicksight`)
- Amazon Redshift (`redshift`)
- Amazon Rekognition (`rekognition`)
- Amazon Relational Database Service (`rds`)
- AWS Resource Groups (`resource-groups`)
- AWS RoboMaker (`robomaker`)
- Amazon Route 53 (`route53`)
- Amazon SageMaker (`sagemaker`)
- AWS Secrets Manager (`secretsmanager`)
- AWS Security Hub (`securityhub`)
- AWS Server Migration Service (`sms`)
- AWS Serverless Application Repository (`serverlessrepo`)
- AWS Service Catalog (`servicecatalog`)
- AWS Shield (`shield`, `DDoSProtection`)
- Amazon Simple Email Service (`ses`)
- Amazon Simple Notification Service (`sns`)
- Amazon Simple Queue Service (`sqs`)
- Amazon Simple Storage Service (`s3`)
- Amazon Simple Workflow Service (`swf`)
- Amazon SimpleDB (`sdb`)
- AWS Snowball (`snowball`)
- AWS Snowball Edge
- AWS Snowmobile
- AWS Step Functions (`states`)
- AWS Storage Gateway (`storagegateway`)
- AWS Systems Manager (`ssm`)
- Amazon Transcribe (`transcribe`)
- AWS Transfer for SFTP (`transfer`)
- Amazon Translate (`translate`)
- Amazon Virtual Private Cloud (Amazon VPC) (`ec2`)
- AWS WAF (`waf`)
- Amazon WorkDocs (`workdocs`)
- Amazon WorkLink (`worklink`)
- Amazon WorkMail (`workmail`)
- Amazon WorkSpaces (`workspaces`)
- AWS X-Ray (`xray`)



The scope of locations covered in this report includes the supporting data centers located in the US East (Northern Virginia), US East (Ohio), US West (Oregon), US West (Northern California), GovCloud (US West), GovCloud (US East), Canada (Montreal), Europe (Ireland), Europe (Frankfurt), Europe (London), Europe (Paris), Europe (Stockholm), Middle East (Bahrain), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Asia Pacific (Osaka)², Asia Pacific (Seoul), Asia Pacific (Mumbai), Asia Pacific (Hong Kong), and South America (São Paulo) Regions, as well as the AWS Edge locations in:

- Canberra, Australia
- Melbourne, Australia
- Perth, Australia
- Sydney, Australia
- Vienna, Austria
- Rio de Janeiro, Brazil
- São Paulo, Brazil
- Montréal, Canada
- Toronto, Canada
- Vancouver, Canada
- Prague, Czech Republic
- Hong Kong, China
- Copenhagen, Denmark
- London, England
- Manchester, England
- Helsinki, Finland
- Marseille, France
- Paris, France
- Berlin, Germany
- Frankfurt, Germany
- Munich, Germany
- Bengaluru, India
- Chennai, India
- Hyderabad, India
- Mumbai, India
- New Delhi, India
- Dublin, Ireland
- Tel Aviv, Israel
- Milan, Italy
- Palermo, Italy
- Osaka, Japan
- Tokyo, Japan
- Seoul, Korea
- Kuala Lumpur, Malaysia
- Amsterdam, Netherlands
- Oslo, Norway
- Manila, Philippines
- Warsaw, Poland
- Lisbon, Portugal
- Singapore
- Cape Town, South Africa
- Johannesburg, South Africa
- Madrid, Spain
- Stockholm, Sweden
- Zurich, Switzerland
- Taipei, Taiwan
- Dubai, United Arab Emirates
- Fujairah, United Arab Emirates
- Arizona, United States
- California, United States
- Colorado, United States
- Florida, United States
- Georgia, United States
- Illinois, United States
- Indiana, United States
- Massachusetts, United States
- Minnesota, United States
- Nevada, United States
- New Jersey, United States
- New York, United States
- Ohio, United States
- Oregon, United States
- Pennsylvania, United States
- Texas, United States
- Utah, United States
- Virginia, United States
- Washington, United States

Infrastructure

AWS operates the cloud infrastructure that customers may use to provision computing resources such as processing and storage. The AWS infrastructure includes the facilities, network, and hardware as well as some operational software (e.g., host operating system, virtualization software, etc.) that support the provisioning and use of these resources. The AWS infrastructure is designed and managed in accordance with security compliance standards and AWS best practices.

² The Asia Pacific (Osaka) Region is a Local Region, which comprises an isolated, fault-tolerant infrastructure design consisting of three virtual Availability Zones located in the same data center and is intended to be used in conjunction with the Asia Pacific (Tokyo) Region. This region requires customers request access through a sales representative.



Components of the System

AWS offers a series of Analytics; Application Integration; Business Productivity; Compute; Customer Engagement; Database; Desktop & App Streaming; Developer Tools; Internet of Things; Management Tools; Media Services; Migration; Mobile Services; Network & Content Delivery; Security, Identity, and Compliance; and Storage services. A description of the AWS services included within the scope of this report is listed below:

AWS Amplify Console ([amplify](#))

AWS Amplify makes it easy to create, configure, and implement scalable mobile and web apps powered by AWS. Amplify seamlessly provisions and manages the mobile backend and provides a simple framework to easily integrate the backend with the iOS, Android, Web, and React Native frontends. Amplify also automates the application release process of both the frontend and backend allowing the customers to deliver features faster.

API Gateway ([apigateway](#))

API Gateway is a fully managed service that makes it easy for developers to publish, maintain, monitor, and secure APIs at any scale. With Amazon API Gateway, customers can create a custom API to code running in AWS Lambda, and then call the Lambda code from customers' API.

Amazon AppStream 2.0 ([appstream](#))

Amazon AppStream 2.0 is a fully managed application streaming service that provides customers instant access to their desktop applications from anywhere. Amazon AppStream 2.0 simplifies application management, improves security, and reduces costs by moving a customer's applications from their users' physical devices to the AWS Cloud. The Amazon AppStream 2.0 streaming protocol provides customers a responsive, fluid performance that is almost indistinguishable from a natively installed application. With Amazon AppStream 2.0, customers can realize the agility to support a broad range of compute and storage requirements for their applications.

AWS AppSync ([appsync](#))

AWS AppSync automatically updates the data in web and mobile applications in real time, and updates data for offline users as soon as they reconnect. AWS AppSync makes it easy to build collaborative mobile and web applications that deliver responsive, collaborative user experiences.

Amazon Athena ([athena](#))

Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure for customers to manage. Athena is highly available; and executes queries using compute resources across multiple facilities and multiple devices in each facility. Amazon Athena uses Amazon S3 as its underlying data store, making customers' data highly available and durable.



AWS Auto Scaling ([autoscaling](#))

Auto Scaling launches/terminates instances on a customer's behalf according to conditions customers define, such as schedule, changing metrics like average CPU utilization, or health of the instance as determined by EC2 or ELB health checks. It allows customers to have balanced compute across multiple availability zones and scale their fleet based on usage.

AWS Backup ([backup](#))

AWS Backup is a fully managed backup service that makes it easy to centralize and automate the back up of data across AWS services in the cloud as well as on premises using the AWS Storage Gateway.

AWS Batch ([batch](#))

AWS Batch enables developers, scientists, and engineers to run batch computing jobs on AWS. AWS Batch dynamically provisions the optimal quantity and type of compute resources (e.g., CPU or memory optimized instances) based on the volume and specific resource requirements of the batch jobs submitted. AWS Batch plans, schedules, and executes customers' batch computing workloads across the full range of AWS compute services and features, such as Amazon EC2 and Spot Instances.

AWS Certificate Manager ([acm](#))

AWS Certificate Manager is a service that lets the customer provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and their internal connected resources.

Amazon Cloud Directory ([clouddirectory](#))

Amazon Cloud Directory enables customers to build flexible cloud-native directories for organizing hierarchies of data along multiple dimensions. Customers also can create directories for a variety of use cases, such as organizational charts, course catalogs, and device registries.

AWS CloudFormation ([cloudformation](#))

AWS CloudFormation enables customers to create and manage a collection of related AWS resources by providing templates to use in the provisioning and updating of AWS services.

Amazon CloudFront ([cloudfront](#))

Amazon CloudFront is a web service that speeds up distribution of customers' static and dynamic web content. CloudFront delivers customers' content through a worldwide network of Edge locations and Amazon Cloud Extension (ACE) locations. ACE is a CloudFront infrastructure (single-rack version) deployed to a non Amazon-controlled facility, namely an internet service provider (ISP) or partner network.

AWS CloudHSM ([cloudhsm](#))

AWS CloudHSM is a service that allows customers to use dedicated hardware security module (HSM) appliances within the AWS cloud. AWS CloudHSM allows customers to store and use encryption keys within HSM appliances in AWS data centers.



AWS CloudTrail ([cloudtrail](#))

AWS CloudTrail is a web service that records AWS activity for customers and delivers log files to a specified Amazon S3 bucket. AWS CloudTrail provides a history of AWS API calls for customer accounts.

Amazon CloudWatch ([cloudwatch](#), [events](#), [logs](#))

Amazon CloudWatch is a monitoring and management service built for developers, system operators, site reliability engineers (SRE), and IT managers. CloudWatch provides the customers with data and actionable insights to monitor their applications, understand and respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health.

CloudWatch SDK Metrics for Enterprise Support ([sdkmetrics](#))

Amazon CloudWatch SDK Metrics for Enterprise Support is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers (SREs), and IT managers. CloudWatch provides customers with data and actionable insights to monitor their applications, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health. CloudWatch collects monitoring and operational data in the form of logs, metrics, and events, providing customers with a unified view of AWS resources, applications, and services that run on AWS and on-premises servers. Customer can use CloudWatch to detect anomalous behavior in their environments, set alarms, visualize logs and metrics side by side, take automated actions, troubleshoot issues, and discover insights to keep your applications running smoothly.

AWS CodeBuild ([codebuild](#))

AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages that are ready to deploy. CodeBuild scales continuously and processes multiple builds concurrently, so that customers' builds are not left waiting in a queue. Customers can use prepackaged build environments or can create custom build environments that use their own build tools. AWS CodeBuild eliminates the need to set up, patch, update, and manage customers' build servers and software.

AWS CodeCommit ([codecommit](#))

AWS CodeCommit is a fully-managed source control service that hosts secure Git-based repositories. It allows teams to collaborate on code in a secure and highly scalable ecosystem. CodeCommit eliminates the need for customers to operate their own source control system or worry about scaling their infrastructure.

AWS CodeDeploy ([codedeploy](#))

AWS CodeDeploy is a fully managed deployment service that automates software deployments to a variety of compute services such as Amazon EC2, AWS Fargate, AWS Lambda, and the customer's on-premises servers.



AWS CodePipeline (`codepipeline`)

AWS CodePipeline is a fully managed continuous delivery service that helps customers automate release pipelines for fast and reliable application and infrastructure updates. CodePipeline automates the build, test, and deploy phases of customers' release process every time there is a code change, based on the release model defined by the customer. This enables customers to rapidly and reliably deliver features and updates. Customers can easily integrate AWS CodePipeline with third-party services such as GitHub or with their own custom plugin.

Amazon Cognito (`cognito-idp`, `cognito-identity`, `cognito-sync`)

Amazon Cognito lets customers add user sign-up, sign-in, and manage permissions for mobile and web applications. Customers can create their own user directory within Amazon Cognito. Customers can also choose to authenticate users through social identity providers such as Facebook, Twitter, or Amazon; with SAML identity solutions; or by using customers' own identity system.

Amazon Comprehend (`comprehend`)

Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text. Amazon Comprehend uses machine learning to help the customers uncover the insights and relationships in their unstructured data.

Amazon Comprehend Medical (`comprehendmedical`)

Amazon Comprehend Medical is a natural language processing service that facilitates the use of machine learning to extract relevant medical information from unstructured text. Using Amazon Comprehend Medical, customers can quickly and accurately gather information, such as medical condition, medication, dosage, strength, and frequency from a variety of sources like doctors' notes, clinical trial reports, and patient health records. Amazon Comprehend Medical uses advanced machine learning models to accurately and quickly identify medical information, such as medical conditions and medications, and determines their relationship to each other, for instance, medicine dosage and strength.

AWS Config (`config`)

AWS Config enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records AWS resource configurations and allows customers to automate the evaluation of recorded configurations against desired configurations.

Amazon Connect (`connect`)

Amazon Connect is a self-service, cloud-based contact center service that enables dynamic, personal, and natural customer engagement at any scale. The self-service graphical interface allows the customers to design contact flows, manage agents, and track performance metrics.



AWS Control Tower ([controlltower](#))

AWS Control Tower provides a way to set up and govern a new, secure, multi-account AWS environment based on best practices established through AWS' experience working with thousands of enterprises as they move to the cloud. With AWS Control Tower, builders can provision new AWS accounts that conforms to customers' policies. If you are building a new AWS environment, starting out on your journey to AWS, starting a new cloud initiative, or are completely new to AWS, Control Tower will help you get started quickly with governance and best practices built-in.

AWS Database Migration Service ([dms](#))

AWS Database Migration Service enables customers to migrate databases between similar and different database programs in the cloud and off-cloud. The service supports homogenous migrations within one database platform, as well as heterogeneous migrations between different database platforms.

AWS DataSync ([datasync](#))

AWS DataSync is a data transfer service that makes it easy for customers to automate moving data between on-premises storage and Amazon S3 or Amazon Elastic File System (Amazon EFS). DataSync automatically handles many of the tasks related to data transfers that can slow down migrations or burden customers' IT operations, including running customers' own instances, handling encryption, managing scripts, network optimization, and data integrity validation.

AWS Direct Connect ([directconnect](#))

AWS Direct Connect enables customers to establish a dedicated network connection between their network and one of the AWS Direct Connect locations. Using AWS Direct Connect, customers can establish private connectivity between AWS and their datacenter, office, or colocation environment.

AWS Directory Service ([ds](#)) – [Excludes Simple Active Directory]

AWS Directory Service for Microsoft Active Directory, also known as AWS Microsoft AD, enables customers' directory-aware workloads and AWS resources to use managed Active Directory in the AWS Cloud. AWS Microsoft AD stores directory content in encrypted Amazon Elastic Block Store volumes using encryption keys that AWS manages.

Amazon DocumentDB (with MongoDB compatibility)

Amazon DocumentDB (with MongoDB compatibility) is a fast, scalable, highly available, and fully managed document database service that supports MongoDB workloads. Amazon DocumentDB is designed from the ground-up to give customers the performance, scalability, and availability customers need when operating mission-critical MongoDB workloads at scale.

Amazon DynamoDB ([dynamodb](#))

Amazon DynamoDB is a managed NoSQL database service. Amazon DynamoDB enables customers to offload to AWS the administrative burdens of operating and scaling distributed databases such as hardware provisioning, setup and configuration, replication, software patching, and cluster scaling.



AWS Elastic Beanstalk ([elasticbeanstalk](#))

AWS Elastic Beanstalk is an application container launch program for customers to launch and scale their applications on top of AWS. Customers can use AWS Elastic Beanstalk to create new environments using Elastic Beanstalk curated programs and their applications, deploy application versions, update application configurations, rebuild environments, update AWS configurations, monitor environment health and availability, and build on top of the scalable infrastructure.

Amazon Elastic Block Store ([ec2](#))

Amazon Elastic Block Store provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. Storage volumes behave like raw, unformatted block devices, with user supplied device names and a block device interface. Customers can create a file system on top of Amazon EBS volumes, or use them in any other way one would use a block device (like a hard drive).

Amazon Elastic Compute Cloud ([ec2](#))

Amazon Elastic Compute Cloud is Amazon's Infrastructure as a Service (IaaS) offering, which provides scalable computing capacity using server instances in AWS' data centers. Amazon EC2 is designed to make web-scale computing easier by enabling customers to obtain and configure capacity with minimal friction. Customers create and launch instances, which are virtual machines that are available in a wide variety of hardware and software configurations.

Amazon Elastic Container Registry ([ecr](#))

Amazon Elastic Container Registry is a fully-managed Docker container registry that makes it easy for developers to store, manage, and deploy Docker container images. Amazon Elastic Container Registry is integrated with Amazon Elastic Container Service.

Amazon Elastic Container Service ([ecs](#)) – [both Fargate and EC2 launch types]

Amazon Elastic Container Service is a highly scalable, high performance container management service that supports Docker containers and allows customers to easily run applications on a managed cluster of Amazon EC2 instances. Amazon Elastic Container Service eliminates the need for customers to install, operate, and scale customers' own cluster management infrastructure.

Amazon Elastic Container Service for Kubernetes ([eks](#))

Amazon Elastic Container Service for Kubernetes (Amazon EKS) makes it easy to deploy, manage, and scale containerized applications using Kubernetes on AWS. Amazon EKS runs the Kubernetes management infrastructure for the customer across multiple AWS availability zones to eliminate a single point of failure.

Amazon Elastic File System ([elasticfilesystem](#))

Amazon Elastic File System provides file storage for Amazon EC2 instances that grows and shrinks elastically as data is added and deleted by users. Amazon EFS spreads data across multiple Availability Zones; in the event that an Availability Zone is not reachable, the structure allows customers to still access their full set of data.



Amazon Elasticsearch Service (es)

Amazon Elasticsearch Service is a fully managed service that makes it easy for the customer to deploy, secure, and operate Elasticsearch at scale with zero down time. Amazon Elasticsearch Service lets the customers pay only for what they use – there are no upfront costs or usage requirements.

Elastic Load Balancing (elasticloadbalancing)

Elastic Load Balancing provides customers with a load balancer that automatically distributes incoming application traffic across multiple Amazon EC2 instances in the cloud. It allows customers to achieve greater levels of fault tolerance for their applications, seamlessly providing the required amount of load balancing capacity needed to distribute application traffic.

Amazon ElastiCache (elasticache)

Amazon ElastiCache automates management tasks for in-memory cache environments, such as patch management, failure detection, and recovery. It works in conjunction with other AWS services to provide a managed in-memory cache.

AWS Elemental MediaConnect (mediaconnect)

AWS Elemental MediaConnect is a high-quality transport service for live video. MediaConnect enables customers to build mission-critical live video workflows in a fraction of the time and cost of satellite or fiber services.

Amazon EMR (elasticmapreduce)

Amazon EMR is a web service that provides managed Hadoop clusters on Amazon EC2 instances running a Linux operating system. Amazon EMR actively manages clusters for customers, replacing failed nodes and adjusting capacity as requested.

AWS Firewall Manager (fms)

AWS Firewall Manager is a security management service that makes it easier to centrally configure and manage AWS WAF rules across customer accounts and applications. Using Firewall Manager, customers can roll out AWS WAF rules for their Application Load Balancers and Amazon CloudFront distributions across accounts in AWS Organizations.

Amazon Forecast (amazonforecast)

Amazon Forecast uses machine learning to combine time series data with additional variables to build forecasts. It is a fully managed service, so there are no servers to provision, and no machine learning models to build, train, or deploy. With Amazon Forecast, customers can import time series data and associated data into Amazon Forecast from their Amazon S3 database. From there, Amazon Forecast automatically loads the data, inspects it, and identifies the key attributes needed for forecasting. Amazon Forecast then trains and optimizes a customer's custom model, and hosts them in a highly available environment where it can be used to generate business forecasts.



Amazon Forecast is protected by encryption. Any content processed by Amazon Forecast is encrypted with customer keys through Amazon Key Management Service, and encrypted at rest in the AWS Region where a customer is using the service. Administrators can also control access to Amazon Forecast through an AWS Identity and Access Management (IAM) permissions policy – ensuring that sensitive information is kept secure and confidential.

Amazon FreeRTOS ([freertos](#))

Amazon FreeRTOS is an operating system for microcontrollers that makes small, low-power edge devices easy to program, deploy, secure, connect, and manage.

Amazon FSx ([fsx](#))

Amazon FSx provides fully managed third-party file systems. Amazon FSx provides the customers with the native compatibility of third-party file systems with feature sets for workloads such as Windows-based storage, high-performance computing (HPC), machine learning, and electronic design automation (EDA).

Amazon Glacier ([glacier](#))

Amazon Glacier is an archival storage solution for data that is infrequently accessed for which retrieval times of several hours are suitable. Amazon Glacier enables customers to set access policies on their vaults for users within their AWS Account.

AWS Global Accelerator ([globalaccelerator](#))

AWS Global Accelerator is a networking service that improves the availability and performance of the applications that customers offer to their global users. AWS Global Accelerator is easy to set up, configure and manage.

AWS Glue ([glue](#))

AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy for customers to prepare and load their data for analytics. The customers can create and run an ETL job with a few clicks in the AWS Management Console.

AWS GuardDuty ([guardduty](#))

AWS GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect the customers' AWS accounts and workloads. With GuardDuty, the customers now have an intelligent and cost-effective option for continuous threat detection in the AWS Cloud.

AWS Identity and Access Management ([iam](#))

AWS Identity and Access Management is a web service that helps customers securely control access to AWS resources for their users. Customers use IAM to control who can use their AWS resources (authentication) and what resources they can use and in what ways (authorization).



VM Import/Export

AWS Import/Export is a service that enables customers to import virtual machine images from their existing environment to Amazon EC2 instances and export them back to their off-cloud environment.

Amazon Inspector (`inspector`)

Amazon Inspector is an automated security assessment service for customers seeking to 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 list of security findings prioritized by level of severity.

AWS IoT Core (`iot`)

AWS IoT Core is a managed cloud service that lets connected devices easily and securely interact with cloud applications and other devices. AWS IoT Core provides secure communication and data processing across different kinds of connected devices and locations so that customers can easily build IoT applications such as industrial solutions and connected home solutions.

AWS IoT Device Management (`iot`)

AWS IoT Device Management provides customers with ability to securely onboard, organize, and remotely manage IoT devices at scale. With AWS IoT Device Management, customer can register their connected devices individually or in bulk, and manage permissions so that devices remain secure.

AWS IoT Events (`iotevents`)

AWS IoT Events is a fully managed service that detects events across thousands of IoT sensors sending different telemetry data, such as temperature from a freezer, humidity from respiratory equipment, and belt speed on a motor. Customers can select the relevant data sources to ingest, define the logic for each event using simple 'if-then-else' statements, and select the alert or custom action to trigger when an event occurs. IoT Events continuously monitors data from multiple IoT sensors and applications, and it integrates with other services, such as AWS IoT Core and AWS IoT Analytics, to enable early detection and unique insights into events. IoT Events automatically triggers alerts and actions in response to events based on the logic defined to resolve issues quickly, reduce maintenance costs, and increase operational efficiency.

AWS IoT Greengrass (`greengrass`)

AWS IoT Greengrass seamlessly extends AWS to edge devices so they can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage.

AWS Key Management Service (`kms`)

AWS Key Management Service allows users to create and manage cryptographic keys. One class of keys, Customer Master Keys (CMKs), are designed to never be exposed in plaintext outside the service. CMKs can be used to encrypt data directly submitted to the service. CMKs can also be used to protect other types of keys, Data Encryption Keys (DEKs), which are created by the service and returned to the user's application for local use. AWS KMS only creates and returns DEKs to users; the service does not store or manage DEKs.



Amazon Kinesis Data Analytics ([kinesisanalytics](#))

Amazon Kinesis Data Analytics is the easiest way for customers to analyze streaming data, gain actionable insights, and respond to business and customer needs in real time. Amazon Kinesis Data Analytics reduces the complexity of building, managing, and integrating streaming applications with other AWS services.

Amazon Kinesis Data Firehose ([firehose](#))

Amazon Kinesis Data Firehose is the easiest way to reliably load streaming data into data stores and analytics tools. It can capture, transform, and load streaming data into Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk, enabling near real-time analytics with existing business intelligence tools and dashboards customers are already using today.

Amazon Kinesis Data Streams ([kinesis](#))

Amazon Kinesis Streams is a platform for streaming data on AWS, so customers can load and analyze streaming data. Amazon Kinesis Streams also provides the ability to build custom streaming data applications for specialized needs.

Amazon Kinesis Video Streams ([kinesisvideo](#))

Amazon Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), playback, and other processing. Kinesis Video Streams automatically provisions and elastically scales the infrastructure needed to ingest streaming video data from millions of devices.

AWS Lambda ([lambda](#))

AWS Lambda lets customers run code without provisioning or managing servers on their own. AWS Lambda uses a compute fleet of Amazon Elastic Compute Cloud (Amazon EC2) instances across multiple Availability Zones in a region, which provides the high availability, security, performance, and scalability of the AWS infrastructure.

Amazon Lex ([runtime.lex](#), [models.lex](#))

Amazon Lex is a service for building conversational interfaces into any application using voice and text. Amazon Lex provides the advanced deep learning functionalities of automatic speech recognition (ASR) for converting speech to text, and natural language understanding (NLU) to recognize the intent of the text, to enable customers to build applications with highly engaging user experiences and lifelike conversational interactions. As a fully managed service, Amazon Lex scales automatically, so customers do not need to worry about managing infrastructure.

AWS License Manager ([license-manager](#))

AWS License Manager makes it easier to manage licenses in AWS and on-premises servers from software vendors. AWS License Manager allows customer's administrators to create customized licensing rules that emulate the terms of their licensing agreements, and then enforces these rules when an instance of EC2 gets launched. Customer administrators can use these rules to limit licensing violations, such as using



more licenses than an agreement stipulates or reassigning licenses to different servers on a short-term basis. The rules in AWS License Manager also enable customers to limit a licensing breach by stopping the instance from launching or by notifying the customer administrators about the infringement. Customer administrators gain control and visibility of all their licenses with the AWS License Manager dashboard and reduce the risk of non-compliance, misreporting, and additional costs due to licensing overages.

AWS License Manager integrates with AWS services to simplify the management of licenses across multiple AWS accounts, IT catalogs, and on-premises, through a single AWS account.

Amazon Macie ([macie](#))

Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes sensitive data such as personally identifiable information (PII) or intellectual property, and provides customers with dashboards and alerts that give visibility into how this data is being accessed or moved.

AWS Managed Services

AWS Managed Services provides ongoing management of a customer's AWS infrastructure. AWS Managed Services automates common activities such as change requests, monitoring, patch management, security, and backup services, and provides full-lifecycle services to provision, run, and support a customer's infrastructure.

Amazon Managed Streaming for Kafka ([kafka](#))

Amazon MSK is a fully managed service that makes it easy for customers to build and run applications that use Apache Kafka to process streaming data. Apache Kafka is an open-source platform for building real-time streaming data pipelines and applications. With Amazon MSK, customers can use Apache Kafka APIs to populate data lakes, stream changes to and from databases, and power machine learning and analytics applications.

Amazon MQ ([mq](#))

Amazon MQ is a managed message broker service for Apache ActiveMQ that sets up and operates message brokers in the cloud. Message brokers allow different software systems – often using different programming languages, and on different platforms – to communicate and exchange information. Amazon MQ manages the administration and maintenance of ActiveMQ, a popular open-source message broker.

Amazon Neptune ([neptune-db](#))

Amazon Neptune is a fast, reliable, fully managed graph database service that makes it easy to build and run applications that work with highly connected datasets. The core of Amazon Neptune is a purpose-built, high-performance graph database engine optimized for storing billions of relationships and querying the graph with milliseconds latency.



AWS OpsWorks ([opsworks](#))

AWS OpsWorks Stacks is an application and server management service. OpsWorks Stacks lets customers manage applications and servers on AWS and on-premises. With OpsWorks Stacks, customers can model their application as a stack containing different layers, such as load balancing, database, and application server. They can deploy and configure Amazon EC2 instances in each layer or connect other resources such as Amazon RDS databases.

AWS OpsWorks for Chef Automate or AWS OpsWorks for Puppet Enterprise ([opsworks-cm](#))

AWS OpsWorks for Chef Automate is a fully managed configuration management service that hosts Chef Automate, a suite of automation tools from Chef for configuration management, compliance and security, and continuous deployment. OpsWorks also maintains customers' Chef server by automatically patching, updating, and backing up customers' server.

AWS Organizations ([organizations](#))

AWS Organizations helps customers centrally govern their environment as customers' grow and scale their workloads on AWS. Whether customers are a growing startup or a large enterprise, Organizations helps customers to centrally manage billing; control access, compliance, and security; and share resources across customer AWS accounts.

Amazon Personalize ([personalize](#))

Amazon Personalize is a machine learning service that makes it easy for developers to create individualized recommendations for customers using their applications. With Amazon Personalize, customers provide an activity stream from their application – clicks, page views, signups, purchases, and so forth – as well as an inventory of the items they want to recommend, such as articles, products, videos, or music. Customers can also choose to provide Amazon Personalize with additional demographic information from their users such as age, or geographic location. Amazon Personalize will process and examine the data, identify what is meaningful, select the right algorithms, and train and optimize a personalization model that is customized for customer data. All data analyzed by Amazon Personalize is kept private and secure, and only used for your customized recommendations.

Amazon Pinpoint ([mobiletargeting](#))

Amazon Pinpoint helps customers engage with their customers by sending email, SMS, and mobile push messages. The customers can use Amazon Pinpoint to send targeted messages (such as promotional alerts and customer retention campaigns), as well as direct messages (such as order confirmations and password reset messages) to their customers.

Amazon Polly ([polly](#))

Amazon Polly is a service that turns text into lifelike speech, allowing customers to create applications that talk, and build entirely new categories of speech-enabled products. Amazon Polly is a Text-to-Speech service that uses advanced deep learning technologies to synthesize speech that sounds like a human voice.



Amazon QuickSight (`quicksight`)

Amazon QuickSight is a fast, cloud-powered business analytics service that makes it easy to build visualizations, perform ad-hoc analysis, and quickly get business insights from customers' data. Using this cloud-based service customers can connect to their data, perform advanced analysis, and create visualizations and dashboards that can be accessed from any browser or mobile device.

Amazon Redshift (`redshift`)

Amazon Redshift is a data warehouse service to analyze data using a customer's existing Business Intelligence (BI) tools. Amazon Redshift also includes Redshift Spectrum, allowing customers to directly run SQL queries against Exabytes of unstructured data in Amazon S3.

Amazon Rekognition (`rekognition`)

The easy-to-use Rekognition API allows customers to automatically identify objects, people, text, scenes, and activities, as well as detect any inappropriate content. Developers can quickly build a searchable content library to optimize media workflows, enrich recommendation engines by extracting text in images, or integrate secondary authentication into existing applications to enhance end-user security.

Amazon Relational Database Service (`rds`)

Amazon Relational Database Service enables customers to set up, operate, and scale a relational database in the cloud. Amazon RDS manages backups, software patching, automatic failure detection, and recovery. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching and backups.

AWS Resource Groups (`resource-groups`)

AWS Resource Groups is service that helps customers organize AWS resources into logical groupings. These groups can represent an application, a software component, or an environment.

AWS RoboMaker (`robomaker`)

AWS RoboMaker is a service that makes_it easy to develop, test, and deploy intelligent robotics applications at scale. RoboMaker extends the most widely used open-source robotics software framework, Robot Operating System (ROS), with connectivity to cloud services.

Amazon Route 53 (`route53`)

Amazon Route 53 provides managed Domain Name System (DNS) web service. Amazon Route 53 connects user requests to infrastructure running both inside and outside of AWS. Customers can use Amazon Route 53 to configure DNS health checks to route traffic to healthy endpoints or to independently monitor the health of their application and its endpoints.



Amazon SageMaker ([sagemaker](#))

Amazon SageMaker is a fully-managed platform that enables developers and data scientists to quickly and easily build, train, and deploy machine learning models at any scale. Amazon SageMaker removes the barriers that typically “slow down” developers who want to use machine learning.

AWS Secrets Manager ([secretsmanager](#))

AWS Secrets Manager helps customers protect secrets needed to access their applications, services, and IT resources. The service enables customers to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to Secrets Manager APIs, eliminating the need to hardcode sensitive information in plain text.

AWS Security Hub ([securityhub](#))

AWS Security Hub gives customers a comprehensive view of their high-priority security alerts and compliance status across AWS accounts. With Security Hub, customers can now have a single place that aggregates, organizes, and prioritizes their security alerts, or findings, from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, and Amazon Macie, as well as from AWS Partner solutions. Findings are visually summarized on integrated dashboards with actionable graphs and tables.

AWS Server Migration Service ([sms](#))

AWS Server Migration Service (SMS) is an agentless service which makes it easier and faster for customers to migrate thousands of on-premises workloads to AWS. AWS SMS allows customers to automate, schedule, and track incremental replications of live server volumes, making it easier for customers to coordinate large-scale server migrations.

AWS Serverless Application Repository ([serverlessrepo](#))

The AWS Serverless Application Repository is a managed repository for serverless applications. It enables teams, organizations, and individual developers to store and share reusable applications, and easily assemble and deploy serverless architectures in powerful new ways.

AWS Service Catalog ([servicecatalog](#))

AWS Service Catalog allows customers to create and manage catalogs of IT services that are approved for use on AWS. AWS Service Catalog allows customers to centrally manage commonly deployed IT services, and helps customers achieve consistent governance and meet their compliance requirements, while enabling users to quickly deploy only the approved IT services they need.

AWS Shield ([shield](#), [DDoSProtection](#))

AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards web applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection.



Amazon Simple Email Service (ses)

Amazon Simple Email Service is an email service that allows customers to send transactional email, marketing messages, or any other type of content. The main Amazon SES sending components are the frontend request router, backend control planes for feature configuration and access management, and a sending Mail Transfer Agent (MTA). Customers can also use Amazon SES to receive messages. The main Amazon SES receiving components are the receiving MTA, backend control planes for feature configuration and access management, and a rule-based message processor.

Amazon Simple Notification Service (sns)

Amazon Simple Notification Service is a web service to set up, operate, and send notifications. It provides developers the capability to publish messages from an application and deliver them to subscribers or other applications. Amazon SNS follows the “publish-subscribe” (pub-sub) messaging paradigm, with notifications being delivered to clients using a “push” mechanism.

Amazon Simple Queue Service (sqs)

Amazon Simple Queue Service offers a distributed hosted queue for storing messages as they travel between computers. By using Amazon SQS, developers can move data between distributed components of their applications that perform different tasks, without losing messages or requiring each component to be always available. Amazon SQS allows customers to build an automated workflow, working in close conjunction with Amazon EC2 and the other AWS infrastructure web services.

Amazon Simple Storage Service (s3)

Amazon Simple Storage Service provides a web services interface that can be used to store and retrieve data from anywhere on the web. To provide customers with the flexibility to determine how, when, and to whom they wish to expose the information they store in AWS, Amazon S3 APIs provide both bucket and object-level access controls, with defaults that only permit authenticated access by the bucket and/or object creator.

Amazon Simple Workflow Service (swf)

Amazon Simple Workflow Service is an orchestration service for building scalable distributed applications. Amazon SWF enables developers to architect and implement these tasks, run them in the cloud or on-premise and coordinate their flow.

Amazon SimpleDB (sdb)

Amazon SimpleDB is a non-relational data store that allows customers to store and query data items via web services requests. Amazon SimpleDB then creates and manages multiple geographically distributed replicas of data automatically to enable high availability and data durability.



AWS Snowball (snowball)

Snowball is a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of the AWS cloud. Using Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple and secure.

AWS Snowball Edge

AWS Snowball Edge is a 100TB data transfer device with on-board storage and compute capabilities. Customers can use Snowball Edge to move large amounts of data into and out of AWS, as a temporary storage tier for large local datasets, or to support local workloads in remote or offline locations. Snowball Edge connects to customers' existing applications and infrastructure using standard storage interfaces, streamlining the data transfer process and minimizing setup and integration.

AWS Snowmobile

AWS Snowmobile is an Exabyte-scale data transfer service used to move extremely large amounts of data to AWS. Customers can transfer their Exabyte data via a 45-foot long ruggedized shipping container, pulled by a semi-trailer truck. Snowmobile makes it easy to move massive volumes of data to the cloud, including video libraries, image repositories, or even a complete data center migration.

AWS Step Functions (states)

AWS Step Functions is a web service that enables customers to coordinate the components of distributed applications and microservices using visual workflows. Customers can build applications from individual components that each perform a discrete function, or task, allowing them to scale and change applications quickly.

AWS Storage Gateway (storagegateway)

The AWS Storage Gateway service connects customers' off-cloud software appliances with cloud-based storage. The service enables organizations to store data in AWS's highly durable cloud storage services: Amazon S3 and Amazon Glacier.

Amazon Systems Manager (ssm)

AWS Systems Manager formerly known as "Amazon EC2 Systems Manager" and "Amazon Simple Systems Manager", gives customers the visibility and control to their infrastructure on AWS. AWS Systems Manager provides customers a unified user interface so customers can view their operational data from multiple AWS services, and allows customers to automate operational tasks across the AWS resources.

Amazon Transcribe (transcribe)

Amazon Transcribe makes it easy for customers to add speech-to-text capability to their applications. Audio data is virtually impossible for computers to search and analyze. Therefore, recorded speech needs to be converted to text before it can be used in applications.



Amazon Transcribe uses a deep learning process called automatic speech recognition (ASR) to convert speech to text quickly. Amazon Transcribe can be used to transcribe customer service calls, to automate closed captioning and subtitling, and to generate metadata for media assets to create a fully searchable archive.

Amazon Transcribe automatically adds punctuation and formatting so that the output closely matches the quality of manual transcription at a fraction of the time and expense.

AWS Transfer for SFTP (`transfer`)

AWS Transfer for SFTP is a fully managed service that enables the transfer of files directly into and out of Amazon S3 using the Secure File Transfer Protocol (SFTP)—also known as Secure Shell (SSH) File Transfer Protocol.

Amazon Translate (`translate`)

Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. Amazon Translate allows customers to localize content - such as websites and applications - for international users, and to easily translate large volumes of text efficiently.

Amazon Virtual Private Cloud (Amazon VPC) (`ec2`)

Amazon Virtual Private Cloud enables customers to provision a logically isolated section of the AWS cloud where AWS resources can be launched in a virtual network defined by the customer. The VPN service provides end-to-end network isolation by using an IP address range of a customer's choice, and routing all of their network traffic between their Amazon VPC and another network designated by the customer via an encrypted Internet Protocol security (IPsec) VPN.

AWS WAF (`waf`)

AWS Web Application Firewall is a web application firewall that helps protect customer web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources.

Amazon WorkDocs (`workdocs`)

Amazon WorkDocs lets customers store all their files on one service. Users can share files, provide rich feedback, and access their files on WorkDocs from any device. WorkDocs encrypts data in transit and at rest, and offers powerful management controls, active directory integration, and near real-time visibility into file and user actions. The WorkDocs SDK allows users to use the same AWS tools they are already familiar with to integrate WorkDocs with AWS products and services, their existing solutions, third-party applications, or build their own.



Amazon WorkLink ([worklink](#))

Amazon WorkLink is a fully managed service that lets the customers provide their employees with secure, easy access to their internal corporate websites and web apps using their mobile phones. With Amazon WorkLink, employees can access internal web content as easily as they access any public website, without the hassle of connecting to their corporate network.

Amazon WorkMail ([workmail](#))

Amazon WorkMail is a managed business email and calendaring service with support for existing desktop and mobile email clients. It allows access to email, contacts, and calendars using Microsoft Outlook, a browser, or native iOS and Android email applications.

Amazon WorkSpaces ([workspaces](#))

Amazon WorkSpaces is a managed desktop computing service in the cloud. Amazon WorkSpaces enables customers to deliver a high quality desktop experience to end-users as well as help meet compliance and security policy requirements.

AWS X-Ray ([xray](#))

AWS X-ray helps developers analyze and debug production, distributed applications, such as those built using a microservices architecture. With X-ray, customers/developers can understand how their application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors. X-ray provides an end-to-end view of requests as they travel through the customers' application and shows a map of the application's underlying components. Customers/developers can use X-ray to analyze both applications in development and in production.



Attachment B – Principle Service Commitments and System Requirements

Overview

Amazon Web Services (AWS) designs its processes and procedures to meet its objectives for the AWS System. Those objectives are based on the service commitments that AWS makes to user entities (customers), the laws and regulations that govern the provision of the AWS System, and the financial, operational and compliance requirements that AWS has established for the services.

The AWS services are subject to relevant regulations, as well as state privacy security laws and regulations in the jurisdictions in which AWS operates.

Security, Availability and Confidentiality commitments to customers are documented and communicated in Service Level Agreements (SLAs) and other customer agreements, as well as in the description of the service offering provided on the AWS website. Security, Availability and Confidentiality commitments are standardized and include, but are not limited to, the following:

- Security and confidentiality principles inherent to the fundamental design of the AWS System are designed to appropriately restrict unauthorized internal and external access to data and customer data is appropriately segregated from other customers.
- Security and confidentiality principles inherent to the fundamental design of the AWS System are designed to safeguard data from within and outside of the boundaries of environments which store a customer's content to meet the service commitments.
- Availability principles inherent to the fundamental design of the AWS System are designed to replicate critical system components across multiple Availability Zones and authoritative backups are maintained and monitored to ensure successful replication to meet the service commitments.

Amazon Web Services establishes operational requirements that support the achievement of security, availability and confidentiality commitments, relevant laws and regulations, and other system requirements. Such requirements are communicated in AWS' system policies and procedures, system design documentation, and contracts with customers. Information security policies define an organization-wide approach to how systems and data are protected. These include policies around how the service is designed and developed, how the system is operated, how the internal business systems and networks are managed, and how employees are hired and trained. In addition to these policies, standard operating procedures have been documented on how to carry out specific manual and automated processes required in the operation and development of the Amazon Web Services System.

As an Infrastructure as a Service (IaaS) System, the AWS System is designed based on a shared responsibility model where both AWS and the customers are responsible for aspects of security, availability and confidentiality. Details of the responsibilities of customers can be found on the AWS website and in the Customer Agreement.



People

Amazon Web Services' organizational structure provides a framework for planning, executing and controlling business operations. Executive and senior leadership play important roles in establishing the Company's tone and core values. The organizational structure assigns roles and responsibilities to provide for adequate staffing, security, efficiency of operations, and segregation of duties. Management has also established authority and appropriate lines of reporting for key personnel.

The Company follows a structured on-boarding process to familiarize new employees with Amazon tools, processes, systems, security practices, policies and procedures. Employees are provided with the Company's Code of Business Conduct and Ethics and additionally complete annual Security & Awareness training to educate them as to their responsibilities concerning information security. Compliance audits are performed so that employees understand and follow established policies.

Data

AWS customers retain control and ownership of their own data. Customers are responsible for the development, operation, maintenance, and use of their content. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software.

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses techniques detailed in NIST 800-88 ("Guidelines for Media Sanitization") as part of the decommissioning process. All decommissioned hardware is sanitized and physically destroyed in accordance with industry-standard practices.

Availability

AWS is architected in a manner to maintain availability of its services through defined programs, processes, and procedures. The AWS Resiliency Program encompasses the processes and procedures by which AWS identifies, responds to, and recovers from a major event or incident within the environment. This program builds upon the traditional approach of addressing contingency management, incorporating elements of business continuity and disaster recovery plans while expanding to consider critical elements of proactive risk mitigation strategies. These strategies include engineering physically separate Availability Zones (AZs) and continuous infrastructure capacity planning.

Contingency plans and incident response playbooks are maintained to reflect emerging continuity risks and lessons learned. Plans are tested and updated through the course of business and the AWS Resiliency Program is annually reviewed and approved by senior leadership.



AWS has identified critical system components required to maintain the availability of the system and recover services in the event of an outage. These components are replicated across multiple availability zones; authoritative backups are maintained and monitored to ensure successful replication.

Service usage is continuously monitored, protecting infrastructure needs and supporting availability commitments and requirements. Additionally, AWS maintains a capacity planning model to assess infrastructure usage and demands.

Confidentiality

AWS is committed to protecting the security and confidentiality of its customers' content, defined as "Your Content" at <https://aws.amazon.com/agreement/>. AWS communicates its confidentiality commitment to customers in the [AWS Customer Agreement](#). AWS' systems and services are designed to enable authenticated AWS customers to access and manage their content by design through tools that allow customers to determine where content is stored, secure content in transit or at rest, initiate actions to remove or delete content, and manage access to AWS services and resources. AWS has also implemented technical and physical controls designed to prevent unauthorized access to or disclosure of content.

Internally, confidentiality requirements are communicated to employees through training and policies. Employees are required to attend security awareness training, which includes information, policies, and procedures related to protecting customers' content. AWS monitors the performance of third parties through periodic reviews, which evaluate performance against contractual obligations, including confidentiality commitments.

