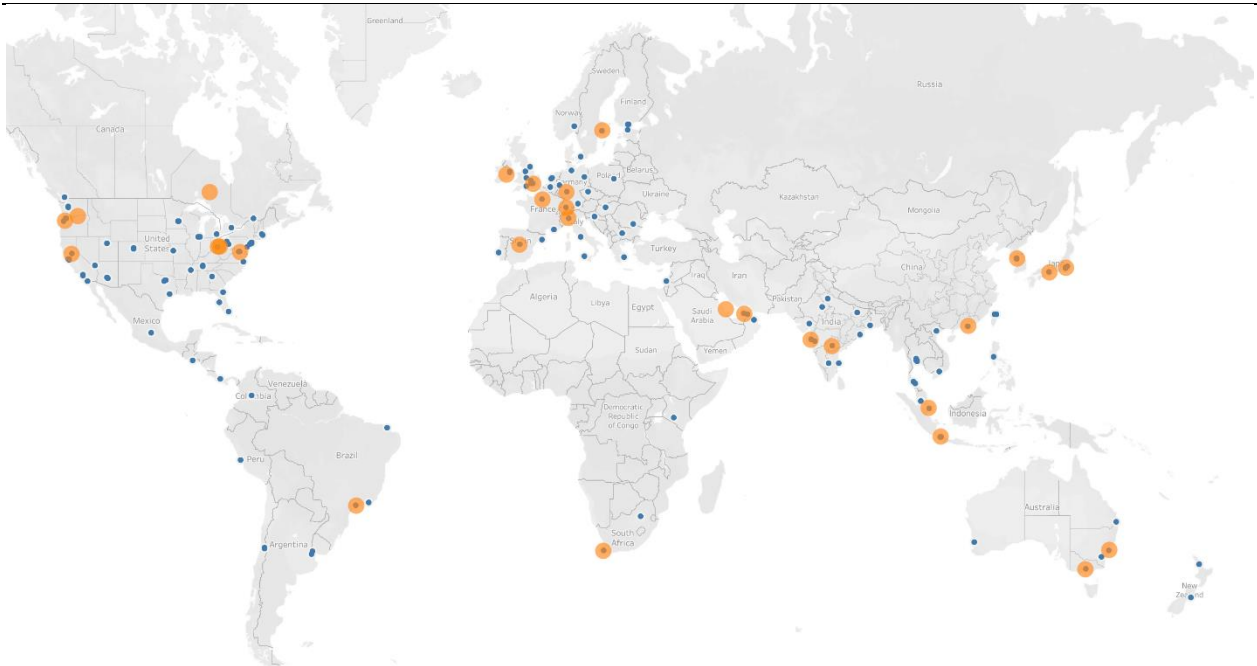




System and Organization Controls 3 (SOC 3) Report

Report on the Amazon Web Services System Relevant to Security, Availability, Confidentiality, and Privacy

For the Period October 1, 2022 – March 31, 2023



Report of Independent Accountants

To the Management of Amazon Web Services, Inc.,

Scope

We have examined management's assertion, contained within the accompanying "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, Confidentiality, and Privacy" (Assertion), that Amazon Web Services, Inc.'s (AWS) controls over the Amazon Web Services System (System) were effective throughout the period October 1, 2022 to March 31, 2023 to provide reasonable assurance that its principal service commitments and system requirements were achieved based on the criteria relevant to security, availability, confidentiality, and privacy (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 requirements

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, confidentiality, and privacy 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.

We are required to be independent of AWS and to meet our other ethical responsibilities, in accordance with the relevant ethical requirements related to our examination engagement.

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 October 1, 2022 to March 31, 2023, to provide reasonable assurance that its principal service commitments and system requirements were achieved based on the applicable trust services criteria.

A handwritten signature in black ink that reads 'Ernst & Young LLP'. The signature is written in a cursive, flowing style.

May 12, 2023



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, Confidentiality, and Privacy**

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

- Identifying the Amazon 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 system 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 requirements
- 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 October 1, 2022 to March 31, 2023 to provide reasonable assurance that the principal service commitments and system requirements were achieved based on the criteria relevant to security, availability, confidentiality, and privacy 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 Overview

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 in 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 of this system description includes the following services:

- AWS Amplify
- Amazon API Gateway
- Amazon AppFlow
- AWS Application Migration Service
- AWS App Mesh
- AWS App Runner
- Amazon AppStream 2.0
- AWS AppSync
- Amazon Athena
- AWS Audit Manager
- Amazon Augmented AI [Excludes Public Workforce and Vendor Workforce for all features]
- Amazon EC2 Auto Scaling
- AWS Backup
- AWS Batch
- Amazon Braket
- AWS Certificate Manager (ACM)
- AWS Chatbot
- Amazon Chime
- Amazon Chime SDK
- AWS Cloud9
- Amazon Cloud Directory
- AWS Cloud Map
- AWS CloudFormation
- Amazon CloudFront
- AWS CloudHSM
- AWS CloudShell
- AWS CloudTrail
- Amazon CloudWatch
- Amazon CloudWatch Logs
- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- AWS CodePipeline
- Amazon Cognito
- AWS IoT Events
- AWS IoT Greengrass
- AWS IoT SiteWise
- Amazon Kendra
- AWS Key Management Service (KMS)
- Amazon Keyspaces (for Apache Cassandra)
- Amazon Kinesis Data Analytics
- Amazon Kinesis Data Firehose
- Amazon Kinesis Data Streams
- Amazon Kinesis Video Streams
- AWS Lake Formation
- AWS Lambda
- Amazon Lex
- AWS License Manager
- Amazon Location Service
- Amazon Macie
- Amazon Managed Grafana
- AWS Managed Services
- Amazon Managed Streaming for Apache Kafka
- Amazon Managed Service for Prometheus
- Amazon Managed Workflows for Apache Airflow
- Amazon MemoryDB for Redis
- Amazon MQ
- Amazon Neptune
- AWS Network Firewall
- Amazon OpenSearch Service
- AWS OpsWorks Stacks
- AWS OpsWorks [includes Chef Automate, Puppet Enterprise]
- AWS Organizations
- AWS Outposts
- Amazon Personalize



- Amazon Comprehend
- Amazon Comprehend Medical
- AWS Config
- Amazon Connect
- AWS Control Tower
- AWS Data Exchange
- AWS Database Migration Service (DMS)
- AWS DataSync
- Amazon Detective
- Amazon DevOps Guru
- AWS Direct Connect
- AWS Directory Service [Excludes Simple AD]
- Amazon DocumentDB [with MongoDB compatibility]
- Amazon DynamoDB
- EC2 Image Builder
- AWS Elastic Beanstalk
- Amazon Elastic Block Store (EBS)
- Amazon Elastic Compute Cloud (EC2)
- Amazon Elastic Container Registry (ECR)
- Amazon Elastic Container Service – (both Fargate and EC2 launch types)
- AWS Elastic Disaster Recovery
- Amazon Elastic Kubernetes Service (EKS) [both Fargate and EC2 launch types]
- Amazon Elastic File System (EFS)
- Elastic Load Balancing (ELB)
- Amazon ElastiCache
- AWS Elemental MediaConnect
- AWS Elemental MediaConvert
- AWS Elemental MediaLive
- Amazon Elastic MapReduce (EMR)
- Amazon EventBridge
- Amazon FinSpace
- AWS Firewall Manager
- Amazon Forecast
- Amazon Fraud Detector
- FreeRTOS
- Amazon FSx
- Amazon S3 Glacier
- AWS Global Accelerator
- AWS Glue
- AWS Glue DataBrew
- Amazon GuardDuty
- AWS Health Dashboard
- Amazon Pinpoint
- Amazon Polly
- AWS Private Certificate Authority
- Amazon Quantum Ledger Database (QLDB)
- Amazon QuickSight
- Amazon Redshift
- Amazon Rekognition
- Amazon Relational Database Service (RDS)
- AWS Resource Access Manager (RAM)
- AWS Resource Groups
- AWS RoboMaker
- Amazon Route 53
- Amazon SageMaker [Excludes Studio Lab, Public Workforce and Vendor Workforce for all features]
- AWS Secrets Manager
- AWS Security Hub
- AWS Server Migration Service (SMS)
- AWS Serverless Application Repository
- AWS Service Catalog
- AWS Shield
- AWS Signer
- 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
- AWS IAM Identity Center (successor to AWS Single Sign-On)
- AWS Snowball
- AWS Snowball Edge
- AWS Snowmobile
- AWS Step Functions
- AWS Storage Gateway
- AWS Systems Manager
- Amazon Textract
- Amazon Timestream
- Amazon Transcribe
- AWS Transfer Family
- Amazon Translate
- Amazon Virtual Private Cloud (VPC)
- VM Import/Export



- Amazon HealthLake
- AWS Identity and Access Management (IAM)
- Amazon Inspector Classic
- AWS IoT Core
- AWS IoT Device Management
- AWS WAF
- Amazon WorkDocs
- Amazon WorkMail
- Amazon WorkSpaces
- Amazon WorkSpaces Web
- AWS X-Ray

More information about the in-scope services, can be found at <https://aws.amazon.com/compliance/services-in-scope/>

The scope of locations covered in this report includes the supporting data centers located in the following regions:

- **Australia:** Asia Pacific (Sydney) (ap-southeast-2), Asia Pacific (Melbourne) (ap-southeast-4)
- **Bahrain:** Middle East (Bahrain) (me-south-1)
- **Brazil:** South America (São Paulo) (sa-east-1)
- **Canada:** Canada (Central) (ca-central-1)
- **England:** Europe (London) (eu-west-2)
- **France:** Europe (Paris) (eu-west-3)
- **Germany:** Europe (Frankfurt) (eu-central-1)
- **Hong Kong:** Asia Pacific (ap-east-1)
- **India:** Asia Pacific (Mumbai) (ap-south-1), Asia Pacific (Hyderabad) (ap-south-2)
- **Ireland:** Europe (Ireland) (eu-west-1)
- **Italy:** Europe (Milan) (eu-south-1)
- **Indonesia:** Asia Pacific (Jakarta) (ap-southwest-3)
- **Japan:** Asia Pacific (Tokyo) (ap-northeast-1), Asia Pacific (Osaka) (ap-northeast-3)
- **Singapore:** Asia Pacific (Singapore) (ap-southeast-1)
- **South Africa:** Africa (Cape Town) (af-south-1)
- **South Korea:** Asia Pacific (Seoul) (ap-northeast-2)
- **Spain:** Europe (Spain) (eu-south-2)
- **Sweden:** Europe (Stockholm) (eu-north-1)
- **Switzerland:** Europe (Zurich) (eu-central-2)
- **United Arab Emirates:** Middle East (UAE) (me-central-1)
- **United States:** US East (Northern Virginia) (us-east-1), US East (Ohio) (us-east-2), US West (Oregon) (us-west-2), US West (Northern California) (us-west-1), AWS GovCloud (US-East) (us-gov-east-1), AWS GovCloud (US-West) (us-gov-west-1)

and the following AWS Edge locations in:

- Caba, Argentina
- General Pacheco, Argentina
- Koto City, Japan
- Osaka, Japan
- Shinagawa, Japan
- Wiltshire, United Kingdom
- Ashburn, United States
- Atlanta, United States



- Brisbane, Australia
- Canberra, Australia
- Melbourne, Australia
- Perth, Australia
- Sydney, Australia
- Vienna, Austria
- Brussels, Belgium
- Fortaleza, Brazil
- Rio de Janeiro, Brazil
- São Paulo, Brazil
- Sofia, Bulgaria
- Montreal, Canada
- Toronto, Canada
- Vancouver, Canada
- Huechuraba, Chile
- Santiago de Chile, Chile
- Bogotá, Colombia
- Zagreb, Croatia
- Prague, Czech Republic
- Ballerup, Denmark
- Tallinn, Estonia
- Espoo, Finland
- Helsinki, Finland
- Marseille, France
- Paris, France
- Berlin, Germany
- Dusseldorf, Germany
- Frankfurt, Germany
- Hamburg, Germany
- Munich, Germany
- Kropia, Greece
- Hong Kong, SAR
- Budapest, Hungary
- Bangalore, India
- Bhubaneswar, India
- Changodar, India
- Chennai, India
- Hyderabad, India
- Jaipur, India
- Kolkata, India
- Mumbai, India
- Navi Mumbai, India
- Delhi, India
- Patna, India
- Pune, India
- Nairobi, Kenya
- Anyang-si, Republic of Korea
- Seoul, Republic of Korea
- Kuala Lumpur, Malaysia
- Santiago de Querétaro, Mexico
- Amsterdam, Netherlands
- Schiphol-Rijk, Netherlands
- Auckland, New Zealand
- Christchurch, New Zealand
- Rosedale, New Zealand
- Oslo, Norway
- Barka, Oman
- Pueblo Nuevo, Panama
- Estación Terrena, Peru
- Santiago de Surco, Peru
- Manila, Philippines
- Warsaw, Poland
- Lisbon, Portugal
- Bucharest, Romania
- Singapore, Singapore
- Cape Town, South Africa
- Johannesburg, South Africa
- Barcelona, Spain
- Madrid, Spain
- Stockholm, Sweden
- Zurich, Switzerland
- Taipei, Taiwan
- New Taipei City, Taiwan
- Bangkok, Thailand
- Bannmai, Thailand
- Khlong Nueng, Thailand
- Pakkret, Thailand
- Tambon Klong Tamru, Thailand
- Thung Song Hong, Thailand
- Dubai, United Arab Emirates
- Fujairah, United Arab Emirates
- Birmingham, United Kingdom
- Brentford, United Kingdom
- Hull, United Kingdom
- Billerica, United States
- Boston, United States
- Chicago, United States
- Columbus, United States
- Dallas, United States
- Denver, United States
- Eden Prairie, United States
- El Segundo, United States
- Elk Grove Village, United States
- Franklin, United States
- Garland, United States
- Greenwood Village, United States
- Houston, United States
- Hillsboro, United States
- Irving, United States
- Itasca, United States
- Jacksonville, United States
- Jersey City, United States
- Kansas City, United States
- Las Vegas, United States
- Los Angeles, United States
- Memphis, United States
- Miami, United States
- Milpitas, United States
- Minneapolis, United States
- Nashville, United States
- New York City, United States
- Newark, United States
- Norfolk, United States
- North Las Vegas, United States
- Northlake, United States
- Portland, United States
- Palo Alto, United States
- Philadelphia, United States
- Phoenix, United States
- Piscataway, United States
- Pittsburgh, United States
- Rancho Cordova, United States
- Reston, United States
- Richardson, United States
- San Diego, United States
- San Jose, United States
- Seattle, United States
- Secaucus, United States



- Bekasi, Indonesia
- Jakarta, Indonesia
- Clonsaugh, Ireland
- Dublin, Ireland
- Haifa, Israel
- Milan, Italy
- Palermo, Italy
- Rome, Italy
- Inzai, Japan
- London, United Kingdom
- Manchester, United Kingdom
- Milton Keynes, United Kingdom
- Slough, United Kingdom
- Surrey, United Kingdom
- Swinton, United Kingdom
- Southfield, United States
- Tampa, United States
- Tempe, United States
- Tukwila, United States
- Vienna, United States
- West Valley City, United States
- HaNoi, Vietnam
- Ho Chi Minh, Vietnam

and the following Wavelength locations in:

- Toronto, Canada
- Berlin, Germany
- Dortmund, Germany
- Munich, Germany
- Osaka, Japan
- Tama, Japan
- Daejeon, South Korea
- Seoul, South Korea
- London, United Kingdom
- Salford, United Kingdom
- Alpharetta, United States
- Annapolis Junction, United States
- Aurora, United States
- Azusa, United States
- Charlotte, United States
- Eufless, United States
- Houston, United States
- Knoxville, United States
- Las Vegas, United States
- Minneapolis, United States
- New Berlin, United States
- Pembroke Pines, United States
- Plant City, United States
- Redmond, United States
- Rocklin, United States
- Southfield, United States
- Tempe, United States
- Wall Township, United States
- Westborough, United States

as well as Local Zone locations in:

- Buenos Aires, Argentina
- Perth, Australia
- Santiago, Chile
- Copenhagen, Denmark
- Helsinki, Finland
- Hamburg, Germany
- Kolkata, India**
- Delhi, India**
- Queretaro, Mexico
- Muscat, Oman
- Lima, Peru
- Warsaw, Poland
- Taipei, Taiwan
- Bangkok, Thailand
- Atlanta, United States
- Boston, United States
- Chicago, United States
- El Segundo, United States
- Greenwood Village, United States
- Hillsboro, United States
- Houston, United States
- Irvine, United States
- Kansas City, United States**
- Las Vegas, United States
- Lee's Summit, United States*
- Miami, United States
- Minneapolis, United States
- Philadelphia, United States
- Phoenix, United States
- Piscataway, United States
- Richardson, United States
- Seattle, United States

* This Local Zone is a private Local Zone and may not be available to all customers.

** This location has a public and private Local Zone. Private Local Zones may not be available to all customers.

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.

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

AWS Amplify is a set of tools and services that can be used together or on their own, to help front-end web and mobile developers build scalable full stack applications, powered by AWS. With Amplify, customers can configure app backend and connect applications in minutes, deploy static web apps in a few clicks and easily manage app content outside of AWS console. Amplify supports popular web frameworks including JavaScript, React, Angular, Vue, Next.js, and mobile platforms including Android, iOS, React Native, Ionic, and Flutter.

AWS Application Migration Service

AWS Application Migration Service is the primary service that AWS recommends for lift-and-shift applications to AWS. The service minimizes time-intensive, error-prone manual processes by automatically converting customers' source servers from physical, virtual, or cloud infrastructure to run natively on AWS. Customers are able to use the same automated process to migrate a wide range of applications to AWS without making changes to applications, their architecture, or the migrated servers.

Amazon API Gateway

Amazon API Gateway is a 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 API Gateway can execute AWS Lambda code in a customer's account, start AWS Step Functions state machines, or make calls to AWS Elastic Beanstalk, Amazon EC2, or web services outside of AWS with publicly accessible HTTP endpoints. Using the Amazon API Gateway console, customers can define customers' REST API and its associated resources and methods, manage customers' API lifecycle, generate customers' client SDKs, and view API metrics.

Amazon AppFlow

Amazon AppFlow is an integration service that enables customers to securely transfer data between Software-as-a-Service (SaaS) applications like Salesforce, SAP, Zendesk, Slack, and ServiceNow, and AWS services like Amazon S3 and Amazon Redshift. With AppFlow, customers can run data flows at enterprise scale at the frequency they choose - on a schedule, in response to a business event, or on demand. Customers are able to configure data transformation capabilities like filtering and validation to generate rich, ready-to-use data as part of the flow itself, without additional steps.



AWS App Mesh

AWS App Mesh is a service mesh that provides application-level networking which allows customer services to communicate with each other across multiple types of compute infrastructure. App Mesh gives customers end-to-end visibility and high availability for their applications. AWS App Mesh makes it easy to run services by providing consistent visibility and network traffic controls, which helps to deliver secure services. App Mesh removes the need to update application code to change how monitoring data is collected or traffic is routed between services. App Mesh configures each service to export monitoring data and implements consistent communications control logic across applications.

AWS App Runner

AWS App Runner is a service that makes it easy for developers to quickly deploy containerized web applications and APIs, at scale and with no prior infrastructure experience required. The service provides a simplified infrastructure-less abstraction for multi-concurrent web applications and API-based services. With App Runner, infrastructure components like build, load balancers, certificates and application replicas are managed by AWS. Customers simply provide their source-code (or a pre-built container image) and get a service endpoint URL in return against which requests can be made.

Amazon AppStream 2.0

Amazon AppStream 2.0 is an 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

AWS AppSync is a service that allows customers to easily develop and manage GraphQL APIs. Once deployed, AWS AppSync automatically scales the API execution engine up and down to meet API request volumes. AWS AppSync offers GraphQL setup, administration, and maintenance, with high availability serverless infrastructure built in.

Amazon 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 Audit Manager

AWS Audit Manager helps customers continuously audit AWS usage to simplify how customers manage risk and compliance with regulations and industry standards. AWS Audit Manager makes it easier to evaluate whether policies, procedures, and activities—also known as controls—are operating as intended. The service offers prebuilt frameworks with controls that are mapped to well-known industry standards and regulations, full customization of frameworks and controls, and automated collection and organization of evidence as designed by each control requirement.



Amazon Augmented AI (excludes Public Workforce and Vendor Workforce for all features)

Amazon Augmented AI is a machine learning service which makes it easy to build the workflows required for human review. Amazon A2I brings human review to all developers, removing the undifferentiated heavy lifting associated with building human review systems or managing large numbers of human reviewers whether it runs on AWS or not. The public and vendor workforce options of this service are not in scope for purposes of this report.

Amazon EC2 Auto Scaling

Amazon EC2 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

AWS Backup is a 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. Using AWS Backup, the customers can centrally configure backup policies and monitor backup activity for AWS resources, such as Amazon EBS volumes, Amazon RDS databases, Amazon DynamoDB tables, Amazon EFS file systems, and AWS Storage Gateway volumes. AWS Backup automates and consolidates backup tasks previously performed service-by-service, removing the need to create custom scripts and manual processes.

AWS 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.

Amazon Braket

Amazon Braket, the quantum computing service of AWS, is designed to help accelerate scientific research and software development for quantum computing. Amazon Braket provides everything customers need to build, test, and run quantum programs on AWS, including access to different types of quantum computers and classical circuit simulators and a unified development environment for building and executing quantum circuits. Amazon Braket also manages the classical infrastructure required for the execution of hybrid quantum-classical algorithms. When customers choose to interact with quantum computers provided by third-parties, Amazon Braket anonymizes the content, so that only content necessary to process the quantum task is sent to the quantum hardware provider. No AWS account information is shared and customer data is not stored outside of AWS.

AWS Certificate Manager (ACM)

AWS Certificate Manager (ACM) 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. SSL/TLS certificates are used to secure network communications and establish the identity of websites over the Internet as well as resources on private networks. AWS Certificate Manager removes the manual process of purchasing, uploading, and renewing SSL/TLS certificates.



AWS Chatbot

AWS Chatbot is an AWS service that enables DevOps and software development teams to use Slack or Amazon Chime chat rooms to monitor and respond to operational events in their AWS Cloud. AWS Chatbot processes AWS service notifications from Amazon Simple Notification Service (Amazon SNS), and forwards them to Slack or Amazon Chime chat rooms so teams can analyze and act on them. Teams can respond to AWS service events from a chat room where the entire team can collaborate, regardless of location.

Amazon Chime

Amazon Chime is a communications service that lets customers meet, chat, and place business calls inside and outside organizations, all using a single application. With Amazon Chime, customers can conduct and attend online meetings with HD video, audio, screen sharing, meeting chat, dial—in numbers, and in-room video conference support. Customer can use chat and chat rooms for persistent communications across desktop and mobile devices. Customers are also able to administer enterprise users, manage policies, and set up SSO or other advanced features in minutes using Amazon Chime management console.

Amazon Chime SDK

The Amazon Chime SDK is a set of real-time communications components that customers can use to quickly add messaging, audio, video, and screen sharing capabilities to their web or mobile applications. Customers can use the Amazon Chime SDK to build real-time media applications that can send and receive audio and video and allow content sharing. The Amazon Chime SDK works independently of any Amazon Chime administrator accounts and does not affect meetings hosted on Amazon Chime.

AWS Cloud9

AWS Cloud9 is an integrated development environment, or IDE. The AWS Cloud9 IDE offers a rich code-editing experience with support for several programming languages and runtime debuggers, and a built-in terminal. It contains a collection of tools that customers use to code, build, run, test, and debug software, and helps customers release software to the cloud. Customers access the AWS Cloud9 IDE through a web browser. Customers can configure the IDE to their preferences. Customers can switch color themes, bind shortcut keys, enable programming language-specific syntax coloring and code formatting, and more.

Amazon Cloud Directory

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. For example, customers can create an organizational chart that can be navigated through separate hierarchies for reporting structure, location, and cost center.

AWS Cloud Map

AWS Cloud Map is a cloud resource discovery service which allows customers to define custom names for their application resources. Cloud Map maintains the location of these changing resources to increase application availability.

Customers can register any application resource, such as databases, queues, microservices, and other cloud resources, with custom names. Cloud Map then constantly checks the health of resources to make



sure the location is up-to-date. The application can then query the registry for the location of the resources needed based on the application version and deployment environment.

AWS CloudFormation

AWS CloudFormation is a service to simplify provisioning of AWS resources such as Auto Scaling groups, ELBs, Amazon EC2, Amazon VPC, Amazon Route 53, and others. Customers author templates of the infrastructure and applications they want to run on AWS, and the AWS CloudFormation service automatically provisions the required AWS resources and their relationships as defined in these templates.

Amazon CloudFront

Amazon CloudFront is a fast content delivery network (CDN) web service that securely delivers data, videos, applications and APIs to customers globally with low latency and high-transfer speeds. CloudFront offers the most advanced security capabilities, including field level encryption and HTTPS support, seamlessly integrated with AWS Shield, AWS Web Application Firewall and Route 53 to protect against multiple types of attacks including network and application layer DDoS attacks. These services co-reside at edge networking locations – globally scaled and connected via the AWS network backbone – providing a more secure, performant, and available experience for the users.

CloudFront delivers customers' content through a worldwide network of Edge locations. When an end user requests content that customers serve with CloudFront, the user is routed to the Edge location that provides the lowest latency, so content is delivered with the best possible performance. If the content is already in that Edge location, CloudFront delivers it immediately.

In addition to Edge locations, CloudFront also uses Amazon Cloud Extension (ACE). ACE is a CloudFront infrastructure (single-rack version) deployed to a non-Amazon controlled facility, namely an internet service provider (ISP) or partner network. Qualifying Network Operators can deliver CloudFront content efficiently and cost effectively from within their network by deploying ACE in their data centers.

AWS CloudHSM

AWS CloudHSM is a service that allows customers to use dedicated hardware security module (HSM) appliances within the AWS cloud. AWS CloudHSM is designed for applications where the use of HSM appliances for encryption and key storage is mandatory.

AWS acquires these production HSM devices securely using the tamper evident authenticable bags from the vendors. These tamper evident authenticable bag serial numbers and production HSM serial numbers are verified against data provided out-of-band by the manufacturer and logged by approved individuals in tracking systems.

AWS CloudHSM allows customers to store and use encryption keys within HSM appliances in AWS data centers. With AWS CloudHSM, customers maintain full ownership, control, and access to keys and sensitive data while Amazon manages the HSM appliances in close proximity to customer applications and data. All HSM media is securely decommissioned and physically destroyed, verified by two personnel, prior to leaving AWS Secure Zones.



AWS CloudShell

AWS CloudShell is a browser-based shell used to securely manage, explore, and interact with your AWS resources. CloudShell is pre-authenticated with customer console credentials. Common development and operations tools are pre-installed, so no local installation or configuration is required. With CloudShell, customers can run scripts with the AWS Command Line Interface (AWS CLI), experiment with AWS service APIs using the AWS SDKs, or use a range of other tools to be productive. Customers can use CloudShell right from their browser.

AWS CloudTrail

AWS CloudTrail is a web service that records AWS activity for customers and delivers log files to a specified Amazon S3 bucket. 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 returned by the AWS service.

AWS CloudTrail provides a history of AWS API calls for customer accounts, including API calls made via the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services (such as AWS CloudFormation). The AWS API call history produced by AWS CloudTrail enables security analysis, resource change tracking, and compliance auditing.

Amazon CloudWatch

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 collects monitoring and operational data in the form of logs, metrics, and events, providing the customers with a unified view of AWS resources, applications and services that run on AWS, and on-premises servers.

Amazon CloudWatch Logs

Amazon CloudWatch Logs is a service used to monitor, store, and access log files from Amazon Elastic Compute Cloud (EC2) instances, AWS CloudTrail, Route 53 and other sources. CloudWatch Logs enables customers to centralize the logs from systems, applications and AWS services used in a single, highly scalable service. Customers can easily view them, search for patterns, filter on specific fields or archive them securely for future analysis. CloudWatch Logs enables customers to view logs, regardless of their source, as a single and consistent flow of events ordered by time, and to query them based on specific criteria.

AWS CodeBuild

AWS CodeBuild is a 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

AWS CodeCommit is a 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.



CodeCommit can be used to securely store anything from source code to binaries, and it works seamlessly with the existing Git tools.

AWS CodeDeploy

AWS CodeDeploy is a 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 CodeDeploy allows customers to rapidly release new features, helps avoid downtime during application deployment, and handles the complexity of updating the applications.

AWS CodePipeline

AWS CodePipeline is a 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

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. In addition, Amazon Cognito enables customers to save data locally on users' devices, allowing customers' applications to work even when the devices are offline. Customers can then synchronize data across users' devices so that their app experience remains consistent regardless of the device they use.

Amazon 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 insights and relationships in their unstructured data without machine learning experience. The service identifies the language of the text; extracts key phrases, places, people, brands, or events; understands how positive or negative the text is; analyzes text using tokenization and parts of speech; and automatically organizes a collection of text files by topic.

Amazon Comprehend Medical

Amazon Comprehend Medical is a HIPAA-eligible natural language processing (NLP) 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

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. With AWS Config,



customers can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine overall compliance against the configurations specified within the customers' internal guidelines. This enables customers to simplify compliance auditing, security analysis, change management, and operational troubleshooting.

Amazon Connect

Amazon Connect is an easy-to-use omnichannel cloud contact center that helps customers provide superior customer service across voice, chat, and tasks at lower cost than traditional contact center systems. Amazon Connect simplifies contact center operations, improves agent efficiency and lowers costs. Customers can setup a contact center in minutes that can scale to support millions of customers from the office or as a virtual contact center.

AWS Control Tower

AWS Control Tower provides the easiest way to set up and govern a new, secure, multi-account AWS environment based on AWS' 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 conform to customer policies. If customers are building a new AWS environment, starting out on the journey to AWS, starting a new cloud initiative, or are completely new to AWS, Control Tower will help customers get started quickly with governance and AWS' best practices built-in.

AWS Data Exchange

AWS Data Exchange makes it easy to find, subscribe to, and use third-party data in the cloud. Qualified data providers include category-leading brands. Once subscribed to a data product, customers can use the AWS Data Exchange API to load data directly into Amazon S3 and then analyze it with a wide variety of AWS analytics and machine learning services. For data providers, AWS Data Exchange makes it easy to reach the millions of AWS customers migrating to the cloud by removing the need to build and maintain infrastructure for data storage, delivery, billing, and entitlement.

AWS Database Migration Service (DMS)

AWS Database Migration Service (DMS) is a cloud service that enables customers to migrate relational databases, data warehouses, NoSQL databases, and other types of data stores. AWS DMS can be used to migrate data into the AWS Cloud, between on-premises instances (through AWS Cloud setup), or between combinations of cloud and on-premises setups. The service supports homogenous migrations within one database platform, as well as heterogeneous migrations between different database platforms. AWS Database Migration Service can also be used for continuous data replication with high-availability.

AWS DataSync

AWS DataSync is an online data transfer service that simplifies, automates and accelerates moving data between on-premises storage and AWS Storage services, as well as between AWS Storage services. DataSync can copy data between Network File System (NFS), Server Message Block (SMB) file servers, self-managed object storage, AWS Snowcone, Amazon Simple Storage Service (Amazon S3) buckets, Amazon EFS file systems and Amazon FSx for Windows File Server file systems. 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.



Amazon Detective

Amazon Detective allows customers to easily analyze, investigate, and quickly identify the root cause of potential security issues or suspicious activity. Amazon Detective collects log data from customer's AWS resources and uses machine learning, statistical analysis, and graph theory to build a linked set of data that enables customers to conduct faster and more efficient security investigations. AWS Security services can be used to identify potential security issues or findings.

Amazon Detective can analyze trillions of events from multiple data sources and automatically creates a unified, interactive view of the resources, users, and the interactions between them over time. With this unified view, customers can visualize all the details and context in one place to identify the underlying reasons for the findings, drill down into relevant historical activities, and quickly determine the root cause.

Amazon DevOps Guru

Amazon DevOps Guru is a service powered by machine learning (ML) that is designed to improve an application's operational performance and availability. DevOps Guru helps detect behaviors that deviate from normal operating patterns so customers can identify operational issues before they impact them.

DevOps Guru uses ML models informed by years of Amazon.com and AWS operational excellence to identify anomalous application behavior (for example, increased latency, error rates, resource constraints, and others) and helps surface critical issues that could cause potential outages or service disruptions. When DevOps Guru identifies a critical issue, it automatically sends an alert and provides a summary of related anomalies, the likely root cause, and context for when and where the issue occurred. When possible, DevOps Guru also helps provide recommendations on how to remediate the issue.

AWS Direct Connect

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 data center, office, or colocation environment.

AWS Directory Service (excludes Simple AD)

AWS Directory Service for Microsoft Active Directory, also known as AWS Managed Microsoft Active Directory (AD), enables customers' directory-aware workloads and AWS resources to use managed Active Directory in the AWS Cloud. AWS Managed Microsoft AD stores directory content in encrypted Amazon Elastic Block Store volumes using encryption keys. Data in transit to and from Active Directory clients is encrypted when it travels through Lightweight Directory Access Protocol (LDAP) over customers' Amazon Virtual Private Cloud (VPC) network. If an Active Directory client resides in an off-cloud network, the traffic travels to customers' VPC by a virtual private network link or an AWS Direct Connect link.

Amazon DocumentDB (with MongoDB compatibility)

Amazon DocumentDB (with MongoDB compatibility) is a fast, scalable, and highly available 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 DocumentDB implements the Apache 2.0 open source MongoDB 3.6 API by emulating the responses that a MongoDB client expects from a MongoDB server, allowing customers to use their existing MongoDB drivers and tools with Amazon DocumentDB.



Amazon DocumentDB uses a distributed, fault-tolerant, self-healing storage system that auto-scales up to 64 TB per database cluster.

Amazon 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.

Customers can create a database table that can store and retrieve data and serve any requested traffic. Amazon DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity specified and the amount of data stored, while maintaining consistent, fast performance. All data items are stored on Solid State Drives (SSDs) and are automatically replicated across multiple availability zones in a region.

EC2 Image Builder

EC2 Image Builder makes it easier to automate the creation, management, and deployment of customized, secure, and up-to-date “golden” server images that are pre-installed and pre-configured with software and settings to meet specific IT standards.

AWS Elastic Beanstalk

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 provided by underlying services such as Auto Scaling, Elastic Load Balancing, Amazon EC2, Amazon VPC, Amazon Route 53, and others.

Amazon Elastic Block Store (EBS)

Amazon Elastic Block Store (EBS) provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect customers from component failure. Amazon EBS allows customers to create storage volumes from 1 GB to 16 TB that can be mounted as devices by Amazon EC2 instances. 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 (e.g., a hard drive).

Amazon EBS volumes are presented as raw unformatted block devices that have been wiped prior to being made available for use. Wiping occurs before reuse. If customers have procedures requiring that all data be wiped via a specific method, customers can conduct a wipe procedure prior to deleting the volume for compliance with customer requirements. Amazon EBS includes Data Lifecycle Manager, which provides a simple, automated way to back up data stored on Amazon EBS volumes.

Amazon Elastic Compute Cloud (EC2)

Amazon Elastic Compute Cloud (EC2) 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.

Security within Amazon EC2 is provided on multiple levels: the operating system (OS) of the host layer, the virtual instance OS or guest OS, a firewall, and signed API calls. Each of these items builds on the capabilities of the others. This helps prevent data contained within Amazon EC2 from being intercepted by unauthorized systems or users and to provide Amazon EC2 instances themselves security without sacrificing flexibility of configuration. The Amazon EC2 service utilizes a hypervisor to provide memory and CPU isolation between virtual machines and controls access to network, storage, and other devices, and maintains strong isolation between guest virtual machines. Independent auditors regularly assess the security of Amazon EC2 and penetration teams regularly search for new and existing vulnerabilities and attack vectors.

AWS prevents customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software.

Amazon EC2 provides a complete firewall solution, referred to as a Security Group; this mandatory inbound firewall is configured in a default deny-all mode and Amazon EC2 customers must explicitly open the ports needed to allow inbound traffic.

Amazon provides a Time Sync function for time synchronization in EC2 Linux instances with the Coordinated Universal Time (UTC). It is delivered over the Network Time Protocol (NTP) and uses a fleet of redundant satellite-connected and atomic clocks in each region to provide a highly accurate reference clock via the local 169.254.169.123 IP address. Irregularities in the Earth's rate of rotation that cause UTC to drift with respect to the International Celestial Reference Frame (ICRF), by an extra second, are called leap second. Time Sync addresses this clock drift by smoothing out leap seconds over a period of time (commonly called leap smearing) which makes it easy for customer applications to deal with leap seconds.

Amazon Elastic Container Registry (ECR)

Amazon Elastic Container Registry is a Docker container image 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 (ECS) and Amazon Elastic Kubernetes Service (EKS).

Amazon Elastic Container Service (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. With simple API calls, customers can launch and stop Docker-enabled applications, query the complete state of customers' clusters, and access many familiar features like security groups, Elastic Load Balancing, EBS volumes, and IAM roles. Customers can use Amazon Elastic Container Service to schedule the placement of containers across customers' clusters based on customers' resource needs and availability requirements.

AWS Elastic Disaster Recovery

AWS Elastic Disaster Recovery minimizes downtime and data loss with the recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery. Customers can set up AWS Elastic Disaster Recovery on their source servers to initiate secure data



replication. Customer content is replicated to a staging area subnet in their AWS account, in the AWS Region they select. The staging area design reduces costs by using affordable storage and minimal compute resources to maintain ongoing replication. Customers can perform non-disruptive tests to confirm that implementation is complete. During normal operation, customers can maintain readiness by monitoring replication and periodically performing non-disruptive recovery and failback drills. If customers need to recover applications, they can launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time.

Amazon Elastic Kubernetes Service (EKS) (both Fargate and EC2 launch types)

Amazon Elastic Kubernetes Service (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 EKS is certified Kubernetes conformant so the customers can use existing tooling and plugins from partners and the Kubernetes community. Applications running on any standard Kubernetes environment are fully compatible and can be easily migrated to Amazon EKS.

Amazon Elastic File System (EFS)

Amazon Elastic File System (EFS) provides file storage for Amazon EC2 instances. EFS presents a network attached file system interface via the NFS v4 protocol. EFS file systems grow and shrink 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.

The customer is responsible for choosing which of their Virtual Private Clouds (VPCs) they want a file system to be accessed from by creating resources called mount targets. One mount target exists for each availability zone, which exposes an IP address and DNS name for mounting the customer's file system onto their EC2 instances. Customers then log into their EC2 instance and issue a 'mount' command, pointing at their mount target's IP address or DNS name. A mount target is assigned one or more VPC security groups to which it belongs. The VPC security groups define rules for what VPC traffic can reach the mount targets and in turn can reach the file system.

Elastic Load Balancing (ELB)

Elastic Load Balancing (ELB) 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

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. For example, an application running in Amazon EC2 can securely access an Amazon ElastiCache Cluster in the same region with very slight latency.

Using the Amazon ElastiCache service, customers create a Cache Cluster, which is a collection of one or more Cache Nodes, each running an instance of the Memcached, Redis Engine, or DAX Engine. A Cache Node is a self-contained environment which provides a fixed-size chunk of secure, network-attached RAM. Each Cache Node runs an instance of the Memcached, Redis Engine, or DAX Engine, and has its own DNS



name and port. Multiple types of Cache Nodes are supported, each with varying amounts of associated memory.

AWS Elemental 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. Customers can use MediaConnect to ingest live video from a remote event site (like a stadium), share video with a partner (like a cable TV distributor), or replicate a video stream for processing (like an over-the-top service). MediaConnect combines reliable video transport, highly secure stream sharing, and real-time network traffic and video monitoring that allow customers to focus on their content, not their transport infrastructure.

AWS Elemental MediaConvert

AWS Elemental MediaConvert is a file-based video transcoding service with broadcast-grade features. It allows customers to create video-on-demand (VOD) content for broadcast and multiscreen delivery at scale. The service combines advanced video and audio capabilities with a simple web services interface. With AWS Elemental MediaConvert, customers can focus on delivering media experiences without having to worry about the complexity of building and operating video processing infrastructure.

AWS Elemental MediaLive

AWS Elemental MediaLive is a live video processing service. Customers can create high-quality video streams for delivery to broadcast televisions and internet-connected multiscreen devices, like connected TVs, tablets, smart phones, and set-top boxes. The service works by encoding live video streams in real-time, taking a larger-sized live video source and compressing it into smaller versions for distribution to viewers. AWS Elemental MediaLive enables customers to focus on creating live video experiences for viewers without the complexity of building and operating video processing infrastructure.

Amazon Elastic MapReduce (EMR)

Amazon Elastic MapReduce (EMR) is a web service that provides managed Hadoop clusters on Amazon EC2 instances running a Linux operating system. 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. Amazon EMR actively manages clusters for customers, replacing failed nodes and adjusting capacity as requested. Amazon EMR securely and reliably handles a broad set of big data use cases, including log analysis, web indexing, data transformations (ETL), machine learning, financial analysis, scientific simulation, and bioinformatics.

Amazon EventBridge

Amazon EventBridge delivers a near real-time stream of events that describe changes in AWS resources. Customers can configure routing rules to determine where to send collected data to build application architectures that react in real time to the data sources. Amazon EventBridge becomes aware of operational changes as they occur and responds to these changes by taking corrective action as necessary by sending message to respond to the environment, activating functions, making changes and capturing state information



Amazon FinSpace

Amazon FinSpace is a data management and analytics service that makes it easy to store, catalog, and prepare financial industry data at scale. Amazon FinSpace reduces the time it takes for financial services industry (FSI) customers to find and access all types of financial data for analysis.

AWS Firewall Manager

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. As new applications are created, Firewall Manager also allows customers to bring new applications and resources into compliance with a common set of security rules from day one.

Amazon Forecast

Amazon Forecast uses machine learning to combine time series data with additional variables to build forecasts. 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 Fraud Detector

Amazon Fraud Detector helps detect suspicious online activities such as the creation of fake accounts and online payment fraud. Amazon Fraud Detector uses machine learning (ML) and 20 years of fraud detection expertise from AWS and Amazon.com to automatically identify fraudulent activity to catch more fraud, faster. With Amazon Fraud Detector, customers can create a fraud detection ML model with just a few clicks and use it to evaluate online activities in milliseconds.

FreeRTOS

FreeRTOS is an operating system for microcontrollers that makes small, low-power edge devices easy to program, deploy, secure, connect, and manage. FreeRTOS extends the FreeRTOS kernel, a popular open source operating system for microcontrollers, with software libraries that make it easy to securely connect the small, low-power devices to AWS cloud services like AWS IoT Core or to more powerful edge devices running AWS IoT Greengrass.

Amazon FSx

Amazon FSx provides 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). The customers don't have to worry about managing file servers and storage, as Amazon FSx automates the time-consuming administration tasks such as hardware provisioning, software configuration, patching,



and backups. Amazon FSx integrates the file systems with cloud-native AWS services, making them even more useful for a broader set of workloads.

Amazon S3 Glacier

Amazon S3 Glacier is an archival storage solution for data that is infrequently accessed for which retrieval times of several hours are suitable. Data in Amazon S3 Glacier is stored as an archive. Archives in Amazon S3 Glacier can be created or deleted, but archives cannot be modified. Amazon S3 Glacier archives are organized in vaults. All vaults created have a default permission policy that only permits access by the account creator or users that have been explicitly granted permission. Amazon S3 Glacier enables customers to set access policies on their vaults for users within their AWS Account. User policies can express access criteria for Amazon S3 Glacier on a per vault basis. Customers can enforce Write Once Read Many (WORM) semantics for users through user policies that forbid archive deletion.

AWS Global Accelerator

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 also makes it easier to manage customers' global applications by providing static IP addresses that act as a fixed entry point to customer applications hosted on AWS which eliminates the complexity of managing specific IP addresses for different AWS Regions and Availability Zones.

AWS Glue

AWS Glue is an 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 Glue DataBrew

AWS Glue DataBrew is a visual data preparation tool that makes it easy for data analysts and data scientists to clean and normalize data to prepare it for analytics and machine learning. Customers can choose from pre-built transformations to automate data preparation tasks, all without the need to write any code.

Amazon GuardDuty

Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect the customers' AWS accounts and workloads. With the cloud, the collection and aggregation of account and network activities is simplified, but it can be time consuming for security teams to continuously analyze event log data for potential threats. With GuardDuty, the customers now have an intelligent and cost-effective option for continuous threat detection in the AWS Cloud.

Amazon HealthLake

Amazon HealthLake is a service offering healthcare and life sciences companies a complete view of individual or patient population health data for query and analytics at scale. Using the HealthLake APIs, health organizations can easily copy health data, such as imaging medical reports or patient notes, from on-premises systems to a secure data lake in the cloud. HealthLake uses machine learning (ML) models to automatically understand and extract meaningful medical information from the raw data, such as medications, procedures, and diagnoses. HealthLake organizes and indexes information and stores it in



the Fast Healthcare Interoperability Resources (FHIR) industry standard format to provide a complete view of each patient's medical history.

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). Customers can grant other people permission to administer and use resources in their AWS account without having to share their password or access key. Customers can grant different permissions to different people for different resources. Customers can use IAM features to securely give applications that run on EC2 instances the credentials that they need in order to access other AWS resources, like S3 buckets and RDS or DynamoDB databases.

VM Import/Export

VM 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 on premises environment. This offering allows customers to leverage their existing investments in the virtual machines that customers have built to meet their IT security, configuration management, and compliance requirements by bringing those virtual machines into Amazon EC2 as ready-to-use instances. Customers can also export imported instances back to their off-cloud virtualization infrastructure, allowing them to deploy workloads across their IT infrastructure.

Amazon Inspector Classic

Amazon Inspector Classic is an automated security assessment service for customers seeking to improve the security and compliance of applications deployed on AWS. Amazon Inspector Classic automatically assesses applications for vulnerabilities or deviations from leading practices. After performing an assessment, Amazon Inspector Classic produces a detailed list of security findings prioritized by level of severity.

AWS IoT Core

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

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

Customers can also organize their devices, monitor and troubleshoot device functionality, query the state of any IoT device in the fleet, and send firmware updates over-the-air (OTA). AWS IoT Device Management is agnostic to device type and OS, so customers can manage devices from constrained microcontrollers to connected cars all with the same service. AWS IoT Device Management allows customers to scale their fleets and reduce the cost and effort of managing large and diverse IoT device deployments.



AWS IoT Events

AWS IoT Events is a 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, 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

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. With AWS IoT Greengrass, connected devices can run AWS Lambda functions, execute predictions based on machine learning models, keep device data in sync, and communicate with other devices securely – even when not connected to the Internet.

AWS IoT SiteWise

AWS IoT SiteWise is a service that enables industrial enterprises to collect, store, organize, and visualize thousands of sensor data streams across multiple industrial facilities. AWS IoT SiteWise includes software that runs on a gateway device that sits onsite in a facility, continuously collects the data from a historian or a specialized industrial server, and sends it to the AWS Cloud. With the service, customers can skip months of developing undifferentiated data collection and cataloging solutions, and focus on using their data to detect and fix equipment issues, spot inefficiencies, and improve production output.

Amazon Kendra

Amazon Kendra is an intelligent search service powered by machine learning. Kendra reimagines enterprise search for customer websites and applications so employees and customers can easily find content, even when it's scattered across multiple locations and content repositories.

AWS Key Management Service (KMS)

AWS Key Management Service (KMS) allows users to create and manage cryptographic keys. One class of keys, KMS keys, are designed to never be exposed in plaintext outside the service. KMS keys can be used to encrypt data directly submitted to the service. KMS keys 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.

AWS KMS is integrated with several AWS services so that users can request that resources in those services are encrypted with unique DEKs provisioned by KMS that are protected by a KMS key the user chooses at the time the resource is created. See in-scope services integrated with KMS at <https://aws.amazon.com/kms/>. Integrated services use the plaintext DEK from AWS KMS in volatile memory of service-controlled endpoints; they do not store the plaintext DEK to persistent disk. An encrypted copy of the DEK is stored to persistent disk by the service and passed back to AWS KMS for decryption each time the DEK is needed to decrypt content the customer requests. DEKs provisioned by AWS KMS are encrypted with a 256-bit key unique to the customer's account under a defined mode of AES – Advanced Encryption Standard.



When a customer requests AWS KMS to create a KMS key, the service creates a key ID for the KMS key and (optionally) key material, referred to as a backing key, which is tied to the key ID of the KMS key. The 256-bit backing key can only be used for encrypt or decrypt operations by the service. Customers can choose to have a KMS key ID created and then securely import their own key material to associate with the key ID. If the customer chooses to enable key rotation for a KMS key with a backing key that the service generated, AWS KMS will create a new version of the backing key for each rotation event, but the key ID remains the same. All future encrypt operations under the key ID will use the newest backing key, while all previous versions of backing keys are retained to decrypt ciphertexts created under the previous version of the key. Backing keys and customer-imported keys are encrypted under AWS-controlled keys when created/imported and they are only ever stored on disk in encrypted form.

All requests to AWS KMS APIs are logged and available in the AWS CloudTrail of the requester and the owner of the key. The logged requests provide information about who made the request, under which KMS key, and describes information about the AWS resource that was protected through the use of the KMS key. These log events are visible to the customer after turning on AWS CloudTrail in their account .

AWS KMS creates and manages multiple distributed replicas of KMS keys and key metadata automatically to enable high availability and data durability. KMS keys themselves are regional objects; plaintext versions of the KMS keys can only be used in the AWS region in which they were created. KMS keys are only stored on persistent disk in encrypted form and in two separate storage systems to ensure durability. When a plaintext KMS key is needed to fulfill an authorized customer request, it is retrieved from storage, decrypted on one of many AWS KMS hardened security appliances in the region, then used only in memory to execute the cryptographic operation (e.g., encrypt or decrypt). The plaintext key is then marked for deletion so that it cannot be re-used. Future requests to use the KMS key each require the decryption of the KMS key in memory for another one-time use.

AWS KMS endpoints are only accessible via TLS using the following cipher suites that support forward secrecy:

- ECDHE-RSA-AES256-GCM-SHA384
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-RSA-AES256-SHA384
- ECDHE-RSA-AES256-SHA
- ECDHE-RSA-AES128-SHA256
- ECDHE-RSA-3DES-CBC3-SHA
- DHE-RSA-AES256-SHA256 (ParamSize: 2048)
- DHE-RSA-AES128-SHA256 (ParamSize: 2048)
- DHE-RSA-AES256-SHA (ParamSize: 2048)
- DHE-RSA-AES128-SHA (ParamSize: 2048)

By design, no one can gain access to the plaintext KMS key material. Plaintext KMS keys are only ever present on hardened security appliances for the amount of time needed to perform cryptographic



operations under them. AWS employees have no tools to retrieve plaintext keys from these hardened security appliances. In addition, multi-party access controls are enforced for operations on these hardened security appliances that involve changing the software configuration or introducing new hardened security appliances into the service. These multi-party access controls minimize the possibility of an unauthorized change to the hardened security appliances, exposing plaintext key material outside the service, or allowing unauthorized use of customer keys. Additionally, key material used for disaster recovery processes by KMS are physically secured such that no AWS employee can gain access. Access attempts to recovery key materials are reviewed by authorized operators on a periodic basis. Roles and responsibilities for those cryptographic custodians with access to systems that store or use key material are formally documented and acknowledged.

Amazon Keyspaces (for Apache Cassandra)

Amazon Keyspaces (for Apache Cassandra) is a scalable, highly available Apache Cassandra-compatible database service. With Amazon Keyspaces, customers can run Cassandra workloads on AWS using the same Cassandra application code and developer tools that customers use today. Amazon Keyspaces is serverless and gives customers the performance, elasticity, and enterprise features customers need to operate business-critical Cassandra workloads at scale.

Amazon Kinesis Data Analytics

Amazon Kinesis Data Analytics is an easy 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. SQL users can easily query streaming data or build entire streaming applications using templates and an interactive SQL editor. Java developers can quickly build sophisticated streaming applications using open source Java libraries and AWS integrations to transform and analyze data in real-time.

Amazon Kinesis Data Firehose

Amazon Kinesis Data Firehose is a reliable way to load streaming data into data stores and analytics tools. It can capture, transform, and load streaming data into Amazon S3, Amazon Redshift, and Amazon OpenSearch Service enabling near real-time analytics with existing business intelligence tools and dashboards customers are already using today. The service automatically scales to match the throughput of the customers' data and requires no ongoing administration. It can also batch, compress, transform, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security.

Amazon Kinesis Data Streams

Amazon Kinesis Data Streams is a massively scalable and durable real-time data streaming service. Kinesis Data Streams can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs and location-tracking events. The collected data is available in milliseconds to enable real-time analytics use cases such as real-time dashboards, real-time anomaly detection, dynamic pricing and more.

Amazon Kinesis Video Streams

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. It also durably stores, encrypts, and indexes video data in the streams, and allows the customers



to access their data through easy-to-use APIs. Kinesis Video Streams enables the customers to playback video for live and on-demand viewing, and quickly build applications that take advantage of computer vision and video analytics.

Amazon Location Service

Amazon Location Service makes it easy for developers to add location functionality to applications without compromising data security and user privacy. With Amazon Location Service, customers can build applications that provide maps and points of interest, convert street addresses into geographic coordinates, calculate routes, track resources, and trigger actions based on location. Amazon Location Service uses high-quality geospatial data to provide maps, places, routes, tracking, and geofencing.

AWS Lake Formation

AWS Lake Formation is an integrated data lake service that makes it easy for customers to ingest, clean, catalog, transform, and secure their data and make it available for analysis and ML. AWS Lake Formation gives customers a central console where they can discover data sources, set up transformation jobs to move data to an Amazon Simple Storage Service (S3) data lake, remove duplicates and match records, catalog data for access by analytic tools, configure data access and security policies, and audit and control access from AWS analytic and ML services. Lake Formation automatically manages access to the registered data in Amazon S3 through services including AWS Glue, Amazon Athena, Amazon Redshift, Amazon QuickSight, and Amazon EMR to ensure compliance with customer defined policies. With AWS Lake Formation, customers can configure and manage their data lake without manually integrating multiple underlying AWS services.

AWS 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

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. Amazon Lex scales automatically, so customers do not need to worry about managing infrastructure.

AWS 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

Amazon Macie is a data security and data privacy service that uses machine learning and pattern matching to help customers discover, monitor, and protect their sensitive data in AWS.

Macie automates the discovery of sensitive data, such as personally identifiable information (PII) and financial data, to provide customers with a better understanding of the data that organization stores in Amazon Simple Storage Service (Amazon S3). Macie also provides customers with an inventory of the S3 buckets, and it automatically evaluates and monitors those buckets for security and access control. Within minutes, Macie can identify and report overly permissive or unencrypted buckets for the organization.

If Macie detects sensitive data or potential issues with the security or privacy of customer content, it creates detailed findings for customers to review and remediate as necessary. Customers can review and analyze these findings directly in Macie, or monitor and process them by using other services, applications, and systems.

Amazon Managed Grafana

Amazon Managed Grafana is a service for open source Grafana, providing interactive data visualization for monitoring and operational data. Using Amazon Managed Grafana, customers can visualize, analyze, and alarm on their metrics, logs, and traces collected from multiple data sources in their observability system, including AWS, third-party ISVs, and other resources across their IT portfolio. Amazon Managed Grafana offloads the operational management of Grafana by automatically scaling compute and database infrastructure as usage demands increase, with automated version updates and security patching. Amazon Managed Grafana natively integrates with AWS services so customers can securely add, query, visualize, and analyze their AWS data across multiple accounts and regions with a few clicks in the AWS Console. Amazon Managed Grafana integrates with AWS IAM Identity Center (successor to AWS SSO) and supports Security Assertion Markup Language (SAML) 2.0, so customers can set up user access to specific dashboards and data sources for only certain users in their corporate directory.

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 Service for Prometheus

Amazon Managed Service for Prometheus is a Prometheus-compatible monitoring and alerting service that facilitates monitoring of containerized applications and infrastructure at scale. The Cloud Native Computing Foundation's Prometheus project is an open source monitoring and alerting solution optimized for container environments. With Amazon Managed Service for Prometheus, customers can use the open source Prometheus query language (PromQL) to monitor and alert on the performance of containerized workloads, without having to scale and operate the underlying infrastructure. Amazon Managed Service for Prometheus automatically scales the ingestion, storage, alerting, and querying of



operational metrics as workloads grow or shrink, and it is integrated with AWS security services to enable fast and secure access to data.

Amazon Managed Workflows for Apache Airflow

Amazon Managed Workflows for Apache Airflow is a service for Apache Airflow that lets customers use their current, familiar Apache Airflow platform to orchestrate their workflows. Customers gain improved scalability, availability, and security without the operational burden of managing underlying infrastructure. Amazon Managed Workflows for Apache Airflow orchestrates customer's workflows using Directed Acyclic Graphs (DAGs) written in Python. Customers provide Amazon Managed Workflows for Apache Airflow an Amazon Simple Storage Service (S3) bucket where customer's DAGs, plugins, and Python requirements reside. Then customers can run and monitor their DAGs from the AWS Management Console, a command line interface (CLI), a software development kit (SDK), or the Apache Airflow user interface (UI).

Amazon Managed Streaming for Apache Kafka

Amazon Managed Streaming for Apache Kafka is a 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 MemoryDB for Redis

Amazon MemoryDB for Redis is a Redis-compatible, durable, in-memory database service. It is purpose-built for modern applications with microservices architectures.

Amazon MemoryDB for Redis is compatible with Redis, an open source data store, enabling customers to quickly build applications using the same flexible Redis data structures, APIs, and commands that they already use today. With Amazon MemoryDB for Redis, all of the customer's data is stored in memory, which enables the customer to achieve microsecond read and single-digit millisecond write latency and high throughput. Amazon MemoryDB for Redis also stores data durably across multiple Availability Zones (AZs) using a distributed transactional log to enable fast failover, database recovery, and node restarts. Delivering both in-memory performance and Multi-AZ durability, Amazon MemoryDB for Redis can be used as a high-performance primary database for microservices applications eliminating the need to separately manage both a cache and durable database.

Amazon 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. Messaging is the communications backbone that connects and integrates the components of distributed applications, such as order processing, inventory management, and order fulfillment for e-commerce. Amazon MQ manages the administration and maintenance of ActiveMQ, a popular open-source message broker.



Amazon Neptune

Amazon Neptune is a fast and reliable 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. Amazon Neptune supports popular graph models, Property Graph, and W3C's RDF, and their respective query languages Apache, TinkerPop Gremlin, and SPARQL, allowing customers to easily build queries that efficiently navigate highly connected datasets. Neptune powers graph use cases such as recommendation engines, fraud detection, knowledge graphs, drug discovery, and network security.

AWS Network Firewall

AWS Network Firewall is a stateful, managed, network firewall and intrusion detection and prevention service for customer virtual private cloud (VPC). With Network Firewall, customers can filter traffic at the perimeter of customer VPC. This includes filtering traffic going to and coming from an internet gateway, NAT gateway, or over VPN or AWS Direct Connect.

Amazon OpenSearch Service

Amazon OpenSearch Service is a service that makes it easy for the customer to deploy, secure, and operate OpenSearch cost effectively at scale. Amazon OpenSearch Service lets the customers pay only for what they use – there are no upfront costs or usage requirements. With Amazon OpenSearch Service, the customers get the ELK stack they need, without the operational overhead.

AWS OpsWorks Stacks

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. OpsWorks Stacks also lets customers set automatic scaling for their servers based on preset schedules or in response to changing traffic levels, and it uses lifecycle hooks to orchestrate changes as their environment scales.

AWS OpsWorks (includes Chef Automate, Puppet Enterprise)

AWS OpsWorks for Chef Automate is a 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 customer servers. OpsWorks eliminates the need for customers to operate their own configuration management systems or worry about maintaining its infrastructure. OpsWorks gives customers access to all of the Chef Automate features, such as configuration and compliance management, which customers manage through the Chef console or command line tools like Knife. It also works seamlessly with customers' existing Chef cookbooks.

AWS OpsWorks for Puppet Enterprise is a configuration management service that hosts Puppet Enterprise, a set of automation tools from Puppet for infrastructure and application management. OpsWorks also maintains customers' Puppet master server by automatically patching, updating, and backing up customers' servers. OpsWorks eliminates the need for customers to operate their own configuration management systems or worry about maintaining its infrastructure. OpsWorks gives



customers' access to all of the Puppet Enterprise features, which customers manage through the Puppet console. It also works seamlessly with customers' existing Puppet code.

AWS 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.

Using AWS Organizations, customers can automate account creation, create groups of accounts to reflect their business needs, and apply policies for these groups for governance. Customers can also simplify billing by setting up a single payment method for all of their AWS accounts. Through integrations with other AWS services, customers can use Organizations to define central configurations and resource sharing across accounts in their organization.

AWS Outposts

AWS Outposts is a service that extends AWS infrastructure, AWS services, APIs and tools to any data center, co-location space, or an on-premises facility for a consistent hybrid experience. AWS Outposts is ideal for workloads that require low latency access to on-premises systems, local data processing or local data storage. Outposts offer the same AWS hardware infrastructure, services, APIs and tools to build and run applications on premises and in the cloud. AWS compute, storage, database and other services run locally on Outposts and customers can access the full range of AWS services available in the Region to build, manage and scale on-premises applications. Service Link is established between Outposts and the AWS region by use of a secured VPN connection over the public internet or AWS Direct Connect.

AWS Outposts are configured with a Nitro Security Key (NSK) which is designed to encrypt customer content and give customers the ability to mechanically remove content from the device. Customer content is cryptographically shredded if a customer removes the NSK from an Outpost device.

Additional information about Security in AWS Outposts, including the shared responsibility model, can be found in the [AWS Outposts User Guide](#).

AWS Health Dashboard

AWS Health Dashboard provides alerts and remediation guidance when AWS is experiencing events that may impact customers. While the AWS Health Dashboard displays the general status of AWS services, AWS Health Dashboard gives customers a personalized view into the performance and availability of the AWS services underlying customer's AWS resources.

The dashboard displays relevant and timely information to help customers manage events in progress and provides proactive notification to help customers plan for scheduled activities. With AWS Health Dashboard, alerts are triggered by changes in the health of AWS resources, giving event visibility, and guidance to help quickly diagnose and resolve issues.

AWS Private Certificate Authority

AWS Private Certificate Authority (CA) is a managed private CA service enables customers to easily and securely manage the lifecycle of their private certificates. Private CA allows developers to be more agile by providing them APIs to create and deploy private certificates programmatically. Customers also have



the flexibility to create private certificates for applications that require custom certificate lifetimes or resource names. With Private CA, customers can create and manage private certificates for their connected resources in one place with a secure, pay as you go, managed private CA service.

Amazon Personalize

Amazon Personalize is a machine learning service that makes it easy for developers to create individualized recommendations for customers using their applications. Amazon Personalize makes it easy for developers to build applications capable of delivering a wide array of personalization experiences, including specific product recommendations, personalized product re-ranking and customized direct marketing. Amazon Personalize goes beyond rigid static rule-based recommendation systems and trains, tunes, and deploys custom machine learning models to deliver highly customized recommendations to customers across industries such as retail, media and entertainment.

Amazon Pinpoint

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

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 Quantum Ledger Database (QLDB)

Amazon Quantum Ledger Database (QLDB) is a ledger database that provides a transparent, immutable and cryptographically verifiable transaction log owned by a central trusted authority. Amazon QLDB can be used to track each and every application data change and maintains a complete and verifiable history of changes over time.

Amazon 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

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

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.



With a wide variety of use cases, Amazon Rekognition enables the customers to easily add the benefits of computer vision to the business.

Amazon Relational Database Service (RDS)

Amazon Relational Database Service (RDS) 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 Access Manager

AWS Resource Access Manager helps customers securely share their resources across AWS accounts, within their organization or organizational units (OUs) in AWS Organizations, and with IAM roles and IAM users for supported resource types. Customers are able to use AWS Resource Access Manager to share transit gateways, subnets, AWS License Manager license configurations, Amazon Route 53 Resolver rules, and more resource types.

AWS Resource Groups

AWS Resource Groups is a service that helps customers organize AWS resources into logical groupings. These groups can represent an application, a software component, or an environment. Resource groups can include more than fifty additional resource types, bringing the overall number of supported resource types to seventy-seven. Some of these new resource types include Amazon DynamoDB tables, AWS Lambda functions, AWS CloudTrail trails, and many more. Customers can now create resource groups that accurately reflect their applications, and take action against those groups, rather than against individual resources.

AWS 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. This includes AWS machine learning services, monitoring services, and analytics services that enable a robot to stream data, navigate, communicate, comprehend, and learn. RoboMaker provides a robotics development environment for application development, a robotics simulation service to accelerate application testing, and a robotics fleet management service for remote application deployment, update, and management.

Amazon Route 53

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 Route 53 enables customers to manage traffic globally through a variety of routing types, including Latency Based Routing, Geo DNS, and Weighted Round Robin, all of these routing types can be combined with DNS Failover. Amazon Route 53 also offers Domain Name Registration; customers can purchase and manage domain names such as example.com and Amazon Route 53 will automatically configure DNS settings for their domains. Amazon Route 53 sends automated requests over the internet to a resource, such as a web server, to verify that it is reachable, available, and functional. Customers also can choose to receive notifications when a resource becomes unavailable and choose to route internet traffic away from unhealthy resources.



Amazon SageMaker (excludes Studio Lab, Public Workforce and Vendor Workforce for all features)

Amazon SageMaker is a 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.

Amazon SageMaker removes the complexity that holds back developer success with the process of building, training, and deploying machine learning models at scale. Amazon SageMaker includes modules that can be used together or independently to build, train, and deploy a customer’s machine learning models.

AWS Secrets Manager

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. Secrets Manager offers secret rotation with built-in integration for Amazon RDS, Amazon Redshift, and Amazon DocumentDB. The service is also extensible to other types of secrets, including API keys and OAuth tokens. In addition, Secrets Manager allows customers to control access to secrets using fine-grained permissions and audit secret rotation centrally for resources in the AWS Cloud, third-party services, and on-premises.

AWS Security Hub

AWS Security Hub gives customers a comprehensive view of their high-priority security alerts and compliance status across AWS accounts. There are a range of powerful security tools at customers’ disposal, from firewalls and endpoint protection to vulnerability and compliance scanners. 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 Classic, 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

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. Using the Serverless Application Repository, customers do not need to clone, build, package, or publish source code to AWS before deploying it. Instead, customers can use pre-built applications from the Serverless Application Repository in their serverless architectures, helping customers reduce duplicated work, ensure organizational best practices, and get to market faster. Integration with AWS Identity and Access Management (IAM) provides resource-level control of each application, enabling customers to publicly share applications with everyone or privately share them with specific AWS accounts.



AWS Service Catalog

AWS Service Catalog allows customers to 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 multi-tier application architectures. 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

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 (SES) is a cost-effective, flexible and scalable email service that enables developers to send mail from within any application. Customers can configure Amazon SES to support several email use cases including transactional, marketing, or mass email communications. Amazon SES' flexible IP deployment and email authentication options help drive higher deliverability and protect sender reputation, while sending analytics to measure impact of each email. With Amazon SES, customers can send email securely, globally and at scale.

Amazon Simple Notification Service (SNS)

Amazon Simple Notification Service (SNS) 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. Using SNS requires defining a "Topic", setting policies on access and delivery of the Topic, subscribing consumers and designating delivery endpoints, and publishing messages to a Topic. Administrators define a Topic as an access point for publishing messages and allowing customers to subscribe to notifications. Security policies are applied to Topics to determine who can publish, who can subscribe, and to designate protocols supported.

Amazon Simple Queue Service (SQS)

Amazon Simple Queue Service (SQS) is a message queuing service that 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 SQS' main components consist of a frontend request-router fleet, a backend data-storage fleet, a metadata cache fleet, and a dynamic workload management fleet. User queues are mapped to one or more backend clusters. Requests to read, write, or delete messages come into the frontends. The frontends contact the metadata cache to find out which backend cluster hosts that queue and then connect to nodes in that cluster to service the request.



For authorization, Amazon SQS has its own resource-based permissions system that uses policies written in the same language used for AWS IAM policies. User permissions for any Amazon SQS resource can be given either through the Amazon SQS policy system or the AWS IAM policy system, which is authorized by AWS Identity and Access Management Service. Such policies with a queue are used to specify which AWS Accounts have access to the queue as well as the type of access and conditions.

Amazon Simple Storage Service (S3)

Amazon Simple Storage Service (S3) 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. Unless a customer grants anonymous access, the first step before a user can access Amazon S3 is to be authenticated with a request signed using the user's secret access key.

An authenticated user can read an object only if the user has been granted read permissions in an Access Control List (ACL) at the object level. An authenticated user can list the keys and create or overwrite objects in a bucket only if the user has been granted read and write permissions in an ACL at the bucket level. Bucket and object-level ACLs are independent; an object does not inherit ACLs from its bucket. Permissions to read or modify the bucket or object ACLs are themselves controlled by ACLs that default to creator-only access. Therefore, the customer maintains full control over who has access to its data. Customers can grant access to their Amazon S3 data to other AWS users by AWS Account ID or email, or DevPay Product ID. Customers can also grant access to their Amazon S3 data to all AWS users or to everyone (enabling anonymous access).

Network devices supporting Amazon S3 are configured to only allow access to specific ports on other Amazon S3 server systems. External access to data stored in Amazon S3 is logged and the logs are retained for at least 90 days, including relevant access request information, such as the data accessor IP address, object, and operation.

Amazon Simple Workflow Service (SWF)

Amazon Simple Workflow Service (SWF) is an orchestration service for building scalable distributed applications. Often an application consists of several different tasks to be performed in a particular sequence driven by a set of dynamic conditions. Amazon SWF enables developers to architect and implement these tasks, run them in the cloud or on-premise and coordinate their flow. Amazon SWF manages the execution flow such that tasks are load balanced across the workers, inter-task dependencies are respected, concurrency is handled appropriately, and child workflows are executed.

Amazon SWF enables applications to be built by orchestrating tasks coordinated by a decider process. Tasks represent logical units of work and are performed by application components that can take any form, including executable code, scripts, web service calls, and human actions.

Developers implement workers to perform tasks. They run their workers either on cloud infrastructure, such as Amazon EC2, or off-cloud. Tasks can be long-running, may fail, may timeout and may complete with varying throughputs and latencies. Amazon SWF stores tasks for workers, assigns them when workers are ready, tracks their progress, and keeps their latest state, including details on their completion. To orchestrate tasks, developers write programs that get the latest state of tasks from Amazon SWF and use



it to initiate subsequent tasks in an ongoing manner. Amazon SWF maintains an application's execution state durably so that the application can be resilient to failures in individual application components.

Amazon SWF provides auditability by giving customers visibility into the execution of each step in the application. The Management Console and APIs let customers monitor all running executions of the application. The customer can zoom in on any execution to see the status of each task and its input and output data. To facilitate troubleshooting and historical analysis, Amazon SWF retains the history of executions for any number of days that the customer can specify, up to a maximum of 90 days.

The actual processing of tasks happens on compute resources owned by the end customer. Customers are responsible for securing these compute resources, for example if a customer uses Amazon EC2 for workers then they can restrict access to their instances in Amazon EC2 to specific AWS IAM users. In addition, customers are responsible for encrypting sensitive data before it is passed to their workflows and decrypting it in their workers.

Amazon SimpleDB

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.

Data in Amazon SimpleDB is stored in domains, which are similar to database tables except that functions cannot be performed across multiple domains. Amazon SimpleDB APIs provide domain-level controls that only permit authenticated access by the domain creator.

Data stored in Amazon SimpleDB is redundantly stored in multiple physical locations as part of normal operation of those services. Amazon SimpleDB provides object durability by protecting data across multiple availability zones on the initial write and then actively doing further replication in the event of device unavailability or detected bit-rot.

AWS IAM Identity Center (successor to AWS Single Sign-On)

AWS IAM Identity Center (successor to AWS Single Sign-On) is a cloud-based service that simplifies managing SSO access to AWS accounts and business applications. Customers can control SSO access and user permissions across all AWS accounts in AWS Organizations. Customers can also administer access to popular business applications and custom applications that support Security Assertion Markup Language (SAML) 2.0. In addition, AWS IAM Identity Center offers a user portal where users can find all their assigned AWS accounts, business applications, and custom applications in one place.

AWS Signer

AWS Signer is a managed code-signing service to ensure the trust and integrity of customer code. Customers validate code against a digital signature to confirm that the code is unaltered and from a trusted publisher. With AWS Signer, customer security administrators have a single place to define their signing environment, including what AWS Identity and Access Management (IAM) role can sign code and in what regions. AWS Signer manages the code-signing certificate public and private keys and enables central management of the code-signing lifecycle.



AWS 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. Snowball Edge can cluster together to form a local storage tier and process customers' data on-premises, helping ensure their applications continue to run even when they are not able to access the cloud.

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. After a customer's data is loaded, Snowmobile is driven back to AWS where their data is imported into Amazon S3 or Amazon Glacier.

AWS Step Functions

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. Step Functions provides a reliable way to coordinate components and step through the functions of a customer's application. Step Functions provides a graphical console to visualize the components of a customer's application as a series of steps. It automatically triggers and tracks each step, and retries when there are errors, so the customer's application executes in order and as expected, every time. Step Functions logs the state of each step, so when things do go wrong, customers can diagnose and debug problems quickly.

AWS Storage Gateway

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

AWS Storage Gateway backs up data off-site to Amazon S3 in the form of Amazon EBS snapshots. AWS Storage Gateway transfers data to AWS and stores this data in either Amazon S3 or Amazon Glacier, depending on the use case and type of gateway used. There are three types of gateways: Tape, File, and Volume Gateways. The Tape Gateway allows customers to store more frequently accessed data in Amazon S3 and less frequently accessed data in Amazon Glacier.

The File Gateway allows customers to copy data to S3 and have those files appear as individual objects in S3. Volume gateways store data directly in Amazon S3 and allow customers to snapshot their data so that



they can access previous versions of their data. These snapshots are captured as Amazon EBS Snapshots, which are also stored in Amazon S3. Both Amazon S3 and Amazon Glacier redundantly store these snapshots on multiple devices across multiple facilities, detecting and repairing any lost redundancy. The Amazon EBS snapshot provides a point-in-time backup that can be restored off-cloud or on a gateway running in Amazon EC2, or used to instantiate new Amazon EBS volumes. Data is stored within a single region that customers specify.

AWS Systems Manager

AWS Systems Manager gives customers the visibility and control to their infrastructure on AWS. AWS Systems Manager provides customers a unified user interface so that customers can view their operational data from multiple AWS services, and it allows customers to automate operational tasks across the AWS resources.

With AWS Systems manager, customers can group resources, like Amazon EC2 instances, Amazon S3 buckets, or Amazon RDS instances, by application, view operational data for monitoring and troubleshooting, and take action on groups of resources.

Amazon Textract

Amazon Textract automatically extracts text and data from scanned documents. With Textract customers can quickly automate document workflows, enabling customers to process large volumes of document pages in a short period of time. Once the information is captured, customers can take action on it within their business applications to initiate next steps for a loan application or medical claims processing. Additionally, customers can create search indexes, build automated approval workflows, and better maintain compliance with document archival rules by flagging data that may require redaction.

Amazon Timestream

Amazon Timestream is a fast, scalable, and serverless time series database service for IoT and operational applications that makes it easy to store and analyze trillions of events per day up to 1,000 times faster and at as little as 1/10th the cost of relational databases. Amazon Timestream saves customers time and cost in managing the lifecycle of time series data by keeping recent data in memory and moving historical data to a cost optimized storage tier based upon user defined policies. Amazon Timestream's purpose-built query engine lets customers access and analyze recent and historical data together, without needing to specify explicitly in the query whether the data resides in the in-memory or cost-optimized tier. Amazon Timestream has built-in time series analytics functions, helping customers identify trends and patterns in data in real-time.

Amazon 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 Family

AWS Transfer Family enables the transfer of files directly into and out of Amazon S3. With the support for Secure File Transfer Protocol (SFTP)—also known as Secure Shell (SSH) File Transfer Protocol, the File Transfer Protocol over SSL (FTPS) and the File Transfer Protocol (FTP), the AWS Transfer Family helps the customers seamlessly migrate their file transfer workflows to AWS by integrating with existing authentication systems and providing DNS routing with Amazon Route 53.

Amazon Translate

Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. Neural machine translation is a form of language translation automation that uses deep learning models to deliver more accurate and more natural sounding translation than traditional statistical and rule-based translation algorithms. 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 (VPC)

Amazon Virtual Private Cloud (VPC) 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. Customers can connect their existing infrastructure to the network isolated Amazon EC2 instances within their Amazon VPC, including extending their existing management capabilities, such as security services, firewalls and intrusion detection systems, to include their instances via a Virtual Private Network (VPN) connection. 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.

Customers can optionally connect their VPC to the Internet by adding an Internet Gateway (IGW) or a NAT Gateway. An IGW allows bi-directional access to and from the internet for some instances in the VPC based on the routes a customer defines, which specify which IP address traffic should be routable from the internet, Security Groups, and Network ACLs (NACLs) which limit which instances can accept or send this traffic. Customers can also optionally configure a NAT Gateway which allows egress-only traffic initiated from a VPC instance to reach the internet, but not allow traffic initiated from the internet to reach VPC instances. This is accomplished by mapping the private IP addresses to a public address on the way out, and then map the public IP address to the private address on the return trip.

The objective of this architecture is to isolate AWS resources and data in one Amazon VPC from another Amazon VPC, and to help prevent data transferred from outside the Amazon network except where the customer has specifically configured internet connectivity options or via an IPsec VPN connection to their off-cloud network.

Further details are provided below:

- **Virtual Private Cloud (VPC):** An Amazon VPC is an isolated portion of the AWS cloud within which customers can deploy Amazon EC2 instances into subnets that segment the VPC's IP address range (as designated by the customer) and isolate Amazon EC2 instances in one subnet from



another. Amazon EC2 instances within an Amazon VPC are accessible to customers via Internal Gateway (IGW), Virtual Gateway (VGW), or VPC Peerings established to the Amazon VPC

- **IPsec VPN:** An IPsec VPN connection connects a customer's Amazon VPC to another network designated by the customer. IPsec is a protocol suite for securing Internet Protocol (IP) communications by authenticating and encrypting each IP packet of a data stream. Amazon VPC customers can create an IPsec VPN connection to their Amazon VPC by first establishing an Internet Key Exchange (IKE) security association between their Amazon VPC VPN gateway and another network gateway using a pre-shared key as the authenticator. Upon establishment, IKE negotiates an ephemeral key to secure future IKE messages. An IKE security association cannot be established unless there is complete agreement among the parameters. Next, using the IKE ephemeral key, two keys in total are established between the VPN gateway and customer gateway to form an IPsec security association. Traffic between gateways is encrypted and decrypted using this security association. IKE automatically rotates the ephemeral keys used to encrypt traffic within the IPsec security association on a regular basis to ensure confidentiality of communications.

AWS WAF

AWS WAF 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.

Customers can use AWS WAF to create custom rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that are designed for their specific application. New rules can be deployed within minutes, letting customers respond quickly to changing traffic patterns. Also, AWS WAF includes a full-featured API that customers can use to automate the creation, deployment, and maintenance of web security rules.

Amazon WorkDocs

Amazon WorkDocs is a secure content creation, storage and collaboration 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 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 WorkMail can be integrated with a customer's existing corporate directory and the customer controls both the keys that encrypt the data and the location (AWS Region) under which the data is stored.

Customers can create an organization in Amazon WorkMail, select the Active Directory they wish to integrate with, and choose their encryption key to apply to all customer content. After setup and validation of their mail domain, users from the Active Directory are selected or added, enabled for Amazon WorkMail, and given an email address identity inside the customer owned mail domain.



Amazon 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. When using Amazon WorkSpaces, an organization's data is neither sent to nor stored on end-user devices. The PCoIP protocol used by Amazon WorkSpaces uses an interactive video stream to provide the desktop experience to the user while the data remains in the AWS cloud or in the organization's off-cloud environment.

When Amazon WorkSpaces is integrated with a corporate Active Directory, each Workspace joins the Active Directory domain, and can be managed like any other desktop in the organization. This means that customers can use Active Directory Group Policies to manage their Amazon WorkSpaces and can specify configuration options that control the desktop, including those that restrict users' abilities to use local storage on their devices. Amazon WorkSpaces also integrates with customers' existing RADIUS server to enable multi-factor authentication (MFA).

Amazon WorkSpaces Web

Amazon WorkSpaces Web is an on-demand, managed service designed to facilitate secure browser access to internal websites and software-as-a-service (SaaS) applications. Customers can access the service from existing web browsers without infrastructure management, specialized client software, or virtual private network (VPN) solutions.

AWS X-Ray

AWS X-Ray helps developers analyze and debug production, distributed applications, such as those built using a microservices architecture. With X-Ray, customers or 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 or developers can use X-Ray to analyze both applications in development and in production.



Attachment B – Principal 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.
- Privacy principles inherent to the fundamental design of the AWS System are designed to protect the security and confidentiality of AWS customer content 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 points of 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 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 unauthorized access to assets. AWS uses techniques detailed in NIST 800-88 ("Guidelines for Media Sanitization") as part of the decommissioning process. All production media is securely decommissioned in accordance with industry-standard practices. Production media is not removed from AWS control until it has been securely decommissioned.

Availability

The AWS Resiliency Program encompasses the processes and procedures by which AWS identifies, responds to and recovers from a major availability event or incident within the AWS services environment. This program builds upon the traditional approach of addressing contingency management which incorporates elements of business continuity and disaster recovery plans and expands this to consider critical elements of proactive risk mitigation strategies such as engineering physically separate Availability Zones (AZs) and continuous infrastructure capacity planning.

AWS contingency plans and incident response playbooks are maintained and updated to reflect emerging risks and lessons learned from past incidents. Service team response plans are tested and updated through the due course of business, and the AWS Resiliency plan is tested, reviewed, and approved by senior leadership annually.

AWS has identified critical system components required to maintain the availability of the system and recover service in the event of outage. Critical system components (example: code bases) are backed up across multiple, isolated locations known as Availability Zones. Each Availability Zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. Common points of failure like generators and cooling equipment are not shared across Availability Zones. Additionally, Availability Zones are physically separate, and designed such that even extremely uncommon disasters such as fires, tornados or flooding should only affect a single Availability Zone. AWS replicates critical



system components across multiple Availability Zones and authoritative backups are maintained and monitored to ensure successful replication.

AWS continuously monitors service usage to project infrastructure needs to support availability commitments and requirements. AWS maintains a capacity planning model to assess infrastructure usage and demands at least monthly, and usually more frequently (e.g., weekly). In addition, the AWS capacity planning model supports the planning of future demands to acquire and implement additional resources based upon current resources and forecasted requirements

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' systems and services are designed to enable authenticated AWS customers to access and manage their content. AWS notifies customers of third-party access to a customer's content on the third-party access page located at <https://aws.amazon.com/compliance/third-party-access>. AWS may remove a customer's content when compelled to do so by a legal order, or where there is evidence of fraud or abuse as described in the Customer Agreement (<https://aws.amazon.com/agreement/>) and Acceptable Use Policy (<https://aws.amazon.com/aup/>). In executing the removal of a customer's content due to the reasons stated above, employees may render it inaccessible as the situation requires. For clarity, this capability to render customer content inaccessible extends to encrypted content as well.

In the course of AWS system and software design, build, and test of product features, a customer's content is not used and remains in the production environment. A customer's content is not required for the AWS software development life cycle. When content is required for the development or test of a service's software, AWS service teams have tools to generate mock, random data.

AWS knows customers care about privacy and data security. That is why AWS gives customers ownership and control over their content by design through tools that allow customers to determine where their content is stored, secure their content in transit or at rest, and manage access to AWS services and resources. AWS also implements technical and physical controls designed to prevent unauthorized access to or disclosure of a customer's content. As described in the Physical Security and Change Management areas in Section III of this report, AWS employs a number of controls to safeguard data from within and outside of the boundaries of environments which store a customer's content. As a result of these measures, access to a customer's content is restricted to authorized parties.

AWS contingency plans and incident response playbooks have defined and tested tools and processes to detect, mitigate, investigate, and assess security incidents. These plans and playbooks include guidelines for responding to potential data breaches in accordance with contractual and regulatory requirements. AWS security engineers follow a protocol when responding to potential data security incidents. The protocol involves steps, which include validating the presence customer content within the AWS service (without actually viewing the data), determining the encryption status of a customer's content, and determining improper access to a customer's content to the extent possible.

During the course of their response, the security engineers document relevant findings in internal tools used to track the security issue. AWS Security Leadership is regularly apprised of all data security issue investigations. In the event there are positive indicators that customer content was potentially accessed



by an unintended party, a security engineer engages AWS Security Leadership and the AWS Legal team to review the findings. AWS Security Leadership and the Legal team review the findings and determine if a notifiable data breach has occurred pursuant to contractual or regulatory obligations. If confirmed, affected customers are notified in accordance with the applicable reporting requirement.

Vendors and third parties with restricted access, that engage in business with Amazon, are subject to confidentiality commitments as part of their agreements with Amazon. Confidentiality commitments are included in agreements with vendors and third parties with restricted access are reviewed by AWS and the third party at time of contract creation or renewal. AWS monitors the performance of third parties through periodic reviews on a risk-based approach, which evaluate performance against contractual obligations.

AWS communicates its confidentiality commitments to customers on its public website located at <https://aws.amazon.com/compliance/third-party-access/> for contractors and <https://aws.amazon.com/compliance/sub-processors/> for sub-processors. The effective date of the policy is communicated there and updated periodically. Before AWS authorizes and permits any new subcontractor to access any customer content, AWS will update this website to inform customers. Vendor confidentiality commitments are governed by the terms of the contract between AWS and the vendor.

Internally, confidentiality requirements are communicated to employees through training and policies. Employees are required to attend Amazon Security Awareness (ASA) training, which includes policies and procedures related to protecting a customer's content. Confidentiality requirements are included in the Data Handling and Classification Policy. Policies are reviewed and updated at least annually.

Privacy

AWS classifies customer data into two categories: customer content and account information. AWS defines customer content as software (including machine images), data, text, audio, video, or images that a customer or any end user transfers to AWS for processing, storage, or hosting by AWS services in connection with that customer's account, and any computational results that a customer or any end user derives from the foregoing through their use of AWS services. For example, customer content includes content that a customer or any end user stores in Amazon Simple Storage Service (S3). The terms of the AWS Customer Agreement (<https://aws.amazon.com/agreement/>) and AWS Service Terms (<https://aws.amazon.com/service-terms/>) apply to customer content.

Account information is information about a customer that a customer provides to AWS in connection with the creation or administration of a customer account. For example, account information includes names, user names, phone numbers, email addresses, and billing information associated with a customer account. Any information submitted by the customer that AWS needs in order to provide services to the customer or in connection with the administration of customer accounts, is not in-scope for this report.

The AWS Privacy Notice is available from the AWS website at <https://aws.amazon.com/privacy/>. The AWS Privacy Notice is reviewed by the AWS Legal team, and is updated as required to reflect Amazon's current business practices and global regulatory requirements. The Privacy Notice describes how AWS collects and uses a customer's personal information in relation to AWS websites, applications, products, services, events, and experiences. The Privacy Notice does not apply to customer content.



As part of the AWS account creation and activation process, AWS customers are informed of the AWS Privacy Notice and are required to accept the Customer Agreement, including the terms and conditions related to the collection, use, retention, disclosure, and disposal of their data. Customers are responsible for determining what content to store within AWS, which may include personal information. Without the acceptance of the Customer Agreement, customers cannot sign up to use the AWS services.

The AWS Customer Agreement informs customers of the AWS data security and privacy commitments prior to activating an AWS account and is made available to customers to review at any time on the AWS website.

The customer determines what data is entered into AWS services and has the ability to configure the appropriate security and privacy settings for the data, including who can access and use the data. Further, the customer is able to choose not to provide certain data. Additionally, the customer manages notification or consent requirements, and maintains the accuracy of the data.

Additionally, the AWS Customer Agreement notes how AWS shares, secures, and retains customer content. AWS also informs customers of updates to the Customer Agreement by making it available on its website and providing the last updated date. Customers should check the Customer Agreement website frequently for any changes to the Customer Agreement.

AWS does not store any customer cardholder data obtained from customers. Rather, AWS passes the customer cardholder data and sends it immediately to the Amazon Payments Platform, the PCI-certified platform that Amazon uses for all payment processing. This platform returns a unique identifier that AWS stores and uses for all future processing. The Amazon Payments Platform sits completely outside of the AWS boundary and is run by the larger Amazon entity. It is not an AWS service, but it is utilized by the larger Amazon entity for payment processing. As such, the Amazon payment platform is not in-scope for this report.

AWS offers customers the ability to update their communication preferences through the AWS console or via the AWS Email Preference Center. When customers update their communication preferences using their email, their updated preferences are saved. Customers can unsubscribe from AWS marketing emails within the AWS console. AWS Customers will still receive important account-related notifications from AWS, such as monthly billing statements, or if there are significant changes to a service that customers use.

AWS provides authenticated customers the ability to access, update, and confirm their data. Denial of access will be communicated using the AWS console. Customers can sign into to their AWS accounts through the AWS console to view and update their data.

AWS (or Amazon) does not disclose customer information in response to government demands unless we're required to do so to comply with a legally valid and binding order. AWS Legal reviews and maintains records of all the information requests, which lists information on the types and volume of information requested. Unless AWS is prohibited from doing so or there is clear indication of illegal conduct in connection with the use of Amazon products or services, AWS notifies customers before disclosing customer content so they can seek protection from disclosure. AWS shares customer content only as described in the AWS Customer Agreement.



AWS may produce non-content and/or content information in response to valid and binding law enforcement and governmental requests, such as subpoenas, court orders, and search warrants. “Non-content information” means customer information such as name, address, email address, billing information, date of account creation, and service usage information. Content information” includes the content that a customer transfers for processing, storage, or hosting in connection with AWS services and any computational results. AWS records customer information requests to maintain a complete, accurate, and timely record of such requests.

If required, customers are responsible for providing notice to the individuals whose data the customer collects and uses within AWS. AWS is not responsible for providing such notice to or obtaining consent from these individuals and is only responsible for communicating its privacy commitments to AWS customers, which is provided during the account creation and activation process.

AWS has documented an incident response policy and plan which outlines an organized approach for responding to security breaches and incidents. The AWS Security team is responsible for monitoring systems, tracking issues, and documenting findings of security-related events. Records are maintained for security breaches and incidents, which includes status information required for supporting forensic activities, trend analysis, and evaluation of incident details.

As part of the process, potential breaches of customer content are investigated and escalated to AWS Security and AWS Legal. Affected customers and regulators are notified of breaches and incidents where legally required. Customers can subscribe to the AWS Security Bulletins page, which provides information regarding identified security issues. AWS notifies affected customers and regulators of breaches and incidents as legally required in accordance with team processes.

AWS retains and disposes of customer content in accordance with the Customer Agreement and the AWS Data Classification and Handling Policy. When a customer terminates their account or contract with AWS, the account is put under isolation; after which within 90 days, customers can restore their accounts and related content. AWS services hosting customer content are designed to retain customer content until the contractual obligation to retain a customers’ content ends or a customer-initiated action to remove or delete the content is taken. When a customer requests data to be deleted, AWS utilizes automated processes to detect that request and make the content inaccessible. After the deletion is complete, automated actions are taken on deleted content to render the content unreadable.

AWS performs application security reviews for Third-Party systems that integrate with AWS, to ascertain security risks are identified and mitigated. A typical security review considers privacy components such as retention period, use, and collection of data as applicable. The review starts with a system owner initiating a review request to the dedicated AWS Vendor Security (AVS) team, and submitting detailed information about the artifacts being reviewed. A Security review is required if an AWS Team engages with a new external party for collecting data or modifications to existing systems.

During this process, the AVS team determines the granularity of review required based on the artifact’s design, threat model, and impact to AWS’ risk profile. They provide security guidance, validate security assurance material, and meet with external parties to discuss their penetration tests, Software Development Life Cycle, change management processes, and other operating security controls. They work with the system owner to identify, prioritize, and remediate security findings. The AVS team collaborates with AWS Legal as needed to validate that changes are in-line with AWS privacy policies. The AVS team



provides their final approval after they have adequately assessed the risks and worked with the requester to implement security controls to mitigate identified risks.