GLOBAL CLOUD INFRASTRUCTURE OF AMAZON WEB SERVICES

Summary by Damian Ndunda © May 2019

TABLE OF CONTENTS

FOREWORD

This booklet is intended to be a short and simple summary of Amazon Web Services. I hope it will truly assist someone who yearns to learn about AWS. For further awareness read the books I used in the reference section. I've tried to use the original author's words as much as possible.

I have written many technology related summaries since I wrote my first one about MSDOS (Microsoft Disk Operating System).The days when COBOL (Computer Business Oriented Programming Language) and Fortran (Formula Translation) were the main languages to learn. These Languages are not common today but still run business softwares for Fortune 500 companies despite them being over 60 years old. The summaries never made it online because internet technology was different back then. Through that whole time there is a book that advices me to Love God and not to kill or steal or be immoral among other instructions that has remained relevant. As technology went obsolete or was updated this instructions remain true and unchanged.

In a few years new technologies come up and the old is upgraded or outdated, but if you knew the old you probably can understand the new better. Just as if you can understand Jesus in the Old Testament prophesies, then you probably understand him better in the New Testament. Otherwise one day one is called intelligent then if they fail to continue learning they fail to know about the new. Example of technology changing would be IPV4 and IPV6. LANS, VLANS, VPN, VPC. Microsoft NetBIOS/NetBEUI, TCP, NAT and AWS Route54. B, C, C++, Java programming languages. Major Certifications from Amazon, Oracle, Microsoft, Cisco, Linux, Google, Other are upgraded every four years on average. Now (today) we have Web 2.0 heading to Web 3.0 making possible not just to share easily but also to run our applications online from the cloud. Facebook (developed with PHP) is another good example of a software cloud run corporation. Others like Bitrix24, Salesforce, Zoho, Weebly, may not be common names to many.

AWS cloud combines these technologies as Developers will run applications from the cloud, network administrators will configure the cloud, Relational Database Administrators will CRUD (Create, Read, Update, Delete) in the cloud, Hardware engineers will maintain the servers, Sales managers will move products from the cloud, Web developers will have stronger development tools that allow customer feedback tracking and reporting, Software run from cloud curbs piracy and introduces a new level of security encryption, CEOs get up-to-date company information from all there international branches at their finger tips for decision making, and many others.

INTRODUCTION

Amazon Web Services (AWS) is a platform of web services offering solutions for computing, storing, and networking, at different layers of abstraction provided by Amazon inc. Amazon inc is an American multinational technology company based in Seattle, Washington, that focuses on e-commerce, cloud computing, digital streaming and artificial intelligence. One can use cloud services to host web sites, run enterprise applications, and mine tremendous amounts of data. The term *web service* means services can be controlled via a web interface.

Virtualization has become a widely accepted way to reduce operating costs and increase the reliability of enterprise IT. In addition, grid computing makes a completely new class of analytics, data crunching, and business intelligence tasks possible that were previously cost and time prohibitive.

The data centers of AWS are distributed throughout the United States, Europe, Asia, and South America.

AWS DATA CENTER LOCATIONS

Figure 1.1 AWS data center locations

Bernard Golden, “Amazon Web Services (AWS) Hardware,” For Dummies, http://mng.bz/k6lT.

Witting A, Witting M, (2016) p 4

WHAT IS "CLOUD COMPUTING"?

Cloud computing, or the cloud, is a metaphor for supply and consumption of IT resources.

It builds on many of the advances in the IT industry over the past decade and presents significant opportunities for organizations to shorten time to market and reduce costs. With cloud computing, organizations can consume shared computing and storage resources rather than building, operating, and improving infrastructure on their own. The speed of change in markets creates significant pressure on the enterprise IT infrastructure to adapt and deliver. Cloud computing provides fresh solutions to address these changes. As defined by Gartner1, "Cloud computing is a style of computing where scalable and elastic IT enabled capabilities are delivered as a service to external customers using Internet technologies. " Flexible, secure, and cost-effective IT infrastructure,

Gartner IT Glossary, http://www.gartner.com/it-glossary/cloud-computing.

> *Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.*
>
> —The NIST Definition of Cloud Computing,
> National Institute of Standards and Technology

Clouds are often divided into the following types:

- ***Public***—A cloud managed by an organization and open to use by the general public
- ***Private***—A cloud that virtualizes and shares the IT infrastructure within a single organization
- ***Hybrid***—A mixture of a public and a private cloud

Cloud computing resembles the trend of business outsourcing because it provides the benefits of leveraging the expertise of others and being cost efficient. However, cloud computing also provides flexibility, scalability, elasticity, and reliability.

AMAZON AND CLOUD COMPUTING

AWS has been operating since 2006, and today serves hundreds of thousands of customers worldwide.

Using AWS, you can requisition compute power, storage, and other services in minutes and have the flexibility to choose the development platform or programming model that makes the most sense for the problems they're trying to solve. You pay only for what you use, with no up-front expenses or long-term commitments, making AWS a cost-effective way to deliver applications.

AWS is a public cloud. Cloud computing services also have several classifications:

■ ***Infrastructure as a service (IaaS)—***Offers fundamental resources like computing, storage, and networking capabilities, using virtual servers such as Amazon EC2, Google Compute Engine, and Microsoft Azure virtual machines

■ ***Platform as a service (PaaS)—***Provides platforms to deploy custom applications to the cloud, such as AWS Elastic Beanstalk, Google App Engine, and Heroku

■ ***Software as a service (SaaS)—***Combines infrastructure and software running in the cloud, including office applications like Amazon WorkSpaces, Google Apps for Work, and Microsoft Office 365

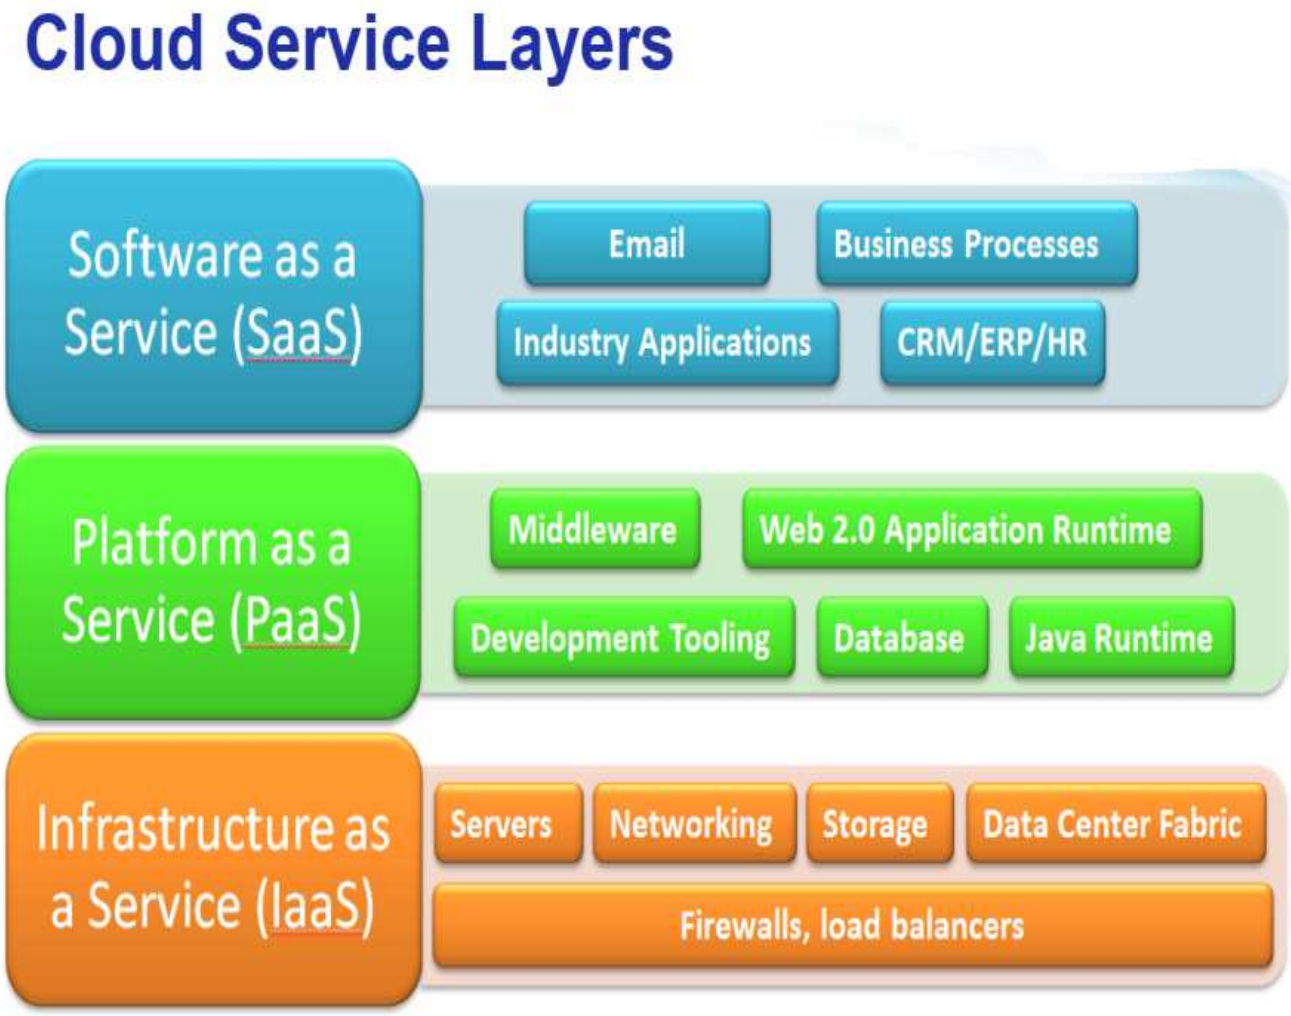

CLOUD SERVICE LAYERS DIAGRAM

Kvreddi, p14

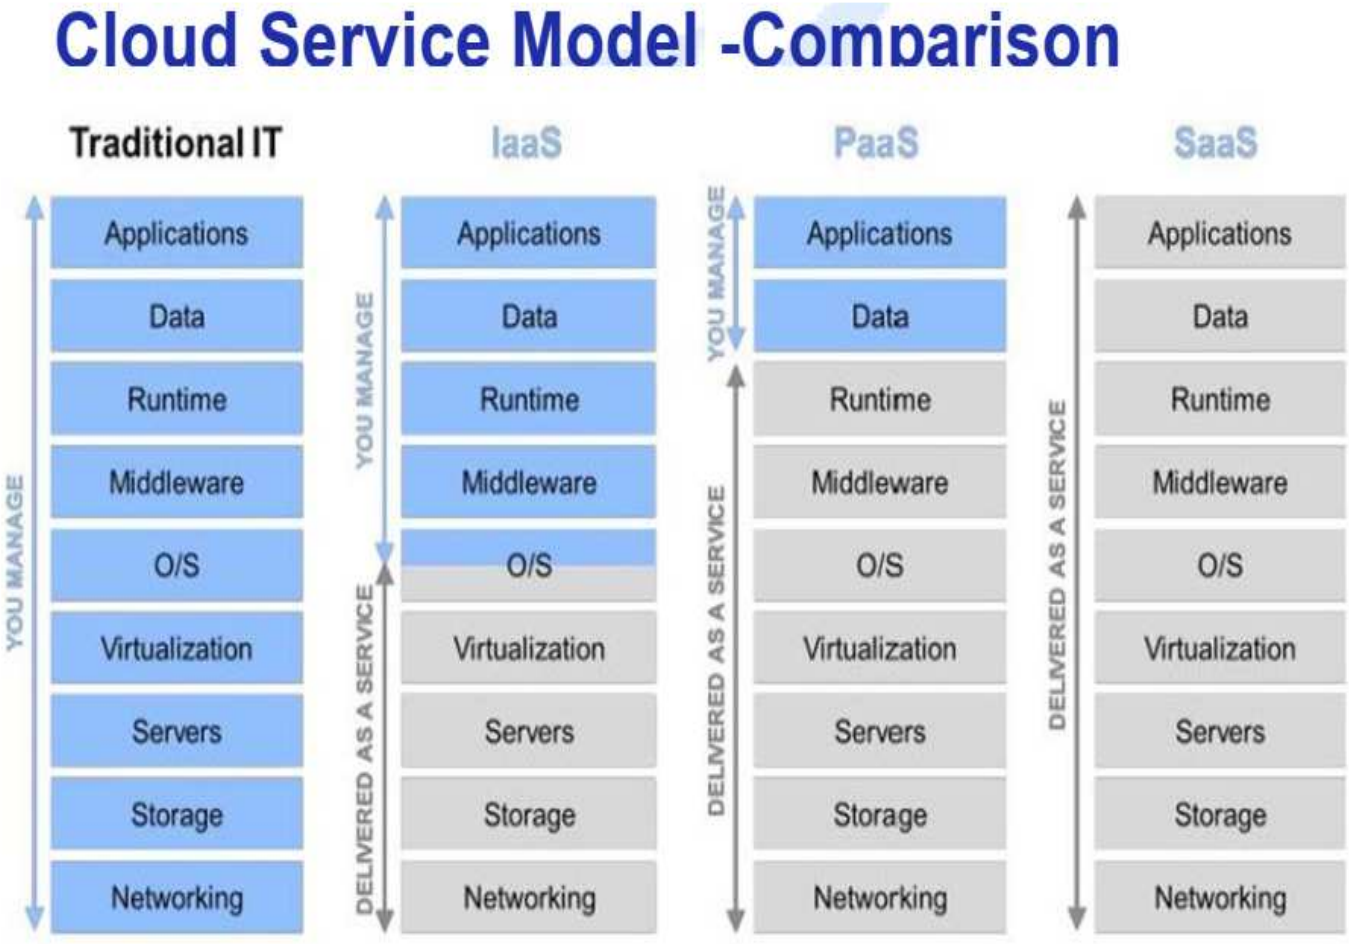

CLOUD SERVICE MODEL COMPARISON DIAGRAM

Kvreddi, p14

HOW YOU CAN BENEFIT FROM USING AWS

- Innovative and fast-growing platform
- Services solve common problems
- Enabling automation
- Flexible capacity (scalability)
- Built for failure (reliability)
- Reducing time to market
- Benefiting from economies of scale
- Its Worldwide
- One Gets Professional partners

THE DIFFERENCES THAT DISTINGUISH AWS

The Differences that Distinguish AWS

AWS is readily distinguished from other vendors in the traditional IT computing landscape because it is:

Flexible. AWS enables organizations to use the programming models, operating systems, databases, and architectures with which they are already familiar. In addition, this flexibility helps organizations mix and match architectures in order to serve their diverse business needs.

Cost-effective. With AWS, organizations pay only for what they use, without up-front or long-term commitments.

Scalable and elastic. Organizations can quickly add and subtract AWS resources to their applications in order to meet customer demand and manage costs.

Secure. In order to provide end-to-end security and end-to-end privacy, AWS builds services in accordance with security best practices, provides the appropriate security features in those services, and documents how to use those features.

Experienced. When using AWS, organizations can leverage Amazon's more than fifteen years of experience delivering large-scale, global infrastructure in a reliable, secure fashion.

Flexible

Using traditional models to deliver IT solutions often requires large investments in new architectures, programming languages, and operating systems. In contrast, the flexibility of AWS allows you to keep the programming models, languages, and operating systems that you are already using or choose others that are better suited for their project. Instead of re-writing applications, you can easily move them to the AWS cloud and tap into advanced computing

capabilities. AWS run almost anything—from full web applications to batch processing to offsite data back-ups.

AWS provides you flexibility when provisioning new services. Instead of the weeks and months it takes to plan, budget, procure, set up, deploy, operate, and hire for a new project, you can simply sign up for AWS and immediately begin deployment on the cloud the equivalent of 1, 10, 100, or 1,000 servers. Many customers find the flexibility of AWS to be a great asset in improving time to market and overall organizational productivity.

Cost-Effective

Developing and deploying an e-commerce application can be a low-cost effort, but a successful deployment can increase the need for hardware and bandwidth. Furthermore, owning and operating your own infrastructure can incur considerable costs, including power, cooling, real estate, and staff. In contrast, the cloud provides an on-demand IT infrastructure that lets you consume only the amount of resources that you actually need. You are not limited to a set amount of storage, bandwidth, or computing resources

You can get started through a completely self-service experience online, scale up and down as needed, and terminate your relationship with AWS at any time.

Scalable and Elastic

AWS uses the term elastic to describe the ability to scale computing resources up and down easily, with minimal friction. Elasticity helps you avoid provisioning resources up front for projects with variable consumption rates or short lifetimes. Instead of acquiring hardware, setting it up, and maintaining it in order to allocate resources to your applications, you use AWS to allocate resources using simple API calls.

Elastic Load Balancing and Auto Scaling can automatically scale your AWS cloud-based resources up to meet unexpected demand, and then scale those resources down as demand decreases. aws.amazon.com/architecture.

Secure

Ensuring the confidentiality, integrity, and availability of your data is of the utmost importance to AWS, as is maintaining your trust and confidence. AWS takes the following approaches to secure the cloud infrastructure:

Physical security. Amazon has many years of experience designing, constructing, and operating large-scale data centers. The AWS infrastructure is located in Amazon-controlled data centers throughout the world.

Secure services. Each service in the AWS cloud is architected to be secure. The services contain a number of capabilities that restrict unauthorized access or usage without sacrificing the flexibility that customers demand.

Data privacy. You can encrypt personal and business data in the AWS cloud, and publish backup and redundancy procedures for services so that your customers can protect their data and keep their applications running.
aws.amazon.com/security.

Experienced

The AWS cloud provides levels of scale, security, reliability, and privacy that are often cost-prohibitive for many organizations to meet or exceed. AWS has built an infrastructure based on lessons learned from over sixteen years' experience managing the multi-billion dollar Amazon.com business.

In addition to new services, AWS constantly hones its operational expertise to ensure ongoing dependability, and we continue to incorporate both industry best practices and proprietary advances into the cloud infrastructure.

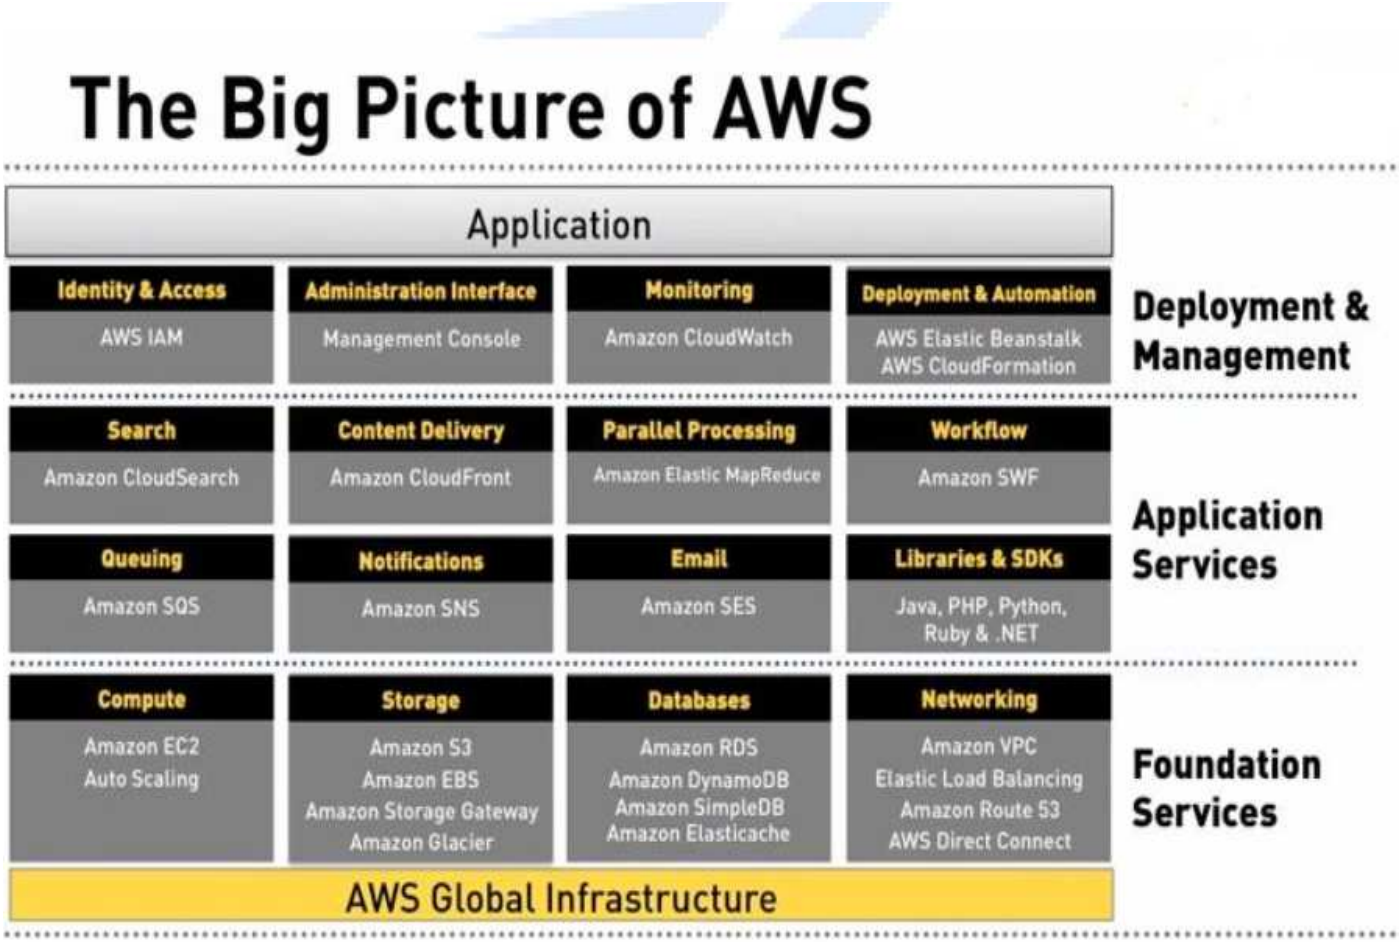

AWS GLOBAL INFRASTRUCTURE PICTURE

Kvreddi, p27

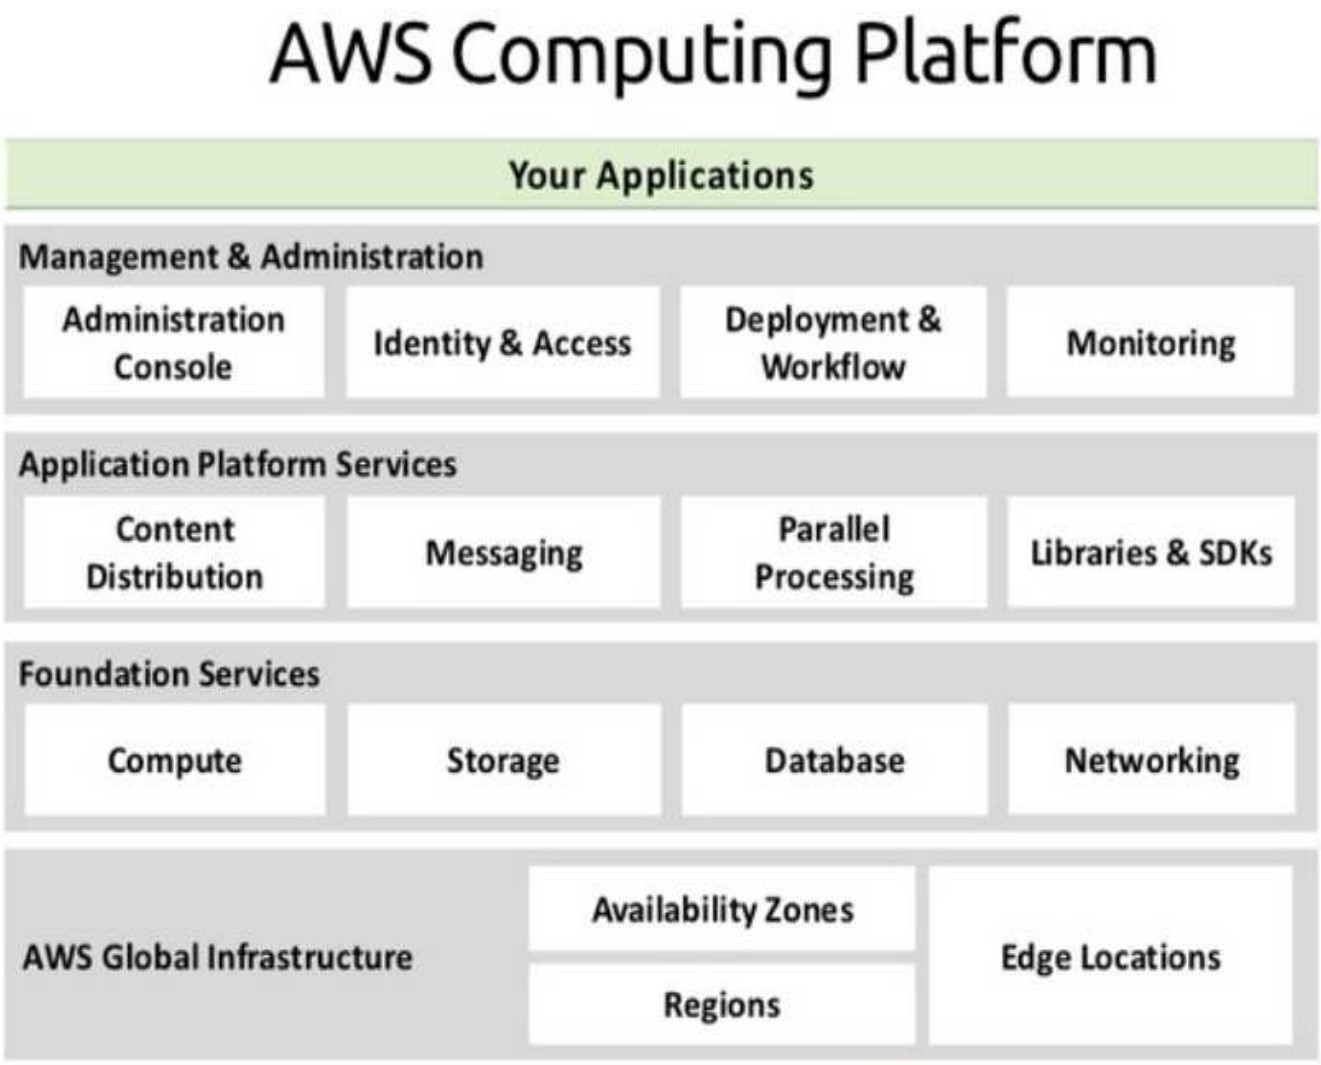

Kvreddi, p27

AWS COMPUTING PLATFORM DIAGRAM

GLOBAL INFRASTRUCTURE

Amazon Inc (2019) Amazon Web Services Whitepaper

The AWS Cloud infrastructure is built around AWS Regions and Availability Zones. An AWS Region is a physical location in the world where we have multiple Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities.

Each Amazon Region is designed to be completely isolated from the other Amazon Regions. This achieves the greatest possible fault tolerance and stability. Each Availability Zone is isolated, but the Availability Zones in a Region are connected through low-latency links. Each Availability Zone is designed as an independent failure zone. Availability Zones are all redundantly connected to multiple tier-1 transit providers.

SECURITY AND COMPLIANCE

SECURITY

The AWS Cloud enables a shared responsibility model. While AWS manages security **of** the cloud, you are responsible for security **in** the cloud. This means that you retain control of the security you choose to implement to protect your own content, platform, applications, systems, and networks no differently than you would in an on-site data center.

You get access to hundreds of tools and features to help you to meet your security objectives. AWS provides security-specific tools and features across network security, configuration management, access control, and data encryption. Finally, AWS environments are continuously audited, with certifications from accreditation bodies across geographies and verticals.

Benefits of AWS Security

• *Keep Your Data Safe:* The AWS infrastructure puts strong safeguards in place to help protect your privacy. All data is stored in highly secure AWS data centers.

• *Meet Compliance Requirements:* AWS manages dozens of compliance programs in its infrastructure. This means that segments of your compliance have already been completed.

• *Save Money:* Cut costs by using AWS data centers. Maintain the highest standard of security without having to manage your own facility

• *Scale Quickly:* Security scales with your AWS Cloud usage. No matter the size of your business, the AWS infrastructure is designed to keep your data safe.

COMPLIANCE

By tying together governance-focused, audit-friendly service features with applicable compliance or audit standards, AWS Compliance enablers build on traditional programs. This helps customers to establish and operate in an AWS security control environment.

The following is a partial list of assurance programs with which AWS complies:

- SOC 1/ISAE 3402, SOC 2, SOC 3
- FISMA, DIACAP, and FedRAMP
- PCI DSS Level 1
- ISO 9001, ISO 27001, ISO 27017, ISO 27018

AMAZON WEB SERVICES CLOUD PLATFORM

Topics

AWS MANAGEMENT CONSOLE

Access and manage Amazon Web Services through the AWS Management Console. Also use the AWS Console Mobile Application to quickly view resources on the go.

AWS COMMAND LINE INTERFACE

With just one tool to download and configure, you can control multiple AWS services from the command line and automate them through scripts.

SOFTWARE DEVELOPMENT KITS

Simplify using AWS services in your applications with an Application Program Interface (API) tailored to your programming language or platform.

ANALYTICS

Topics

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 to manage, and you pay only for the queries that you run.

Athena is out-of-the-box integrated with AWS Glue Data Catalog, allowing you to create a unified metadata repository across various services, crawl data sources to discover schemas and populate your Catalog with new and modified table and partition definitions, and maintain schema versioning.

Amazon EMR

Amazon EMR provides a managed Hadoop framework that makes it easy, fast, and cost-effective to process vast amounts of data across dynamically scalable Amazon EC2 instances. You can also run other popular distributed frameworks such as Apache Spark, HBase, Presto, and Flink in Amazon EMR, and interact with data in other AWS data stores such as Amazon S3 and Amazon DynamoDB.

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 CloudSearch

Amazon CloudSearch
supports 34 languages and popular search features such as highlighting, autocomplete, and geospatial search.

Amazon Elasticsearch Service

With Amazon Elasticsearch Service, you get easy-to-use APIs
and real-time analytics capabilities to power use-cases such as log analytics, full-text search, application monitoring, and clickstream analytics, with enterprise-grade availability, scalability, and security. The service offers integrations with open-source tools like Kibana and Logstash for data ingestion and visualization.

Amazon Kinesis

Amazon Kinesis makes it easy to collect, process, and analyze real-time, streaming data so you can get timely insights and react quickly to new information.

With Amazon Kinesis, you can ingest real-time data such as video, audio, application logs, website clickstreams, and IoT telemetry data for machine learning, analytics, and other applications.

Amazon Kinesis currently offers four services: Kinesis Data Firehose, Kinesis Data Analytics, Kinesis Data Streams, and Kinesis Video Streams.

Amazon Kinesis Data Firehose

It can capture, transform, and load streaming data into Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today.

Amazon Kinesis Data Analytics

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

Amazon Kinesis Data Streams

KDS 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 data collected 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.

It also durably stores, encrypts, and indexes video data in your streams, and allows you to access your data through easy-to-use APIs.

Amazon Redshift

Amazon Redshift is a fast, scalable data warehouse that makes it simple and cost-effective to analyze all your data across your data warehouse and data lake. Redshift delivers ten times faster performance than other data warehouses by using machine learning, massively parallel query execution, and columnar storage on high-performance disk.

Amazon QuickSight

Amazon QuickSight is a fast, cloud-powered business intelligence (BI) service that makes it easy for you to deliver insights to everyone in your organization. QuickSight lets you create and publish interactive dashboards that can be accessed from browsers or mobile devices.

AWS Data Pipeline

AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services, as well as on-premises data sources, at specified intervals.

AWS Glue

AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy for customers to prepare and load their data for analytics.

You simply point AWS Glue to your data stored on AWS, and AWS Glue discovers your data and stores the associated metadata (e.g. table definition and schema) in the AWS Glue Data Catalog. Once cataloged, your data is immediately searchable, queryable, and available for ETL.

AWS Lake Formation

AWS Lake Formation is a service that makes it easy to set up a secure data lake in days. A data lake is a centralized, curated, and secured repository that stores all your data, both in its original form and prepared for analysis. A data lake enables you to break down data silos and combine different types of analytics to gain insights and guide better business decisions.

Amazon Managed Streaming for Kafka (MSK)

build and run applications that use Apache Kafka to process streaming data. Apache Kafka is an opensource platform for building real-time streaming data pipelines and applications. With Amazon MSK, you can use Apache Kafka APIs to populate data lakes, stream changes to and from databases, and power machine learning and analytics applications.

Amazon Managed Streaming for Kafka makes it easy for you to build and run production applications on Apache Kafka without needing Apache Kafka infrastructure management expertise.

APPLICATION INTEGRATION

Topics

- AWS Step Functions
- Amazon MQ
- Amazon SQS
- Amazon SNS
- Amazon SWF

AWS Step Functions

AWS Step Functions lets you coordinate multiple AWS services into serverless workflows so you can build and update apps quickly. Using Step Functions, you can design and run workflows that stitch together services such as AWS Lambda and Amazon ECS into feature-rich applications

Amazon MQ

Amazon MQ is a managed message broker service for Apache ActiveMQ that makes it easy to set up and operate message brokers in the cloud. Message brokers allow different software systems–often using different programming languages, and on different platforms–to communicate and exchange information. Amazon MQ reduces your operational load by managing the provisioning, setup, and maintenance of ActiveMQ, a popular open-source message broker

Amazon SQS

Amazon Simple Queue Service (Amazon SQS) is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS eliminates the complexity and overhead associated with managing and operating message oriented middleware, and empowers developers to focus on differentiating work.

SQS offers two types of message queues. Standard queues offer maximum throughput, best-effort ordering, and at-least-once delivery. SQS FIFO queues are designed to guarantee that messages are processed exactly once, in the exact order that they are sent.

Amazon SNS

Amazon Simple Notification Service (Amazon SNS) is a highly available, durable, secure, fully managed pub/sub messaging service that enables you to decouple microservices, distributed systems, and serverless applications. Amazon SNS provides topics for high-throughput, push-based, many-to-many messaging.

Amazon SWF

Amazon Simple Workflow (Amazon SWF) helps developers build, run, and scale background jobs that have parallel or sequential steps.

AR AND VR

Topics

- Amazon Sumerian

Amazon Sumerian

Amazon Sumerian lets you create and run virtual reality (VR), augmented reality (AR), and 3D applications quickly and easily without requiring any specialized programming or 3D graphics expertise. With Sumerian, you can build highly immersive and interactive scenes that run on popular hardware such as Oculus Go, Oculus Rift, HTC Vive, HTC Vive Pro, Google Daydream, and Lenovo Mirage as well as Android and iOS mobile devices. For example, you can build a virtual classroom that lets you train new employees around the world, or you can build a virtual environment that enables people to tour a building remotely. Sumerian makes it easy to create all the building blocks needed to build highly immersive and interactive 3D experiences including adding objects (e.g. characters, furniture, and landscape), and designing, animating, and scripting environments. Sumerian does not require specialized expertise and you can design scenes directly from your browser

AWS COST MANAGEMENT

Topics

- AWS Cost Explorer
- AWS Budgets
- AWS Cost & Usage Report
- Reserved Instance (RI) Reporting

AWS Cost Explorer

Get started quickly by creating custom reports (including charts and

tabular data) that analyze cost and usage data, both at a high level (e.g., total costs and usage across all accounts) and for highly-specific requests (e.g., m2.2xlarge costs within account Y that are tagged “project: secretProject”).

AWS Budgets

AWS Budgets gives you the ability to set custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. You can also use AWS Budgets to set RI utilization or coverage targets and receive alerts when your utilization drops below the threshold you define. RI alerts support Amazon EC2, Amazon RDS, Amazon Redshift, and Amazon ElastiCache reservations.

AWS Cost & Usage Report

The AWS Cost & Usage Report lists AWS usage for each service category used by an account and its IAM users in hourly or daily line items, as well as any tags that you have activated for cost allocation purposes.

Reserved Instance (RI) Reporting

Using the RI Utilization and Coverage reports available in AWS Cost Explorer, you can visualize your RI data at an aggregate level or inspect a particular RI subscription.

The AWS Cost & Usage Report lists AWS usage for each service category used by an account and its IAM users in hourly or daily line items, as well as any tags that you have activated for cost allocation purposes.

BLOCKCHAIN

Amazon Managed Blockchain is a fully managed service that makes it easy to create and manage scalable blockchain networks using the popular open source frameworks Hyperledger Fabric and Ethereum.

Blockchain makes it possible to build applications where multiple parties can execute transactions without the need for a trusted, central authority. It manages your certificates, lets you easily invite new members to join the network, and tracks operational metrics such as usage of compute, memory, and storage resources. In addition, Managed Blockchain can replicate an immutable copy of your blockchain network activity into Amazon Quantum Ledger Database (QLDB), a fully managed ledger database. This allows you to easily analyze the network activity outside the network and gain insights into trends.

BUSINESS APPLICATIONS

Topics

- Alexa for Business
- Amazon WorkDocs
- Amazon WorkMail
- Amazon Chime

Alexa for Business

employees can use Alexa as their intelligent assistant to be more productive in meeting rooms, at their desks, and even with the Alexa devices they already have at home.

Amazon WorkDocs

Users can comment on files, send them to others for feedback, and upload new versions without having to resort to emailing multiple versions of their files as attachments.

Amazon WorkMail

Amazon WorkMail is a secure, managed business email and calendar service with support for existing desktop and mobile email client applications.

You can integrate Amazon WorkMail with your existing corporate directory, use email journaling to meet compliance requirements, and control both the keys that

Encrypt your data and the location in which your data is stored. You can also set up interoperability with Microsoft Exchange Server, and programmatically manage users, groups, and resources using the Amazon WorkMail SDK.

Amazon Chime

You can use Amazon Chime for online meetings, video conferencing, calls, chat, and to share content, both inside and outside your organization. Amazon Chime works with Alexa for Business, which means you can use Alexa to start your meetings with your voice.

COMPUTE SERVICES

Topics

- Amazon EC2
- Amazon EC2 Auto Scaling
- Amazon Elastic Container Registry
- Amazon Elastic Container Service
- Amazon Elastic Container Service for Kubernetes
- Amazon Lightsail
- AWS Batch
- AWS Elastic Beanstalk
- AWS Fargate
- AWS Lambda
- AWS Serverless Application Repository
- AWS Outposts
- VMware Cloud on AWS

Amazon EC2

The Amazon EC2 simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment.

Instance Types

• *On-Demand Instances*—With On-Demand instances, you pay for compute capacity by the hour with no long-term commitments. You can increase or decrease your compute capacity depending on the demands of your application and only pay the specified hourly rate for the instances you use. The use of On-Demand instances frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs. On-Demand instances also remove the need to buy "safety net" capacity to handle periodic traffic spikes.

• *Reserved Instances*—Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. You have the flexibility to change families, operating system types, and tenancies while benefitting from Reserved Instance pricing when you use Convertible Reserved Instances.

• *Spot Instances*—Spot Instances allow you to bid on spare Amazon EC2 computing capacity. Since Spot instances are often available at a discount compared to On-Demand pricing, you can significantly reduce the cost of running your applications, grow your application's compute capacity and throughput for the same budget, and enable new types of cloud computing applications.

Amazon EC2 Auto Scaling

Dynamic scaling responds to changing demand and predictive scaling automatically schedules the right number of EC2 instances based on predicted demand. Dynamic scaling and predictive scaling can be used together to scale faster.

Amazon Elastic Container Registry

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

Amazon ECR is integrated with Amazon Elastic Container Service (Amazon ECS), simplifying your development to production workflow.

Amazon Elastic Container Service

Amazon ECS eliminates the need for you to install and operate your own container orchestration software, manage and scale a cluster of virtual machines, or schedule containers on those virtual machines.
With simple API calls, you can launch and stop Docker-enabled applications, query the complete state of your application, and access many familiar features such as IAM roles, security groups, load balancers, Amazon CloudWatch Events, AWS CloudFormation templates, and AWS CloudTrail logs.

Amazon Elastic Container Service for Kubernetes

Amazon Elastic Container Service for Kubernetes (Amazon EKS) makes it easy to deploy, manage, and scale containerized applications using Kubernetes on AWS.

Amazon Lightsail

Lightsail plans include everything you need to jumpstart your project – a virtual machine, SSD based storage, data transfer, DNS management, and a static IP address – for a low, predictable price.

AWS Batch

AWS Batch enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of 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 Elastic Beanstalk

AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on familiar servers such as Apache, Nginx, Passenger, and Internet Information Services (IIS)

AWS Fargate

AWS Fargate is a compute engine for Amazon ECS that allows you to run containers without having to manage servers or clusters. With AWS Fargate, you no longer have to provision, configure, and scale clusters of virtual machines to run containers. This removes the need to choose server types, decide when to scale your clusters, or optimize cluster packing. AWS Fargate removes the need for you to interact with or think about servers or clusters.

Amazon ECS has two modes: Fargate launch type and EC2 launch type. With Fargate launch type, all you have to do is package your application in containers, specify the CPU and memory requirements, define networking and IAM policies, and launch the application. EC2 launch type allows you to have server-level, more granular control over the infrastructure that runs your container applications. With EC2 launch type, you can use Amazon ECS to manage a cluster of servers and schedule placement of containers on the servers. Amazon ECS keeps track of all the CPU, memory and other resources in your cluster, and also finds the best server for a container to run on based on your specified resource requirements. You are responsible for provisioning, patching, and scaling clusters of servers. You can decide which type of server to use, which applications and how many containers to run in a cluster to optimize utilization, and when you should add or remove servers from a cluster. EC2 launch type gives you more control of your server clusters and provides a broader range of customization options, which might be required to support some specific applications or possible compliance and government requirements.

AWS Lambda

AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume—there is no charge when your code is not running. With Lambda, you can run code for virtually any type of application or backend service—all with zero administration.

AWS Serverless Application Repository

The AWS Serverless Application Repository enables you to quickly deploy code samples, components, and complete applications for common use cases such as web and mobile back-ends, event and data processing, logging, monitoring, IoT, and more. Each application is packaged with an AWS Serverless Application Model (SAM) template that defines the AWS resources used. Publicly shared applications also include a link to the application's source code.

To share an application you've built, publish it to the AWS Serverless Application Repository

AWS Outposts

AWS Outposts bring native AWS services, infrastructure, and operating models to virtually any data center, co-location space, or on-premises facility. You can use the same APIs, the same tools, the same hardware, and the same functionality across on-premises and the cloud to deliver a truly consistent hybrid experience. Outposts can be used to support workloads that need to remain on-premises due to low latency or local data processing needs.

AWS Outposts come in two variants:

1) VMware Cloud on AWS Outposts allows you to use the same VMware control plane and APIs you use to run your infrastructure,

2) AWS native variant of AWS Outposts allows you to use the same exact APIs and control plane you use to run in the AWS cloud, but on-premises.

VMware Cloud on AWS

VMware Cloud on AWS is an integrated cloud offering jointly developed by AWS and VMware delivering a highly scalable, secure and innovative service that allows organizations to seamlessly migrate and extend

their on-premises VMware vSphere-based environments to the AWS Cloud running on next-generation Amazon Elastic Compute Cloud (Amazon EC2) bare metal infrastructure.

Availability in the following AWS Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Sydney), Asia Pacific (Tokyo), Europe (Frankfurt), Europe (Ireland), and Europe (London).

With VMware Cloud on AWS, organizations can simplify their Hybrid IT operations by using the same VMware Cloud Foundation technologies including vSphere, vSAN, NSX, and vCenter Server across their on-premises data centers and on the AWS Cloud without having to purchase any new or custom hardware, rewrite applications, or modify their operating models. The service automatically provisions infrastructure and provides full VM compatibility and workload portability between your on-premises environments and the AWS Cloud. With VMware Cloud on AWS, you can leverage AWS's breadth of services, including compute, databases, analytics, Internet of Things (IoT), security, mobile, deployment, application services, and more.

CUSTOMER ENGAGEMENT

Topics

- Amazon Connect
- Amazon SES

Amazon Connect

The self-service graphical interface in Amazon Connect makes it easy for non-technical users to design contact flows, manage agents, and track performance metrics – no specialized skills required. There are no up-front payments or long-term commitments and no infrastructure to manage with Amazon Connect; customers pay by the minute for Amazon Connect usage plus any associated telephony services

Amazon SES

Amazon Simple Email Service (Amazon SES) is a cloud-based email sending service designed to help digital marketers and application developers send marketing, notification, and transactional emails

See also Amazon Pinpoint

DATABASE

Topics

- Amazon Aurora
- Amazon Relational Database Service
- Amazon RDS on VMware
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon Neptune
- Amazon Quantum Ledger Database (QLDB)
- Amazon Timestream

Amazon Aurora

Amazon Aurora is a MySQL and PostgreSQL compatible relational database engine that combines the speed and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases.

Amazon Aurora is up to five times faster than standard MySQL databases and three times faster than standard PostgreSQL databases. It provides the security, availability, and reliability of commercial databases at 1/10th the cost. Amazon Aurora is fully managed by Amazon Relational Database Service (RDS), which automates time-consuming administration tasks like hardware provisioning, database setup, patching, and backups.

Amazon Aurora features a distributed, fault-tolerant, self-healing storage system that auto-scales up to 64TB per database instance. It delivers high performance and availability with up to 15 low-latency read replicas, point-in-time recovery, continuous backup to Amazon S3, and replication across three Availability Zones (AZs).

Amazon Relational Database Service

Amazon RDS is available on several database instance types - optimized for memory, performance or I/O - and provides you with six familiar database engines to choose from, including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. You can use the AWS Database Migration Service to easily migrate or replicate your existing databases to Amazon RDS

Amazon RDS on VMware

Amazon Relational Database Service (Amazon RDS) on VMware lets you deploy managed databases in on-premises VMware environments using the Amazon RDS technology enjoyed by hundreds of thousands of AWS customers.

You can easily replicate RDS on VMware
databases to RDS instances in AWS, enabling low-cost hybrid deployments for disaster recovery, read replica bursting, and optional long-term backup retention in Amazon Simple Storage Service (Amazon S3).

Amazon DynamoDB

Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. It's a fully managed, multiregion, multimaster database with built-in security, backup and restore, and in-memory caching for internet-scale applications. DynamoDB can handle more than 10 trillion requests per day and support peaks of more than 20 million requests per second.

Amazon ElastiCache

The service improves the performance of web applications by allowing you to retrieve
information from fast, managed, in-memory caches, instead of relying entirely on slower disk-based databases.

Amazon ElastiCache supports two open-source in-memory caching engines:

• *Redis* - a fast, open source, in-memory data store and cache. Amazon ElastiCache for Redis is a Rediscompatible in-memory service that delivers the ease-of-use and power of Redis along with the availability, reliability, and performance suitable for the most demanding applications. Both singlenode and up to 15-shard clusters are available, enabling scalability to up to 3.55 TiB of in-memory data. ElastiCache for Redis is fully managed, scalable, and secure. This makes it an ideal candidate to power high-performance use cases such as web, mobile apps, gaming, ad-tech, and IoT.

• *Memcached* - a widely adopted memory object caching system. ElastiCache for Memcached is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service.

Amazon Neptune

The core of Amazon Neptune is a purposebuilt, 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 you 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

Amazon Quantum Ledger Database (QLDB)

Amazon QLDB is a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log owned by a central trusted authority. Amazon QLDB tracks. Each and every application data change and maintains a complete and verifiable history of changes over time.

Ledgers are typically used to record a history of economic and financial activity in an organization. Many organizations build applications with ledger-like functionality because they want to maintain an accurate history of their applications' data, for example, tracking the history of credits and debits in banking transactions, verifying the data lineage of an insurance claim, or tracing movement of an item in a supply chain network. Ledger applications are often implemented using custom audit tables or audit trails created in relational databases. However, building audit functionality with relational databases is time consuming and prone to human error. It requires custom development, and since relational databases are not inherently immutable, any unintended changes to the data are hard to track and verify. Alternatively, blockchain frameworks, such as Hyperledger Fabric and Ethereum, can also be used as a ledger. However, this adds complexity as you need to set-up an entire blockchain network with multiple nodes, manage its infrastructure, and require the nodes to validate each transaction before it can be added to the ledger.

Amazon QLDB is a new class of database that eliminates the need to engage in the complex development effort of building your own ledger-like applications. With QLDB, your data's change history is immutable – it cannot be altered or deleted – and using cryptography, you can easily verify that there have been no unintended modifications to your application's data. QLDB uses an immutable transactional log, known as a journal, that tracks each application data change and maintains a complete and verifiable history of changes over time. QLDB is easy to use because it provides developers with a familiar SQL-like API, a flexible document data model, and full support for transactions. QLDB is also serverless, so it automatically scales to support the demands of your application. There are no servers to manage and no read or write limits to configure. With QLDB, you only pay for what you use.

Amazon Timestream

Amazon Timestream is a fast, scalable, fully managed time series database service for IoT and operational applications that makes it easy to store and analyze trillions of events per day at 1/10th the cost of relational databases. Driven by the rise of IoT devices, IT systems, and smart industrial machines, time-series data — data that measures how things change over time — is one of the fastest growing data types. Time-series data has specific characteristics such as typically arriving in time order form, data is append-only, and queries are always over a time interval. While relational databases can store this data, they are inefficient at processing this data as they lack optimizations such as storing and retrieving data by time intervals. Timestream is a purpose-built time series database that efficiently

stores and processes this data by time intervals. With Timestream, you can easily store and analyze log data for DevOps, sensor data for IoT applications, and industrial telemetry data for equipment maintenance. As your data grows over time, Timestream's adaptive query processing engine understands its location and format, making your data simpler and faster to analyze. Timestream also automates rollups, retention, tiering, and compression of data, so you can manage your data at the lowest possible cost. Timestream is serverless, so there are no servers to manage. It manages time-consuming tasks such as server provisioning, software patching, setup, configuration, or data retention and tiering, freeing you to focus on building your applications.

DESKTOP AND APP STREAMING

Topics

- Amazon WorkSpaces
- Amazon AppStream 2.0

Amazon WorkSpaces

Amazon WorkSpaces is a fully managed, secure cloud desktop service. You can use Amazon WorkSpaces to provision either Windows or Linux desktops in just a few minutes and quickly scale to provide thousands of desktops to workers across the globe. With Amazon WorkSpaces, your users get a fast, responsive desktop of their choice that they can access anywhere, anytime, from any supported device.

Amazon AppStream 2.0

Amazon AppStream 2.0 is a fully managed application streaming service. You centrally manage your desktop applications on AppStream 2.0 and securely deliver them to any computer. You can easily scale to any number of users across the globe without acquiring, provisioning, and operating hardware or infrastructure. AppStream 2.0 is built on AWS, so you benefit from a data center and network architecture designed for the most security-sensitive organizations. Each user has a fluid and responsive experience with your applications, including GPU-intensive 3D design and engineering ones, because your applications run on virtual machines (VMs) optimized for specific use cases and each streaming session automatically adjusts to network conditions.

Enterprises can use AppStream 2.0 to simplify application delivery and complete their migration to the cloud. Educational institutions can provide every student access to the applications they need for class on any computer. Software vendors can use AppStream 2.0 to deliver trials, demos, and training for their applications with no downloads or installations. They can also develop a full software-as-a-service (SaaS) solution without rewriting their application.

DEVELOPER TOOLS

Topics

- AWS CodeCommit
- CodeBuild
- CodeDeploy
- CodePipeline
- AWS CodeStar
- Amazon Corretto
- AWS Cloud9
- AWS X-Ray

AWS CodeCommit

AWS CodeCommit is a fully managed source control service that makes it easy for companies to host secure and highly scalable private Git repositories. AWS CodeCommit eliminates the need to operate your own source control system or worry about scaling its infrastructure. You can use AWS CodeCommit to securely store anything from source code to binaries, and it works seamlessly with your existing Git tools.

CodeBuild

CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages that are ready to deploy. CodeBuild scales continuously and processes multiple builds concurrently. You can get started quickly by using prepackaged build environments, or you can create custom build environments that use your own build tools.

CodeDeploy

CodeDeploy makes it easier for you to rapidly release new features,
helps you avoid downtime during application deployment, and handles the complexity of updating your applications.

CodePipeline

CodePipeline is a fully managed continuous delivery service that helps you automate your release pipelines for fast and reliable application and infrastructure updates. CodePipeline automates the build, test, and deploy phases of your release process every time there is a code change, based on the release model you define.

AWS CodeStar

AWS CodeStar enables you to quickly develop, build, and deploy applications on AWS. AWS CodeStar provides a unified user interface, enabling you to easily manage your software development activities in one place. With AWS CodeStar, you can set up your entire continuous delivery toolchain in minutes, allowing you to start releasing code faster. AWS CodeStar makes it easy for your whole team to work together securely, allowing you to easily manage access and add owners, contributors, and viewers to your projects

Amazon Corretto

Amazon Corretto is a no-cost, multiplatform, production-ready distribution of the Open Java Development Kit (OpenJDK). Corretto comes with long-term support that will include performance enhancements and security fixes.

AWS Cloud9

AWS Cloud9 is a cloud-based integrated development environment (IDE) that lets you write, run, and debug your code with just a browser. It includes a code editor, debugger, and terminal. Cloud9 comes prepackaged with essential tools for popular programming languages, including JavaScript, Python, PHP, and more, so you don't need to install files or configure your development machine to start new projects. Since your Cloud9 IDE is cloud-based, you can work on your projects from your office, home, or anywhere using an internet-connected machine. Cloud9 also provides a seamless experience for developing serverless applications enabling you to easily define resources, debug, and switch between local and remote execution of serverless applications.

AWS X-Ray

AWS X-Ray helps developers analyze and debug distributed applications in production or under development, such as those built using a microservices architecture. With X-Ray, you can understand how your application and its underlying services are performing so you can 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 your application, and shows a map of your application's underlying components. You can use XRay to analyze both applications in development and in production, from simple three-tier applications to complex microservices applications consisting of thousands of services.

GAME TECH

Topics

- Amazon GameLift
- Amazon Lumberyard

Amazon GameLift

Amazon GameLift is a managed service for deploying, operating, and scaling dedicated game servers for session-based multiplayer games

Amazon Lumberyard

Amazon Lumberyard is a free, cross-platform, 3D game engine for you to create the highest-quality games, connect your games to the vast compute and storage of the AWS Cloud, and engage fans on Twitch

INTERNET OF THINGS (IOT)

Topics

- AWS IoT Core
- Amazon FreeRTOS
- AWS IoT Greengrass
- AWS IoT 1-Click
- AWS IoT Analytics
- AWS IoT Button
- AWS IoT Device Defender
- AWS IoT Device Management
- AWS IoT Events
- AWS IoT SiteWise
- AWS IoT Things Graph
- AWS Partner Device Catalog

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 can support billions of devices and trillions of messages, and can process and route those messages to AWS endpoints and to other devices reliably and securely. With AWS IoT Core, your applications can keep track of and communicate with all your devices, all the time, even when they aren't connected.

Amazon FreeRTOS

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

A microcontroller (MCU) is a single chip containing a simple processor that can be found in many devices, including appliances, sensors, fitness trackers, industrial automation, and automobiles. Many of these small devices could benefit from connecting to the cloud or locally to other devices. For example, smart electricity meters need to connect to the cloud to report on usage, and building security systems need to communicate locally so that a door will unlock when you badge in. Microcontrollers have limited compute power and memory capacity and typically perform simple, functional tasks. Microcontrollers frequently run operating systems that do not have built-in functionality to connect to local networks or the cloud, making IoT applications a challenge. Amazon FreeRTOS helps solve this problem by providing both the core operating system (to run the edge device) as well as software libraries that make it easy to securely connect to the cloud (or other edge devices) so you can collect data from them for IoT applications and take action.

AWS IoT Greengrass

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.
With AWS IoT Greengrass, you can use familiar languages and programming models to create and test your device software in the cloud, and then deploy it to your devices. AWS IoT Greengrass can be programmed to filter device data and only transmit necessary information back to the cloud.

AWS IoT 1-Click

AWS IoT 1-Click is a service that enables simple devices to trigger AWS Lambda functions that can execute an action. AWS IoT 1-Click supported devices enable you to easily perform actions such as notifying technical support, tracking assets, and replenishing goods or services.

You can easily create device groups and associate them with a Lambda function that executes your desired action when triggered. You can also track device health and activity with the pre-built reports.

AWS IoT Analytics

AWS IoT Analytics is a fully-managed service that makes it easy to run and operationalize sophisticated analytics on massive volumes of IoT data without having to worry about the cost and complexity typically required to build an IoT analytics platform.

IoT data is highly unstructured which makes it difficult to analyze with traditional analytics and business intelligence tools that are designed to process structured data. IoT data comes from devices that often record fairly noisy processes (such as temperature, motion, or sound). The data from these devices can frequently have significant gaps, corrupted messages, and false readings that must be cleaned up before analysis can occur. Also, IoT data is often only meaningful in the context of additional, third party data inputs.

AWS IoT Analytics automates each of the difficult steps that are required to analyze data from IoT devices. AWS IoT Analytics filters, transforms, and enriches IoT data before storing it in a time-series data store for analysis. You can setup the service to collect only the data you need from your devices, apply mathematical transforms to process the data, and enrich the data with device-specific metadata such as device type and location before storing the processed data. Then, you can analyze your data by running ad hoc or scheduled queries using the built-in SQL query engine, or perform more complex analytics and machine learning inference. AWS IoT Analytics makes it easy to get started with machine learning by including pre-built models for common IoT use cases.

AWS IoT Button

The AWS IoT Button is a programmable button based on the Amazon Dash Button hardware. This simple Wi-Fi device is easy to configure, and it's designed for developers to get started with AWS IoT Core, AWS Lambda, Amazon DynamoDB, Amazon SNS, and many other Amazon Web Services without writing device-specific code.

You can code the button's logic in the cloud to configure button clicks to count or track items, call or alert someone, start or stop something, order services, or even provide feedback.

AWS IoT Device Defender

AWS IoT Device Defender is a fully managed service that helps you secure your fleet of IoT devices. AWS IoT Device Defender continuously audits your IoT configurations to make sure that they aren't deviating from security best practices. A configuration is a set of technical controls you set to help keep information secure when devices are communicating with each other and the cloud. AWS IoT Device Defender makes it easy to maintain and enforce IoT configurations, such as ensuring device identity, authenticating and authorizing devices, and encrypting device data.

AWS IoT Device Management

AWS IoT Device Management makes it easy to securely onboard, organize, monitor, and remotely manage IoT devices at scale. With AWS IoT Device Management, you can register your connected devices individually or in bulk, and easily manage permissions so that devices remain secure. You can also organize your devices, monitor and troubleshoot device functionality, query the state of any IoT device in your fleet, and send firmware updates over-the-air (OTA). AWS IoT Device Management is agnostic to device type and OS, so you can manage devices from constrained microcontrollers to connected cars all with the same service.

AWS IoT Events

AWS IoT Events is a fully managed IoT service that makes it easy to detect and respond to events from IoT sensors and applications. Events are patterns of data identifying more complicated circumstances than expected, such as changes in equipment when a belt is stuck or connected motion detectors using movement signals to activate lights and security cameras.

AWS IoT SiteWise

AWS IoT SiteWise is a managed service that makes it easy to collect and organize data from industrial equipment at scale. You can easily monitor equipment across your industrial facilities to identify waste, such as breakdown of equipment and processes, production inefficiencies, and defects in products.

You can use IoT SiteWise to monitor operations across facilities, quickly compute common industrial performance metrics, and build applications to analyze industrial equipment data, prevent costly equipment issues, and reduce production inefficiencies.

AWS IoT Things Graph

AWS IoT Things Graph is a service that makes it easy to visually connect different devices and web services to build IoT applications.

You can get started with AWS IoT Things Graph using these pre-built models for popular device types, such as switches and programmable logic controllers (PLCs), or create your own custom model using a GraphQL-based schema modeling language, and deploy your IoT application to AWS IoT Greengrassenabled devices such as cameras, cable set-top boxes, or robotic arms in just a few clicks. IoT Greengrass is software that provides local compute and secure cloud connectivity so devices can respond quickly to local events even without internet connectivity, and runs on a huge range of devices from a Raspberry Pi to a server-level appliance. IoT Things Graph applications run on IoT Greengrass-enabled devices.

AWS Partner Device Catalog

The AWS Partner Device Catalog helps you find devices and hardware to help you explore, build, and go to market with your IoT solutions. Search for and find hardware that works with AWS, including development kits and embedded systems to build new devices, as well as off-the-shelf-devices such as gateways, edge servers, sensors, and cameras for immediate IoT project integration.

MACHINE LEARNING

Topics

- Amazon SageMaker
- Amazon SageMaker Ground Truth
- Amazon Comprehend
- Amazon Lex
- Amazon Polly
- Amazon Rekognition
- Amazon Translate
- Amazon Transcribe
- Amazon Elastic Inference
- Amazon Forecast
- Amazon Textract
- Amazon Personalize
- Amazon Deep Learning AMIs
- AWS DeepLens
- AWS DeepRacer
- Apache MXNet on AWS
- TenserFlow on AWS
- AWS Inferentia

Amazon SageMaker

Amazon SageMaker is a fully-managed platform that enables developers and data scientists to quickly and easily build, train, and deploy machine learning models at any scale.

Machine learning often feels a lot harder than it should be to most developers because the process to build and train models, and then deploy them into production is too complicated and too slow.

Amazon SageMaker includes modules that can be used together or independently to build, train, and deploy your machine learning models.

Amazon SageMaker Ground Truth

Amazon SageMaker Ground Truth helps you build highly accurate training datasets for machine learning quickly. SageMaker Ground Truth offers easy access to public and private human labelers and provides them with built-in workflows and interfaces for common labeling tasks. Savings are achieved by using machine learning to automatically label data. Over time, SageMaker Ground Truth can label more and more data automatically and substantially speed up the creation of training datasets.

Amazon Comprehend

Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text. No machine learning experience required

Amazon Comprehend uses machine learning to help you uncover the insights and relationships in your unstructured data. 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. You can also use AutoML capabilities in Amazon Comprehend to build a custom set of entities or text classification models that are tailored uniquely to your organization's needs.

Amazon Comprehend Medical. The service can identify medical information, such as medical conditions, medications, dosages, strengths, and frequencies from a variety of sources like doctor's notes, clinical trial reports, and patient health records. Amazon Comprehend Medical also identifies the relationship among the extracted medication and test, treatment and procedure information for easier analysis.

Amazon Lex

Amazon Lex is a service for building conversational interfaces into any application using voice and text. 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 you to build applications with highly engaging user experiences and lifelike conversational interactions.

Amazon Polly

Polly is an Amazon artificial intelligence (AI) service that uses advanced deep learning technologies to synthesize speech that sounds like a human voice. Polly includes 47 lifelike voices spread across 24 languages, so you can select the ideal voice and build speech-enabled applications that work in many different countries.

Amazon Rekognition

Amazon Rekognition is a service that makes it easy to add image analysis to your applications. With Rekognition, you can detect objects, scenes, and faces in images. You can also search and compare faces. The Amazon Rekognition API enables you to quickly add sophisticated deep-learning-based visual search and image classification to your applications.
Amazon Rekognition is based on the same proven, highly scalable, deep learning technology developed by Amazon's computer vision scientists to analyze billions of images daily for Prime Photos. Amazon Rekognition uses deep neural network models to detect and label thousands of objects and scenes in your images, and we are continually adding new labels and facial recognition features to the service.

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 Transcribe

Amazon Transcribe is an automatic speech recognition (ASR) service that makes it easy for developers to add speech-to-text capability to their applications.

Amazon Transcribe can be used for lots of common applications, including the transcription of customer service calls and generating subtitles on audio and video content. The service can transcribe audio files stored in common formats, like WAV and MP3, with time stamps for every word so that you can easily locate the audio in the original source by searching for the text.

Amazon Elastic Inference

Amazon Elastic Inference allows you to attach low-cost GPU-powered acceleration to Amazon EC2 and Amazon SageMaker instances to reduce the cost of running deep learning inference by up to

75%. Amazon Elastic Inference supports TensorFlow, Apache MXNet, and ONNX models, with more frameworks coming soon.

In most deep learning applications, making predictions using a trained model—a process called inference—can drive as much as 90% of the compute costs of the application due to two factors. First, standalone GPU instances are designed for model training and are typically oversized for inference. While training jobs batch process hundreds of data samples in parallel, most inference happens on a single input in real time that consumes only a small amount of GPU compute. Even at peak load, a GPU's compute capacity may not be fully utilized, which is wasteful and costly. Second, different models need different amounts of GPU, CPU, and memory resources. Selecting a GPU instance type that is big enough to satisfy the requirements of the least used resource often results in under-utilization of the other resources and high costs.

Amazon Elastic Inference solves these problems by allowing you to attach just the right amount of GPU-powered inference acceleration to any EC2 or SageMaker instance type with no code changes.

Amazon Forecast

Amazon Forecast is a fully managed service that uses machine learning to deliver highly accurate forecasts.
Companies today use everything from simple spreadsheets to complex financial planning software to attempt to accurately forecast future business outcomes such as product demand, resource needs, or financial performance. These tools build forecasts by looking at a historical series of data, which is called time series data.

Amazon Forecast requires no machine learning experience to get started. You only need to provide historical data, plus any additional data that you believe may impact your forecasts.

Amazon Textract

Amazon Textract is a service that automatically extracts text and data from scanned documents. Amazon Textract goes beyond simple optical character recognition (OCR) to also identify the contents of fields in forms and information stored in tables.

With Textract you can quickly automate document workflows, enabling you to process millions of document pages in hours. Once the information is captured, you can take action on it within your business applications to initiate next steps for a loan application or medical claims processing. Additionally, you can create smart search indexes, build automated approval workflows, and better maintain compliance with document archival rules by flagging data that may require redaction.

Amazon Personalize

Amazon Personalize is a machine learning service that makes it easy for developers to create individualized recommendations for customers using their applications.
Machine learning is being increasingly used to improve customer engagement by powering personalized product and content recommendations, tailored search results, and targeted marketing promotions.

Amazon Deep Learning AMIs

You can quickly launch
Amazon EC2 instances pre-installed with popular deep learning frameworks such as Apache MXNet and Gluon, TensorFlow, Microsoft Cognitive Toolkit, Caffe, Caffe2, Theano, Torch, PyTorch, Chainer, and Keras to train sophisticated, custom AI models, experiment with new algorithms, or to learn new skills and techniques.

AWS DeepLens

AWS DeepLens helps put deep learning in the hands of developers, literally, with a fully programmable video camera, tutorials, code, and pre-trained models designed to expand deep learning skills.

AWS DeepRacer

AWS DeepRacer is a 1/18th scale race car which gives you an interesting and fun way to get started with reinforcement learning (RL). RL is an advanced machine learning (ML) technique which takes a very different approach to training models than other machine learning methods. Its super power is that it learns very complex behaviors without requiring any labeled training data, and can make short term decisions while optimizing for a longer term goal.

Apache MXNet on AWS

Apache MXNet on AWS is a fast and scalable training and inference framework with an easy-to-use, concise API for machine learning.

MXNet includes the Gluon interface that allows developers of all skill levels to get started with deep learning on the cloud, on edge devices, and on mobile apps. In just a few lines of Gluon code, you can build linear regression, convolutional networks and recurrent LSTMs for object detection, speech recognition, recommendation, and personalization.

You can get started with MxNet on AWS with a fully-managed experience using Amazon SageMaker, a platform to build, train, and deploy machine learning models at scale. Or, you can use the AWS Deep Learning AMIs to build custom environments and workflows with MxNet as well as other frameworks including TensorFlow, PyTorch, Chainer, Keras, Caffe, Caffe2, and Microsoft Cognitive Toolkit.

TenserFlow on AWS

TenserFlow™ enables developers to quickly and easily get started with deep learning in the cloud.

The framework has broad support in the industry and has become a popular choice for deep learning research and application development, particularly in areas such as computer vision, natural language understanding and speech translation.

AWS Inferentia

AWS Inferentia is a machine learning inference chip designed to deliver high performance at low cost.

AWS Inferentia will support the TensorFlow, Apache MXNet, and PyTorch deep learning frameworks, as well as models that use the ONNX format.

AWS Inferentia provides high throughput, low latency inference performance at an extremely low cost. Each chip provides hundreds of TOPS (tera operations per second) of inference throughput to allow complex models to make fast predictions. For even more performance, multiple AWS Inferentia chips can be used together to drive thousands of TOPS of throughput.

MANAGEMENT AND GOVERNANCE

Topics

- Amazon CloudWatch
- AWS Auto Scaling
- AWS Control Tower
- AWS Systems Manager
- AWS CloudFormation
- AWS CloudTrail
- AWS Config
- AWS OpsWorks

- AWS Service Catalog
- AWS Trusted Advisor
- AWS Personal Health Dashboard
- AWS Managed Services
- AWS Console Mobile Application
- AWS License Manager
- AWS Well-Architected Tool

Amazon CloudWatch

Amazon CloudWatch is a monitoring and management service built for developers, system operators, site reliability engineers (SRE), and IT managers. CloudWatch provides you with data and actionable insights to monitor your 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 you with a unified view of AWS resources, applications and services that run on AWS, and on-premises servers. You can use CloudWatch to set high resolution alarms, visualize logs and metrics side by side, take automated actions, troubleshoot issues, and discover insights to optimize your applications, and ensure they are running smoothly.

AWS Auto Scaling

AWS Auto Scaling monitors your applications and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost. Using AWS Auto Scaling, it's easy to setup application scaling for multiple resources across multiple services in minutes. AWS Auto Scaling makes scaling simple with recommendations that allow you to optimize performance, costs, or balance between them.

AWS Control Tower

AWS Control Tower automates the set-up of a baseline environment, or landing zone, that is a secure, well-architected multi-account AWS environment. The configuration of the landing zone is based on best practices that have been established by working with thousands of enterprise customers to create a secure environment that makes it easier to govern AWS workloads with rules for security, operations, and compliance.

AWS Systems Manager

AWS Systems Manager gives you visibility and control of your infrastructure on AWS. Systems Manager provides a unified user interface so you can view operational data from multiple AWS services and allows you to automate operational tasks across your AWS resources. With Systems Manager, you 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 your groups of resources. Systems Manager simplifies resource and application management, shortens the time to detect and resolve operational problems, and makes it easy to operate and manage your infrastructure securely at scale.

AWS Systems Manager contains the following tools:

• *Resource groups:* Lets you create a logical group of resources associated with a particular workload such as different layers of an application stack, or production versus development environments. For example, you can group different layers of an application, such as the frontend web layer and the backend data layer. Resource groups can be created, updated, or removed programmatically through the API.

• *Insights Dashboard:* Displays operational data that the AWS Systems Manager automatically aggregates for each resource group. Systems Manager eliminates the need for you to navigate across multiple AWS consoles to view your operational data. With Systems Manager you can view API call logs from AWS CloudTrail, resource configuration changes from AWS Config, software inventory, and patch compliance status by resource group. You can also easily integrate your Amazon CloudWatch Dashboards, AWS Trusted Advisor notifications, and AWS Personal Health Dashboard performance and availability alerts into your Systems Manager dashboard. Systems Manager centralizes all relevant operational data, so you can have a clear view of your infrastructure compliance and performance.

• *Run Command:* Provides a simple way of automating common administrative tasks like remotely executing shell scripts or PowerShell commands, installing software updates, or making changes to the configuration of OS, software, EC2 and instances and servers in your on-premises data center.

• *State Manager:* Helps you define and maintain consistent OS configurations such as firewall settings and anti-malware definitions to comply with your policies. You can monitor the configuration of a large set of instances, specify a configuration policy for the instances, and automatically apply updates or configuration changes.

• *Inventory:* Helps you collect and query configuration and inventory information about your instances and the software installed on them. You can gather details about your instances such as installed applications, DHCP settings, agent detail, and custom items. You can run queries to track and audit your system configurations.

• *Maintenance Window:* Lets you define a recurring window of time to run administrative and maintenance tasks across your instances. This ensures that installing patches and updates, or making other configuration changes does not disrupt business-critical operations. This helps improve your application availability.

• *Patch Manager:* Helps you select and deploy operating system and software patches automatically across large groups of instances. You can define a maintenance window so that patches are applied only during set times that fit your needs. These capabilities help ensure that your software is always up to date and meets your compliance policies.

• *Automation:* Simplifies common maintenance and deployment tasks, such as updating Amazon Machine Images (AMIs). Use the Automation feature to apply patches, update drivers and agents, or bake applications into your AMI using a streamlined, repeatable, and auditable process.

• *Parameter Store:* Provides an encrypted location to store important administrative information such as passwords and database strings. The Parameter Store integrates with AWS KMS to make it easy to encrypt the information you keep in the Parameter Store.

• *Distributor:* Helps you securely distribute and install software packages, such as software agents. Systems Manager Distributor allows you to centrally store and systematically distribute software packages while you maintain control over versioning. You can use Distributor to create and distribute software packages and then install them using Systems Manager Run Command and State Manager. Distributor can also use AWS Identity and Access Management (IAM) policies to control who can create or update packages in your account. You can use the existing IAM policy support for Systems Manager Run Command and State Manager to define who can install packages on your hosts.

• *Session Manager:* Provides a browser-based interactive shell and CLI for managing Windows and Linux EC2 instances, without the need to open inbound ports, manage SSH keys, or use bastion hosts.

Administrators can grant and revoke access to instances through a central location by using AWS Identity and Access Management (IAM) policies. This allows you to control which users can access each instance, including the option to provide non-root access to specified users. Once access is provided, you can audit which user accessed an instance and log each command to Amazon S3 or Amazon CloudWatch Logs using AWS CloudTrail.

AWS CloudFormation

AWS CloudFormation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

You can use the AWS CloudFormation sample templates or create your own templates to describe your AWS resources, and any associated dependencies or runtime parameters, required to run your application. Visualize your templates as diagrams and edit them using a drag-and-drop interface with the AWS CloudFormation Designer.

AWS CloudTrail

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

AWS Config

AWS Config is a fully managed service that provides you with an AWS resource inventory, configuration history, and configuration change notifications to enable security and governance. The Config Rules feature enables you to create rules that automatically check the configuration of AWS resources recorded by AWS Config.

With AWS Config, you can discover existing and deleted AWS resources, determine your overall compliance against rules, and dive into configuration details of a resource at any point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting

AWS OpsWorks

AWS OpsWorks is a configuration management service that provides managed instances of Chef and Puppet. Chef and Puppet are automation platforms that allow you to use code to automate the configurations of your servers. OpsWorks lets you use Chef and Puppet to automate how servers are configured, deployed, and managed across your Amazon EC2 instances or on-premises compute environments. OpsWorks has three offerings, AWS OpsWorks for Chef Automate, AWS OpsWorks for Puppet Enterprise, and AWS OpsWorks Stacks.

AWS Service Catalog

AWS Service Catalog allows organizations 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 you to centrally manage commonly deployed IT services and helps you achieve consistent governance and meet your compliance requirements, while enabling users to quickly deploy only the approved IT services they need.

AWS Trusted Advisor

AWS Trusted Advisor is an online resource to help you reduce cost, increase performance, and improve security by optimizing your AWS environment.

AWS Personal Health Dashboard

AWS Personal Health Dashboard provides alerts and remediation guidance when AWS is experiencing events that might affect you. While the Service Health Dashboard displays the general status of AWS services, Personal Health Dashboard gives you a personalized view into the performance and availability of the AWS services underlying your AWS resources.

AWS Managed Services

AWS Managed Services provides ongoing management of your AWS infrastructure so you can focus on your applications. By implementing best practices to maintain your infrastructure, AWS Managed Services helps to reduce your operational overhead and risk. 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 your infrastructure.

AWS Console Mobile Application

The AWS Console Mobile Application lets customers view and manage a select set of resources to support incident response while on-the-go.

The Console Mobile Application allows AWS customers to monitor resources through a dedicated dashboard and view configuration details, metrics, and alarms for select AWS services. The Dashboard provides permitted users with a single view a resource's status, with real-time data on Amazon CloudWatch, Personal Health Dashboard, and AWS Billing and Cost Management.

AWS License Manager

AWS License Manager makes it easier to manage licenses in AWS and on-premises servers from software vendors such as Microsoft, SAP, Oracle, and IBM. AWS License Manager lets administrators create customized licensing rules that emulate the terms of their licensing agreements, and then enforces these rules when an instance of EC2 gets launched. 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 enable you to limit a licensing breach by physically stopping the instance from launching or by notifying administrators about the infringement.

AWS Well-Architected Tool

The AWS Well-Architected Tool helps you review the state of your workloads and compares them to the latest AWS architectural best practices. The tool is based on the AWS Well-Architected Framework, developed to help cloud architects build secure, high-performing, resilient, and efficient application infrastructure.

MEDIA SERVICES

Topics

- Amazon Elastic Transcoder
- AWS Elemental MediaConnect
- AWS Elemental MediaConvert
- AWS Elemental MediaLive
- AWS Elemental MediaPackage
- AWS Elemental MediaStore
- AWS Elemental MediaTailor

Amazon Elastic Transcoder

It is designed to be a highly scalable, easyto-use, and cost-effective way for developers and businesses to convert (or transcode) media files from their source format into versions that will play back on devices like smartphones, tablets, and PCs.

AWS Elemental MediaConnect

AWS Elemental MediaConnect is a high-quality transport service for live video. Today, broadcasters and content owners rely on satellite networks or fiber connections to send their high-value content into the cloud or to transmit it to partners for distribution.

Now you can get the reliability and security of satellite and fiber combined with the flexibility, agility, and economics of IP-based networks using AWS Elemental MediaConnect. MediaConnect enables you to build mission-critical live video workflows in a fraction of the time and cost of satellite or fiber services

AWS Elemental MediaConvert

AWS Elemental MediaConvert is a file-based video transcoding service with broadcast-grade features. It allows you to easily 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 and pay-as-you-go pricing.

AWS Elemental MediaLive

AWS Elemental MediaLive is a broadcast-grade live video processing service. It lets you create highquality 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 your live video streams in real-time, taking a larger-sized live video source and compressing it into smaller versions for distribution to your viewers.

AWS Elemental MediaPackage

AWS Elemental MediaPackage reliably prepares and protects your video for delivery over the Internet. From a single video input, AWS Elemental MediaPackage creates video streams formatted to play on connected TVs, mobile phones, computers, tablets, and game consoles.

AWS Elemental MediaStore

It gives you the performance, consistency, and low latency required to deliver live streaming video content. AWS Elemental MediaStore acts as the origin store in your video workflow. Its high performance capabilities meet the needs of the most demanding media delivery workloads, combined with long-term, cost-effective storage.

AWS Elemental MediaTailor

AWS Elemental MediaTailor lets video providers insert individually targeted advertising into their video streams without sacrificing broadcast-level quality-of-service. With AWS Elemental MediaTailor, viewers of your live or on-demand video each receive a stream that combines your content with ads personalized to them.

It also improves ad delivery rates, helping you make more money from every video, and it works with a wider variety of content delivery networks, ad decision servers, and client devices.

MIGRATION AND TRANSFER

Topics

- AWS Migration Hub
- AWS Application Discovery Service

- AWS Database Migration Service
- AWS Server Migration Service
- AWS Snowball
- AWS Snowball Edge
- AWS Snowmobile
- AWS DataSync
- AWS Transfer for SFTP

AWS Migration Hub

AWS Migration Hub provides a single location to track the progress of application migrations across multiple AWS and partner solutions. Using Migration Hub allows you to choose the AWS and partner migration tools that best fit your needs, while providing visibility into the status of migrations across your portfolio of applications. Migration Hub also provides key metrics and progress for individual applications, regardless of which tools are being used to migrate them. For example, you might use AWS Database Migration Service, AWS Server Migration Service, and partner migration tools such as ATADATA ATAmotion, CloudEndure Live Migration, or RiverMeadow Server Migration Saas to migrate an application comprised of a database, virtualized web servers, and a bare metal server. Using Migration Hub, you can view the migration progress of all the resources in the application. This allows you to quickly get progress updates across all of your migrations, easily identify and troubleshoot any issues, and reduce the overall time and effort spent on your migration projects.

AWS Application Discovery Service

AWS Application Discovery Service helps enterprise customers plan migration projects by gathering information about their on-premises data centers.

AWS Application Discovery Service collects and presents configuration, usage, and behavior data from your servers to help you better understand your workloads. You can export this data as a CSV file and use it to estimate the Total Cost of Ownership (TCO) of running on AWS and to plan your migration to AWS. In addition, this data is also available in AWS Migration Hub, where you can migrate the discovered servers and track their progress as they get migrated to AWS.

AWS Database Migration Service

AWS Database Migration Service helps you migrate databases to AWS easily and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. The AWS Database Migration Service can migrate your data to and from most widely used commercial and open-source databases. The service supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora or Microsoft SQL Server to MySQL. It also allows you to stream data to Amazon Redshift from any of the supported sources including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle, SAP ASE, and SQL Server, enabling consolidation and easy analysis of data in the petabyte-scale data warehouse. AWS Database Migration Service can also be used for continuous data replication with high availability.

AWS Server Migration Service

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

AWS Snowball

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

Snowball uses multiple layers of security designed to protect your data including tamper-resistant enclosures, 256-bit encryption, and an industry-standard Trusted Platform Module (TPM) designed to ensure both security and full chain of custody of your data. Once the data transfer job has been processed and verified, AWS performs a software erasure of the Snowball appliance.

AWS Snowball Edge

AWS Snowball Edge is a data migration and edge computing device that comes in two options. Snowball Edge Storage Optimized provides 100 TB of capacity and 24 vCPUs and is well suited for local storage and large scale data transfer. Snowball Edge Compute Optimized provides 52 vCPUs and an optional GPU for use cases such as advanced machine learning and full motion video analysis in disconnected environments.

Common use cases include data migration, data transport, image collation, IoT sensor stream capture, and machine learning

AWS Snowmobile

AWS Snowmobile is an exabyte-scale data transfer service used to move extremely large amounts of data to AWS. You can transfer up to 100 PB per Snowmobile, a 45-foot long ruggedized shipping container, pulled by a semi-trailer truck.

AWS DataSync

AWS DataSync is a data transfer service that makes it easy for you to automate moving data between on-premises storage and Amazon S3 or Amazon Elastic File System (Amazon EFS). DataSync automatically handles many of the tasks related to data transfers that can slow down migrations or burden your IT operations, including running your own instances, handling encryption, managing scripts, network optimization, and data integrity validation. You can use DataSync to transfer data at speeds up to 10 times faster than open-source tools.

AWS Transfer for SFTP

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

MOBILE SERVICES

Topics

- AWS Amplify
- Amazon Cognito
- Amazon Pinpoint
- AWS Device Farm
- AWS AppSync

AWS Amplify

AWS Amplify makes it easy to create, configure, and implement scalable mobile applications powered by AWS. Amplify seamlessly provisions and manages your mobile backend and provides a simple framework

to easily integrate your backend with your iOS, Android, Web, and React Native frontends. Amplify also automates the application release process of both your frontend and backend allowing you to deliver features faster.

Amazon Cognito

Amazon Cognito lets you add user sign-up, sign-in, and access control to your web and mobile apps quickly and easily. With Amazon Cognito, you also have the option to authenticate users through social identity providers such as Facebook, Twitter, or Amazon, with SAML identity solutions, or by using your own identity system. In addition, Amazon Cognito enables you to save data locally on users' devices, allowing your applications to work even when the devices are offline. You can then synchronize data across users' devices so that their app experience remains consistent regardless of the device they use. With Amazon Cognito, you can focus on creating great app experiences instead of worrying about building, securing, and scaling a solution to handle user management, authentication, and sync across devices.

Amazon Pinpoint

Amazon Pinpoint makes it easy to send targeted messages to your customers through multiple engagement channels. Examples of targeted campaigns are promotional alerts and customer retention campaigns, and transactional messages are messages such as order confirmations and password reset messages.

AWS Device Farm

AWS Device Farm is an app testing service that lets you test and interact with your Android, iOS, and web apps on many devices at once, or reproduce issues on a device in real time. View video, screenshots, logs, and performance data to pinpoint and fix issues before shipping your app.

AWS AppSync

AWS AppSync is a serverless back-end for mobile, web, and enterprise applications.
AWS AppSync makes it easy to build data driven mobile and web applications by handling securely all the application data management tasks like online and offline data access, data synchronization, and data manipulation across multiple data sources. AWS AppSync uses GraphQL, an API query language designed to build client applications by providing an intuitive and flexible syntax for describing their data requirement.

NETWORKING AND CONTENT DELIVERY

Topics

- Amazon VPC
- Amazon CloudFront
- Amazon Route 53
- AWS PrivateLink
- AWS Direct Connect
- AWS Global Accelerator
- Amazon API Gateway
- AWS Transit Gateway
- AWS App Mesh
- AWS Cloud Map
- Elastic Load Balancing

Amazon VPC

Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range,

creation of subnets, and configuration of route tables and network gateways. You can use both IPv4 and IPv6 in your VPC for secure and easy access to resources and applications.

Amazon CloudFront

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. CloudFront is integrated with AWS – both physical locations that are directly connected to the AWS global infrastructure, as well as other AWS services. CloudFront works seamlessly with services including AWS Shield for DDoS mitigation, Amazon S3, Elastic Load Balancing or Amazon EC2 as origins for your applications, and Lambda@Edge to run custom code closer to customers' users and to customize the user experience.

Amazon Route 53

Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end

Use Amazon Route 53 to configure DNS health checks to route traffic to healthy endpoints or to independently monitor the health of your application and its endpoints. Amazon Route 53 traffic flow makes it easy for you to manage traffic globally through a variety of routing types, including latency-based routing, Geo DNS, and weighted round robin—all of which can be combined with DNS Failover in order to enable a variety of low-latency, fault-tolerant architectures.

AWS PrivateLink

AWS PrivateLink simplifies the security of data shared with cloud-based applications by eliminating the exposure of data to the public Internet. AWS PrivateLink provides private connectivity between VPCs, AWS services, and on-premises applications, securely on the Amazon network.

AWS Direct Connect

AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data center, office, or co-location environment,

Using industry standard 802.1Q virtual LANS (VLANs), this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources, such as objects stored in Amazon S3 using public IP address space, and private resources such as EC2 instances running within a VPC using private IP address space, while maintaining network separation between the public and private environments.

AWS Global Accelerator

AWS Global Accelerator is a networking service that improves the availability and performance of the applications that you offer to your global users.

Amazon API Gateway

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. With a few clicks in the AWS Management Console, you can create an API that acts as a "front door" for applications to access data, business logic, or functionality from your back-end services, such as workloads running on Amazon EC2, code running on AWS Lambda, or any web application. Amazon API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management.

AWS Transit Gateway

AWS Transit Gateway is a service that enables customers to connect their Amazon Virtual Private Clouds (VPCs) and their on-premises networks to a single gateway. As you grow the number of workloads

running on AWS, you need to be able to scale your networks across multiple accounts and Amazon VPCs to keep up with the growth.

AWS App Mesh

AWS App Mesh makes it easy to monitor and control microservices running on AWS. App Mesh standardizes how your microservices communicate, giving you end-to-end visibility and helping to ensure high-availability for your applications.

AWS App Mesh makes it easy to run microservices by providing consistent visibility and network traffic controls for every microservice in an application. App Mesh removes the need to update application code to change how monitoring data is collected or traffic is routed between microservices.

AWS Cloud Map

AWS Cloud Map is a cloud resource discovery service. With Cloud Map, you can define custom names for your application resources, and it maintains the updated location of these dynamically changing resources. This increases your application availability because your web service always discovers the most up-to-date locations of its resources.

Cloud Map allows you to register any application resources such as databases, queues, microservices, and other cloud resources with custom names. The application can then query the registry for the location of the resources needed based on the application version and deployment environment.

Elastic Load Balancing

Elastic Load Balancing (ELB) automatically distributes incoming application traffic across multiple targets, such as Amazon EC2 instances, containers, and IP addresses. It can handle the varying load of your application traffic in a single Availability Zone or across multiple Availability Zones. Elastic Load Balancing offers three types of load balancers that all feature the high availability, automatic scaling, and robust security necessary to make your applications fault tolerant.

• *Application Load Balancer* is best suited for load balancing of HTTP and HTTPS traffic and provides advanced request routing targeted at the delivery of modern application architectures, including microservices and containers. Operating at the individual request level (Layer 7), Application Load Balancer routes traffic to targets within Amazon Virtual Private Cloud (Amazon VPC) based on the content of the request.

• *Network Load Balancer* is best suited for load balancing of TCP traffic where extreme performance is required. Operating at the connection level (Layer 4), Network Load Balancer routes traffic to targets within Amazon Virtual Private Cloud (Amazon VPC) and is capable of handling millions of requests per second while maintaining ultra-low latencies. Network Load Balancer is also optimized to handle sudden and volatile traffic patterns.

• *Classic Load Balancer* provides basic load balancing across multiple Amazon EC2 instances and operates at both the request level and connection level. Classic Load Balancer is intended for applications that were built within the EC2-Classic network.

ROBOTICS

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.

SATELLITE

AWS Ground Station

AWS Ground Station is a fully managed service that lets you control satellite communications, downlink and process satellite data, and scale your satellite operations quickly, easily and cost-effectively without having to worry about building or managing your own ground station infrastructure. Satellites are used for a wide variety of use cases, including weather forecasting, surface imaging, communications, and video broadcasts. Ground stations are at the core of global satellite networks, which are facilities that provide communications between the ground and the satellites by using antennas to receive data and control systems to send radio signals to command and control the satellite.

We provide direct access to AWS services and the AWS Global Infrastructure including our low-latency global fiber network right where your data is downloaded into our AWS Ground Station. This enables you to easily control satellite communications, quickly ingest and process your satellite data, and rapidly integrate that data with your applications and other services running in the AWS Cloud.

SECURITY, IDENTITY, AND COMPLIANCE

Topics

- AWS Security Hub
- Amazon Cloud Directory
- AWS Identity and Access Management
- Amazon GuardDuty
- Amazon Inspector
- Amazon Macie
- AWS Artifact
- AWS Certificate Manager
- AWS CloudHSM
- AWS Directory Service
- AWS Firewall Manager
- AWS Key Management Service
- AWS Organizations
- AWS Secrets Manager
- AWS Shield
- AWS Single Sign-On
- AWS WAF

AWS Security Hub

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

With Security Hub, you now have a single place that aggregates, organizes, and prioritizes your security alerts, or findings, from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, and Amazon Macie, as well as from AWS Partner solutions.

Amazon Cloud Directory

Amazon Cloud Directory enables you to build flexible, cloud-native directories for organizing hierarchies of data along multiple dimensions. With Cloud Directory, you can create directories for a variety of use cases, such as organizational charts, course catalogs, and device registries. While traditional directory solutions, such as Active Directory Lightweight Directory Services (AD LDS) and other LDAP-based directories, limit you to a single hierarchy, Cloud Directory offers you the flexibility to create directories with hierarchies that span multiple dimensions.

AWS Identity and Access Management

AWS Identity and Access Management (IAM) enables you to securely control access to AWS services and resources for your users. Using IAM, you can create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. IAM allows you to do the following:

• Manage IAM users and their access: You can create users in IAM, assign them individual security credentials (access keys, passwords, and multi-factor authentication devices), or request temporary security credentials to provide users access to AWS services and resources. You can manage permissions in order to control which operations a user can perform.
• Manage IAM roles and their permissions: You can create roles in IAM and manage permissions to control which operations can be performed by the entity, or AWS service, that assumes the role. You can also define which entity is allowed to assume the role.
• Manage federated users and their permissions: You can enable identity federation to allow existing identities (users, groups, and roles) in your enterprise to access the AWS Management Console, call AWS APIs, and access resources, without the need to create an IAM user for each identity.

Amazon GuardDuty

Amazon GuardDuty is a threat detection service that continuously monitors for malicious or unauthorized behavior to help you protect your AWS accounts and workloads. It monitors for activity such as unusual API calls or potentially unauthorized deployments that indicate a possible account compromise. GuardDuty also detects potentially compromised instances or reconnaissance by attackers. There is a 30-dayfree trial available for every new account to the service.

Amazon Inspector

Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity. These findings can be reviewed directly or as part of detailed assessment reports which are available via the Amazon Inspector console or API.

Examples of built-in rules include checking for access to your EC2 instances from the internet, remote root login being enabled, or vulnerable software versions installed.

Amazon Macie

Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes sensitive data such as personally identifiable

information (PII) or intellectual property, and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved.

AWS Artifact

It provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA).

AWS Certificate Manager

AWS Certificate Manager is a service that lets you easily provision, manage, and deploy Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and your 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 time-consuming manual process of purchasing, uploading, and renewing SSL/TLS certificates.

AWS CloudHSM

The AWS CloudHSM is a cloud-based hardware security module (HSM) that enables you to easily generate and use your own encryption keys on the AWS Cloud. With CloudHSM, you can manage your own encryption keys using FIPS 140-2 Level 3 validated HSMs. CloudHSM offers you the flexibility to integrate with your applications using industry-standard APIs, such as PKCS#11, Java Cryptography Extensions (JCE), and Microsoft CryptoNG (CNG) libraries.

AWS Directory Service

AWS Directory Service for Microsoft Active Directory, also known as AWS Managed Microsoft AD, enables your directory-aware workloads and AWS resources to use managed Active Directory in the AWS Cloud

AWS Managed Microsoft AD is built on actual Microsoft Active Directory and does not require you to synchronize or replicate data from your existing Active Directory to the cloud. You can use standard Active Directory administration tools and take advantage of built-in Active Directory features such as Group Policy and single sign-on (SSO). With AWS Managed Microsoft AD, you can easily join Amazon EC2 and Amazon RDS for SQL Server instances to a domain, and use AWS Enterprise IT applications such as Amazon WorkSpaces with Active Directory users and groups.

AWS Firewall Manager

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

AWS Key Management Service

AWS Key Management Service (KMS) makes it easy for you to create and manage keys and control the use of encryption across a wide range of AWS services and in your applications. AWS KMS is a secure and resilient service that uses FIPS 140-2 validated hardware security modules to protect your keys. AWS KMS is integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.

AWS Organizations

AWS Organizations offers policy-based management for multiple AWS accounts. With Organizations, you can create groups of accounts, automate account creation, apply and manage policies for those groups. Organizations enables you to centrally manage policies across multiple accounts, without requiring custom scripts and manual processes.

Using AWS Organizations, you can create Service Control Policies (SCPs) that centrally control AWS service use across multiple AWS accounts. You can also use Organizations to help automate the creation of new accounts through APIs. Organizations helps simplify the billing for multiple accounts by enabling you to setup a single payment method for all the accounts in your organization through consolidated billing.

AWS Secrets Manager

AWS Secrets Manager helps you protect secrets needed to access your applications, services, and IT resources. The service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle.

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. There are two tiers of AWS Shield: Standard and Advanced.

All AWS customers benefit from the automatic protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring network and transport layer DDoS attacks that target your website or applications. When you use AWS Shield Standard with Amazon CloudFront and Amazon Route 53, you receive comprehensive availability protection against all known infrastructure (Layer 3 and 4) attacks.

For higher levels of protection against attacks targeting your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, and Amazon Route 53 resources, you can subscribe to AWS Shield Advanced. In addition to the network and transport layer protections that come with Standard, AWS Shield Advanced provides additional detection and mitigation against large and sophisticated DDoS attacks, near real-time visibility into attacks, and integration with AWS WAF, a web application firewall. AWS Shield Advanced also gives you 24x7 access to the AWS DDoS Response Team (DRT) and protection against DDoS related spikes in your Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, and Amazon Route 53 charges.

AWS Single Sign-On

AWS Single Sign-On (SSO) is a cloud SSO service that makes it easy to centrally manage SSO access to multiple AWS accounts and business applications.

AWS SSO also includes built-in SAML integrations to many business applications, such as Salesforce, Box, and Office 365. Further, by using the AWS SSO application configuration wizard, you can create Security Assertion Markup Language (SAML) 2.0 integrations and extend SSO access to any of your SAML-enabled applications.

AWS WAF

AWS WAF is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF gives you control over which traffic to allow or block to your web application by defining customizable web security rules.

STORAGE

Topics

- Amazon S3
- Amazon Elastic Block Store
- Amazon Elastic File System
- Amazon FSx for Lustre
- Amazon FSx for Windows File Server
- Amazon S3 Glacier
- AWS Storage Gateway

Amazon S3

Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance. This means customers of all sizes and industries can use it to store and protect any amount of data for a range of use cases, such as websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics.

Amazon Elastic Block Store

Amazon Elastic Block Store (Amazon 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 you from component failure, offering high availability and durability

Amazon Elastic File System

Amazon Elastic File System (Amazon EFS) provides a simple, scalable, elastic file system for Linux-based workloads for use with AWS Cloud services and on-premises resources. It is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files, so your applications have the storage they need – when they need it.

You can access your file systems across AZs and regions and share files between thousands of Amazon EC2 instances and on-premises servers via AWS Direct Connect or AWS VPN.

Use cases such as lift-and-shift enterprise applications, big data analytics, web serving and content management, application development and testing, media and entertainment workflows, database backups, and container storage.

Amazon FSx for Lustre

Amazon FSx for Lustre is a fully managed file system that is optimized for compute-intensive workloads, such as high performance computing, machine learning, and media data processing workflows. Many of these applications require the high-performance and low latencies of scale-out, parallel file systems you can launch and run a Lustre file system that can process massive data sets at up to hundreds of gigabytes per second of throughput, millions of IOPS, and sub-millisecond latencies. Amazon FSx for Lustre is seamlessly integrated with Amazon S3, making it easy to link your long term data sets with your high performance file systems to run compute-intensive workloads.

Amazon FSx for Windows File Server

Amazon FSx for Windows File Server provides a fully managed native Microsoft Windows file system so you can easily move your Windows-based applications that require file storage to AWS. Built on Windows Server, Amazon FSx provides shared file storage with the compatibility and features that your Windowsbased applications rely on, including full support for the SMB protocol and Windows NTFS, Active Directory (AD) integration, and Distributed File System (DFS). Amazon FSx uses SSD storage to provide the fast performance your Windows applications and users expect, with high levels of throughput and IOPS, and consistent sub-millisecond latencies. This compatibility and performance is particularly

important when moving workloads that require Windows shared file storage, like CRM, ERP, and .NET applications, as well as home directories.
With Amazon FSx, you can launch highly durable and available Windows file systems that can be accessed from up to thousands of compute instances using the industry-standard SMB protocol.

Amazon S3 Glacier

Amazon S3 Glacier is a secure, durable, and extremely low-cost storage service for data archiving and long-term backup.You can store data for as little as $0.004 per gigabyte per month,

AWS Storage Gateway

The AWS Storage Gateway is a hybrid storage service that enables your on-premises applications to seamlessly use AWS cloud storage. You can use the service for backup and archiving, disaster recovery, cloud data processing, storage tiering, and migration. Your applications connect to the service through a virtual machine or hardware gateway appliance using standard storage protocols, such as NFS, SMB and iSCSI. The gateway connects to AWS storage services, such as Amazon S3, Glacier, and Amazon EBS, providing storage for files, volumes, and virtual tapes in AWS. The service includes a highly-optimized data transfer mechanism, with bandwidth management, automated network resilience, and efficient data transfer, along with a local cache for low-latency on-premises access to your most active data.

CREATING AN AWS ACCOUNT

To create an account, one needs the following:

- A telephone number to validate your identity
- A credit card to pay your bills

Signing up

The sign-up process consists of five steps:

1. Provide your login credentials.
2. Provide your contact information.
3. Provide your payment details.
4. Verify your identity.
5. Choose your support plan.

Point your favorite modern web browser to https://aws.amazon.com, and click the Create a Free Account / Create an AWS Account button.

HOW MUCH DOES IT COST?

Use the AWS Simple Monthly Calculator (http://aws.amazon.com/calculator)

FREE TIER

Here is what's included in the Free Tier:

- 750 hours (roughly a month) of a small virtual server running Linux or Windows. This means you can run one virtual server the whole month or you can run 750 virtual servers for one hour.
- 750 hours (or roughly a month) of a load balancer.
- Object store with 5 GB of storage.
- Small database with 20 GB of storage, including backup.

If your Free Tier ends after one year, you pay for all resources you use. You get some additional benefits, as detailed at http://aws.amazon.com/free.

BILLING

- *Based on hours of usage*—If you use a server for 61 minutes, that's usually counted as 2 hours.
- *Based on traffic*—Traffic can be measured in gigabytes or in number of requests.
- *Based on storage usage*—Usage can be either provisioned capacity (for example, 50 GB volume no matter how much you use) or real usage (such as 2.3 GB used)

Pay-per-use opportunities

You no longer need to make upfront investments in infrastructure. You can start servers on demand and only pay per hour of usage; and you can stop using those servers whenever you like and no longer have to pay for them.

REFERENCES

Amazon Inc (2019) Amazon Elastic Compute Cloud - User Guide for Linux Instances, Amazon Web Services Inc.

Amazon Inc (2019) Amazon Web Services Documentation, AWS inc

Amazon Inc (2019) Amazon Web Services Whitepaper, AWS inc

Baron J., Baz H., Bixler T., Gaut B., Kelly K.E., Senior S., Stamper J., (2017) AWS Certified Solutions Architect Official Study Guide, John Wiley & Sons, Inc. Indianapolis, Indiana

Beach B.(2014) Pro Powershell for Amazon Web Services – DevOps for the AWS Cloud, Apress Media, California USA.

Golden B. (2013) Amazon Web Services™ For Dummies, John Wiley & Sons, Inc. New Jersey USA

Hashimoto N. (August 2013) Amazon S3 Cookbook, PACKT Publishing, Birmingham, United Kingdom

Kvreddi, Hand Book For AWS Admin, SUVEN IT, Hyderabad India [Online] Available from: https://www.scribd.com/document/.../Aws-Admin-Guide-by-Suven-It-v1-PDF [Accessed date 24 May 2019 16:24hrs]

Matthew S. (2019) Overview of Amazon Web Services AWS Whitepaper, AWS inc

Mueller J.P., (2017) AWS For Admins (dummies), John Wiley & Sons, Hoboken, New Jersey

Varia J., Mathew S. (2014), Overview of Amazon Web Services [Online] Available from: https://media.amazonwebservices.com/AWS_Overview.pdf [Accessed date 24 May 2019 16:24hrs]

Wadia Y. (February 2016) AWS Administration – The Definitive Guide, Packt Publishing, Birmingham UK

Witting A, Witting M, (2016) Amazon Web services in Action. Manning Publications Co Shelter Island, New York