

Are you Ready for AWS Certified SysOps Administrator - Associate exam? Self-assess yourself with "[Whizlabs FREE TEST](#)"



AWS Certified SysOps Administrator - Associate (SOA-C02) WhizCard

Quick Bytes for you before the exam!

The information provided in WhizCard is for educational purposes only; created in our efforts to help aspirants prepare for the AWS Certified SysOps Administrator - Associate certification exam. Though references have been taken from AWS documentation, it's not intended as a substitute for the official docs. The document can be reused, reproduced, and printed in any form; ensure that appropriate sources are credited and required permissions are received.



Service Names	Page No.
Analytics	
1. Amazon OpenSearch Service	5
Application Integration	
2. AWS EventBridge	7
3. AWS SNS	8
4. AWS SQS	9
Cloud Financial Management	
5. AWS Cost & Usage Report	11
6. AWS Cost Explorer	12
Compute	
7. AWS EC2	14
8. Amazon EC2 Auto Scaling	15
9. Amazon EC2 Image Builder	16
10. AWS Lambda	17
Databases	
11. Amazon Aurora	19
12. Amazon DynamoDB	20
13. Amazon ElastiCache	21
14. Amazon RDS	22

Service Names	Page No.
Migration and Transfer	
15. AWS DataSync	24
16. AWS Transfer Family	25
Networking and Content Delivery	
26. Amazon CloudFront	28
27. Elastic Load Balancing (ELB)	29
28. Amazon Route 53	30
29. Amazon VPC	31
30. AWS Transit Gateway	32
Security, Identity, and Compliance	
31. AWS Certificate Manager	35
32. Amazon Detective	36
33. AWS Directory Service	38
34. Amazon GuardDuty	40
35. Amazon Identity and Access Management (IAM)	42
36. Amazon Inspector	43
37. AWS Key Management Service	44
38. AWS Secrets Manager	45
39. AWS Security Hub	46

Service Names	Page No.
40. AWS WAF	47
Storage	
41. AWS Backup	50
42. Amazon Elastic Block Store	52
43. Amazon Elastic File System (EFS)	54
44. Amazon Simple Storage Service (S3)	55
45. Amazon S3 Glacier	56
46. AWS Storage Gateway	57
47. Types of Storage Gateway	58

Service Names	Page No.
Management and Governance	
48. AWS CloudFormation	61
49. AWS CloudTrail	62
50. Amazon CloudWatch	63
51. Amazon CloudWatch Logs	64
52. AWS Control Tower	66
53. AWS Organizations	68
54. AWS Systems Manager	69
55. AWS Trusted Advisor	70
56. AWS Resource Access Manager	72



Analytics

Amazon OpenSearch Service








OpenSearch Service is a free and open-source search engine for all types of data like textual, numerical, geospatial, structured, and unstructured.



What is Amazon OpenSearch Service?

Amazon OpenSearch Service is a managed service that allows users to deploy, manage, and scale OpenSearch clusters in the AWS Cloud. It provides direct access to the OpenSearch APIs.

Amazon OpenSearch Service can be integrated with following services:

- Amazon CloudWatch 
- Amazon CloudTrail 
- Amazon Kinesis 
- Amazon S3 
- AWS IAM 
- AWS Lambda 
- Amazon DynamoDB 

- ✓ Amazon OpenSearch Service with Kibana (visualization) & Logstash (log ingestion) provides an enhanced search experience for the applications and websites to find relevant data quickly.
- ✓ Amazon OpenSearch Service launches the OpenSearch cluster's resources and detects the failed OpenSearch nodes and replaces them.
- ✓ The OpenSearch Service cluster can be scaled with a few clicks in the console.

Pricing Details:



- Charges are applied for each hour of use of EC2 instances and storage volumes attached to the instances.
- Amazon OpenSearch Service does not charge for data transfer between availability zones.



Application Integration

Amazon EventBridge

What is Amazon EventBridge?







Amazon EventBridge is a serverless event bus service that connects applications with data from multiple sources.



Amazon EventBridge integrates with the following services:

- AWS CloudTrail
- AWS CloudFormation
- AWS Config
- AWS Identity and Access Management (IAM)
- AWS Kinesis Data Streams
- AWS Lambda


Functions of Amazon EventBridge:

-  An event bus is an entity that receives events, and rules get attached to that event bus that matches the events received.
-  It helps to build loosely coupled and distributed event-driven architectures.
-  It connects applications and delivers the events without the need to write custom code.
-  It delivers a stream of real-time data from SaaS applications or other AWS services and routes that data to different targets such as Amazon EC2 instances, Amazon ECS tasks, AWS CodeBuild projects, etc.
-  It sets up routing rules that determine the targets to build application architectures that react according to the data sources.
-  The EventBridge schema registry stores a collection of event structures (schemas) and allows users to download code for those schemas in the IDE representing events as objects in the code.

Amazon SNS

What is Amazon SNS?

Amazon Simple Notification Service (Amazon SNS) is a serverless notification service that offers message delivery from publishers to subscribers.









- ✓ It creates asynchronous communication between publishers and subscribers by sending messages to a 'topic.'
- ✓ It supports application-to-application subscribers that include Amazon SQS and other AWS services and Application-to-person subscribers that include Mobile SMS, Email, etc.

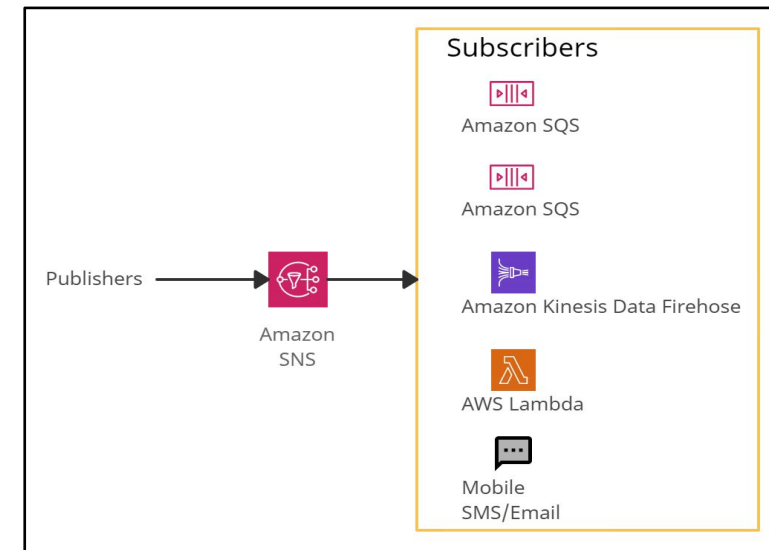
- The producer sends one message to one SNS topic.
- Multiple receivers (subscribers) listen for the notification of messages.
- All the subscribers will receive all the messages.

Example:

1 message, 1 topic, 10 subscribers so that a single message will be notified to 10 different subscribers.

SNS helps to publish messages to many subscriber endpoints:

- Amazon SQS Queues 
- AWS Lambda Functions 
- Email 
- Amazon Kinesis Data Firehose 
- Mobile push 
- SMS 



Amazon SNS

Amazon SQS

What are Amazon Simple Queue Service (SQS)?

Amazon Simple Queue Service (SQS) is a serverless service used to decouple (loose couple) serverless applications and components.

- ❑ The queue represents a temporary repository between the producer and consumer of messages.
- ❑ It can scale up to 1-10000 messages per second.
- ❑ The default retention period of messages is four days and can be extended to fourteen days.
- ❑ SQS messages get automatically deleted after being consumed by the consumers.
- ❑ SQS messages have a fixed size of 256KB.

There are two SQS Queue types:

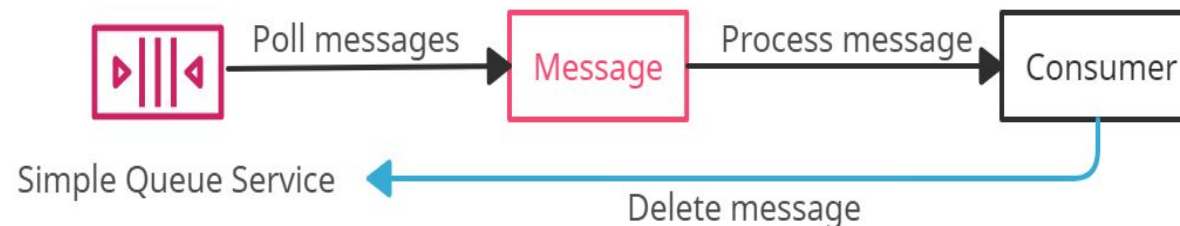
Standard Queue -

- ❖ The unlimited number of transactions per second.
- ❖ Messages get delivered in any order.
- ❖ Messages can be sent twice or multiple times.

FIFO Queue -

- ❖ 300 messages per second.
- ❖ Support batches of 10 messages per operation, results in 3000 messages per second.
- ❖ Messages get consumed only once.

Dead-Letter Queue is a queue for those messages that are not consumed successfully. It is used to handle message failure.



Delay Queue is a queue that allows users to postpone/delay the delivery of messages to a **queue** for a specific number of seconds. Messages can be delayed for 0 seconds (default) -15 (maximum) minutes.

Visibility Timeout is the amount of time during which SQS prevents other consumers from receiving (poll) and processing the messages.

- Default visibility timeout - 30 seconds
- Minimum visibility timeout - 0 seconds
- Maximum visibility timeout - 12 hours



Cloud Financial Management

What is AWS Cost & Usage Report?

AWS Cost & Usage Report (AWS CUR) allows users to access the detailed set of AWS cost and usage data available, including metadata about AWS resources, pricing, Reserved Instances, and Savings Plans.

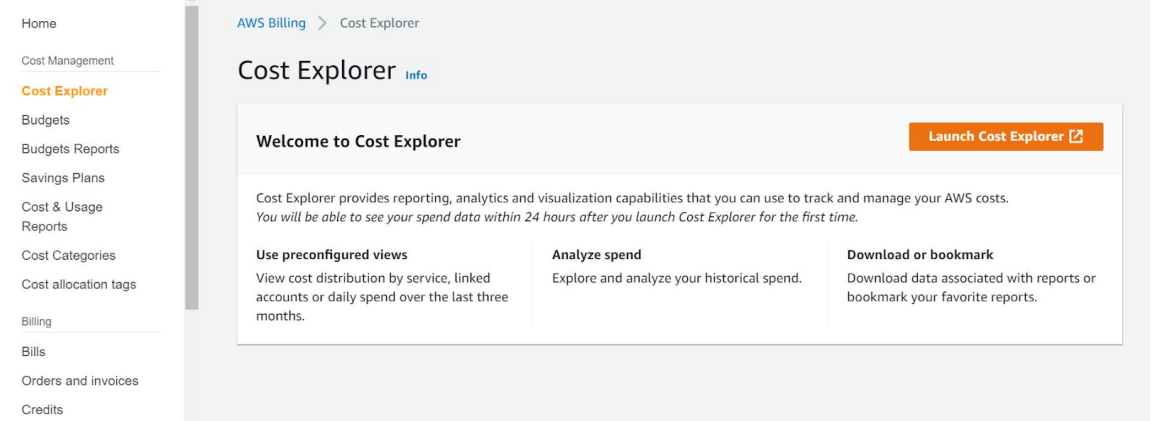
AWS Cost & Usage Report is a part of AWS Cost Explorer.

- AWS Cost and Usage Reports functions:
 - It sends report files to your Amazon S3 bucket.
 - It updates reports up to three times a day.
 - It creates, retrieves, and deletes reports using the AWS CUR API Reference.
- There is a feature of the Data Dictionary that lists the columns added in the report to easily analyze cost and usage in detail.
- For viewing, reports can be downloaded from the Amazon S3 console, for analyzing the report Amazon Athena can be used, or upload the report into Amazon Redshift or Amazon QuickSight.
- Users with IAM permissions or IAM roles can access and view the reports.
- If a member account in an organization owns or creates a Cost and Usage Report, then it can have access only to billing data for the time it has been a member of the Organization.
- If the master account of an AWS Organization wants to block access to the member accounts to set up a Cost and Usage Report, a Service Control Policy (SCP) can be used.

AWS Cost Explorer

What is AWS Cost Explorer?

AWS Cost Explorer is a UI-tool that enables users to analyze the costs and usage with the help of a graph, the Cost Explorer cost and usage reports, and/or the Cost Explorer RI report. It can be accessed from the Billing and Cost Management console.



AWS Cost Explorer

It provides default reports for analysis with some filters and constraints to create the reports. Analysis using Cost Explorer can be saved as a bookmark, CSV file download, or save them as a report.

The default reports provided by Cost Explorer are:

- **Cost and usage reports:**

It provides the following data for understanding the costs:-

- AWS Marketplace
- Daily costs
- Monthly costs by linked account
- Monthly costs by service
- Monthly EC2 running hours costs and usage

- **Reserved Instance reports:**

It provides the following reports for understanding the reservations:-

- **RI utilization reports:** It gives information about how much costs are saved or overspent by using Reserved Instances (RIs).
- **RI coverage reports:** It gives information about how many hours are saved or overspent by using Reserved Instances (RIs).

- The first time that the user signs up for Cost Explorer, it directs through the main parts of the console. It prepares the data regarding costs & usage and displays up to 12 months of historical data (might be less if less used), current month data, and then calculates the forecast data for the next 12 months.
- It uses the same set of data that is used to generate the AWS Cost and Usage Reports and the billing reports.
- It provides a custom time period to view the data at a monthly or daily interval.
- It provides a feature of Savings Plans which provides savings of up to 72% on the AWS compute usage.
- It provides a way to access the data programmatically using the Cost Explorer API.

Price details:

- Analysis of costs and usage using the Cost Explorer can be viewed free of charge.
- The cost of using AWS Cost Explorer API is \$0.01 per API request.




Compute

Amazon EC2

What is Amazon EC2?

Amazon Elastic Compute Cloud (Amazon EC2) is a service that provides secure and scalable compute capacity in the AWS cloud. It falls under the category of Infrastructure as a Service (IAAS).




It provides the different type of instances based on the pricing models:

On-Demand Instances

- ✓ Useful for short-term needs, unpredictable workloads.
- ✓ No advance payment, no prior commitment.

Spot Instances

- ✓ No advance payment, no prior commitment.
- ✓ Useful for cost-sensitive compute workloads.



Reserved Instances

- ✓ Useful for long-running workloads and predictable usage.
- ✓ Offer to choose from No upfront, Partial upfront, or All upfront.

Dedicated Instances

- ✓ Instances run on hardware dedicated to a single user.
- ✓ Other customers can not share the hardware.

Dedicated Hosts

- ✓ A whole physical server with an EC2 instance allocates to an organization.

It provides different compute platforms and instance types based on price, CPU, operating system, storage, and networking, and each instance type consists of one or more instance sizes. Eg., t2.micro, t4g.nano, m4.large, r5a.large, etc.

It provides pre-configured templates that package the operating system and other software for the instances. This template is called Amazon Machine Images (AMIs).

It helps to login into the instances using key-pairs, in which AWS manages the public key, and the user operates the private key.

It also provides firewall-like security by specifying IP ranges, type, protocols (TCP), port range (22, 25, 443) using security groups.

It provides temporary storage volumes known as instance store volumes, which are deleted if the instance gets stopped, hibernated, or terminated. It also offers non-temporary or persistent volumes known as Amazon EBS volumes.


It enables users to write scripts under the option 'User data,' used at the instances' launch.

It offers to choose from three IP addresses, which are Public IP address (Changes when the instance is stopped or refreshed), Private IP address (retained even if the model is stopped), Elastic IP address (static public IP address).

Amazon EC2 Auto Scaling

What is Amazon EC2 Auto Scaling?

Amazon EC2 Auto Scaling is a region-specific service used to maintain application availability and enables users to automatically add or remove EC2 instances according to the compute workloads.



- ❖ The Auto Scaling group is a collection of the minimum number of EC2 used for high availability.
- ❖ It enables users to use Amazon EC2 Auto Scaling features such as fault tolerance, health check, scaling policies, and cost management.
- ❖ The scaling of the Auto Scaling group depends on the size of the desired capacity. It is not necessary to keep DesiredCapacity and MaxSize equal.

E.g.,
DesiredCapacity: '2' - There will be total 2 EC2 instances
MinSize: '1'
MaxSize: '2'

- ❖ EC2 Auto Scaling supports automatic Horizontal Scaling (increases or decreases the number of EC2 instances) rather than Vertical Scaling (increases or decreases EC2 instances like large, small, medium).

Launch Template


A **launch template** is similar to launch configuration with extra features as below

It launches both Spot and On-Demand instances.


It specifies multiple instance types

It specifies multiple launch templates.

It scales across multiple Availability Zones within the same AWS region.



Horizontal Scaling



Vertical Scaling

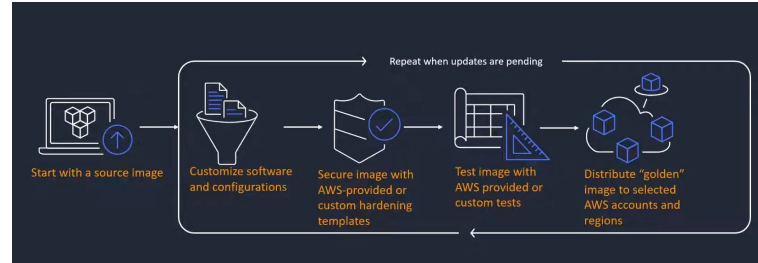
Amazon EC2 Image Builder

What is Amazon EC2 Image Builder?

EC2 Image Builder is a tool that is used for creating, customizing, managing, securing, and distributing Virtual Machine and container images for its use on AWS or on-premises environments.

Organizations can build standardized images, test, validate, and ship them for Production use using EC2 Image Builder.

EC2 Image Builder is a managed service offering that simplifies the end-end Image-building workflow and ensures that the images are distributed across Accounts and regions that need it.



AWS Documentation

Use cases:

- Automation for maintaining up-to-date secure images. Amazon EC2 Image Builder significantly reduces the effort here.
- Improving service uptime of images in production. Amazon EC2 Image Builder allows for testing of built images for validating applications on updated builds.
- Integrated Security - Amazon EC2 Image Builder simplifies securing VMs. Images can be configured to only include essential components.

Best Practices:

- It is advisable to share EC2 Image Builder images with trusted accounts only.
- Images with private or sensitive data should not be made Public.
- It is advisable to apply all Windows/Linux security patches during image builds.
- Base images used as recipes within Image Builder should follow AWS's least security privileges best practice for its Security Groups.

Features:

- EC2 Image Builder uses a Built-in image pipeline to automate the end-end workflow.
- The pipeline can be triggered manually, in a Schedule, when a source image or a component has been updated.
- The image pipeline workflow includes the build & test phases. The test phase validates whether the image will work as expected before shipping it for production use.
- With Image Builder, you can enhance the security of your deployments by applying AWS security settings for creating images that meet security criteria within an organization.
- Image Builder's integration with AWS RAM allows for easy sharing of Image Builder resources with AWS Accounts or through AWS Organizations.

Exam Tip:

- Amazon EC2 Image Builder does not incur any cost. The cost comes with associated AWS resources that are used to create, store & share the images.
- Amazon EC2 Image Builder has integrations with AWS RAM, Amazon ECR & AWS Organizations that enable it to share automation scripts, recipes & images across AWS Accounts.
- The Base Image for Image Builder can be an existing AMI, Custom AMI, or AMI from an imported image from VMDK, VHDX, or OVF formats.

AWS Lambda

What is AWS Lambda?

AWS Lambda is a serverless computing service that allows users to run code as functions without provisioning or managing servers.



It helps to run the code on highly-available infrastructure and performs administrative tasks like server maintenance, logging, capacity provisioning, and automatic scaling and code monitoring.


Using AWS Lambda, one can build serverless applications composed of Lambda functions triggered by events and can be automatically deployed using AWS CodePipeline and AWS CodeBuild.

Amazon EC2	Amazon Lambda
They are termed virtual servers in the AWS cloud.	They are termed virtual functions.
It is limited to instance types (RAM and CPU).	Limited by time (less execution time of 300 seconds).
It runs continuously.	It runs on demand.
Scaling computing resources is manual.	It has automated scaling.

✓ The memory allocated to AWS Lambda for computing is 128MB (minimum) to 3008MB (maximum). Additional memory can be requested in an increment of 64MB between 128MB - 3008MB.

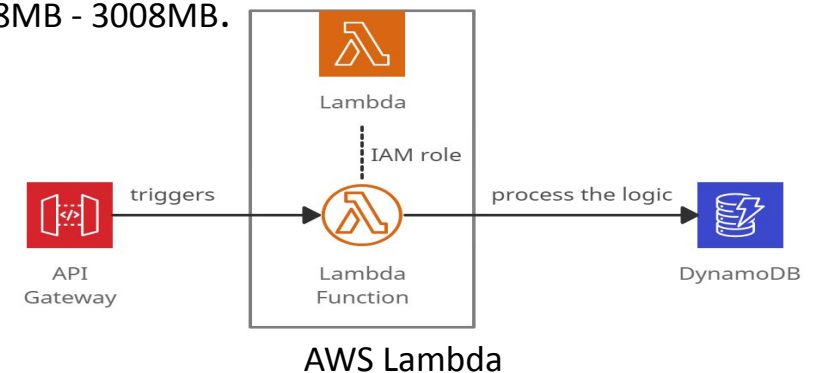
✓ The default execution time for AWS Lambda is 3 seconds, and the maximum is 15 minutes (900 seconds).

Lambda Functions supports the following languages:



Pricing details:

Charges are applied based on the number of requests for the functions and the time taken to execute the code





Databases

Amazon Aurora



What is Aurora?

Amazon Aurora is a MySQL and PostgreSQL-compatible, fully managed relational database engine built to enhance traditional enterprise databases' performance and availability.

- Is a part of the fully managed Amazon Relational Database Service (Amazon RDS).

Features include:

- RDS Management Console
- CLI commands and API operations for patching Backup
- Recovery
- Database Setup
- Failure Detection and repair


Performance



5x greater than



MySQL on RDS




3x greater than



PostgreSQL on RDS

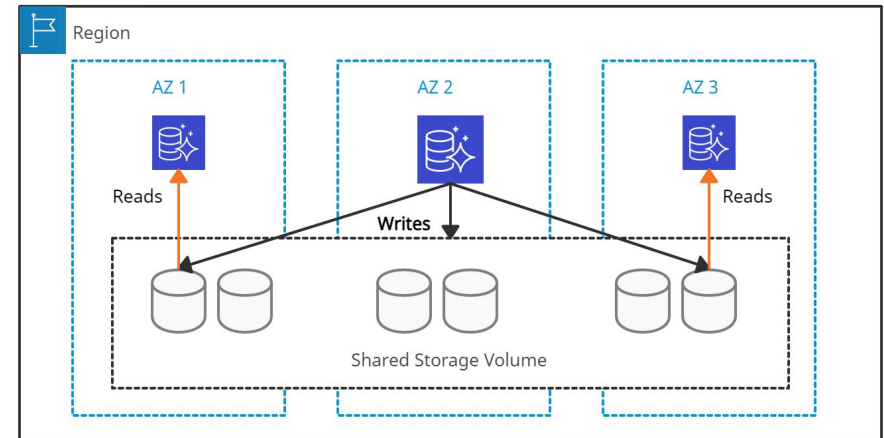
- Amazon Aurora replicates 2 copies of data in each availability zone (minimum of 3 AZ). So a total of 6 copies per region.

Data Replication : 2 Types



- Aurora replica (in-region)** It can provide 15 read replicas.
- MySQL Read Replica (cross-region)** It can provide 5 read replicas.

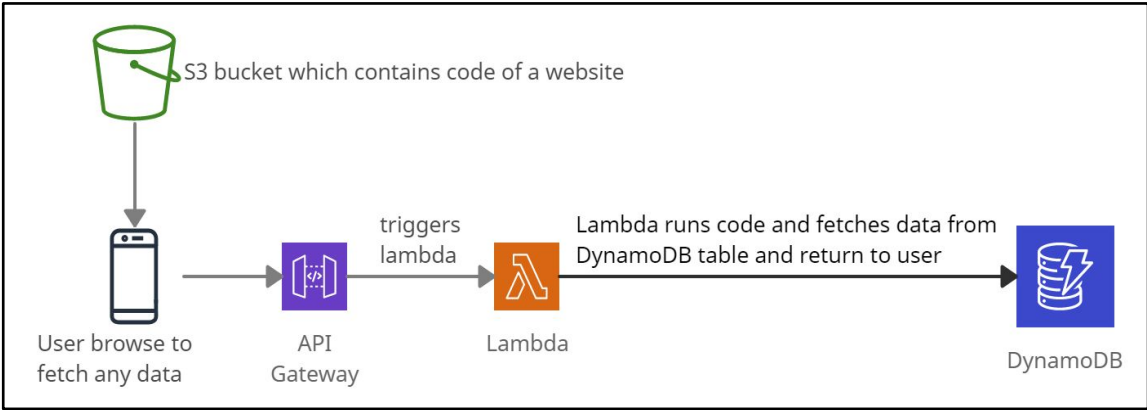
Amazon Aurora Cross-Region read replicas help to improve disaster recovery and provide fast reads in regions closer to the application users.



Amazon DynamoDB

What is Amazon DynamoDB?

Amazon DynamoDB is a serverless NoSQL database service that provides fast and predictable performance with single-digit millisecond latency.



Amazon DynamoDB example

- ▶ It provides a push button scaling feature, signifying that DB can scale without any downtime.
- ▶ It is a multi-region cloud service that supports key-value and document data structure.
- ▶ It provides high availability and data durability by replicating data synchronously on solid-state disks (SSDs) across 3 AZs in a region.
- ▶ It helps to store session states and supports ACID transactions for business-critical application
- ▶ It provides the on-demand backup capability of the tables for long-term retention and enables point-in-time recovery from accidental write or delete operations.
- ▶ Amazon DynamoDB Accelerator (DAX) is a highly available in-memory cache service that provides data from DynamoDB tables. DAX is not used for strongly consistent reads and write-intensive workloads.
- ▶ It supports Cross-Region Replication using DynamoDB Global Tables. Global Tables helps to deploy a multi-region database and provide automatic multi-master replication to AWS regions.

Amazon ElastiCache

What is Amazon ElastiCache?

ElastiCache is a web service used to manage and run in-memory data stores Redis and Memcached in the cloud.

- »» It is best suited for Online Analytical Processing (OLAP) transaction workloads and for storing session states.
- »» It has in-memory caching features to provide sub-millisecond latency for read-heavy application workloads.

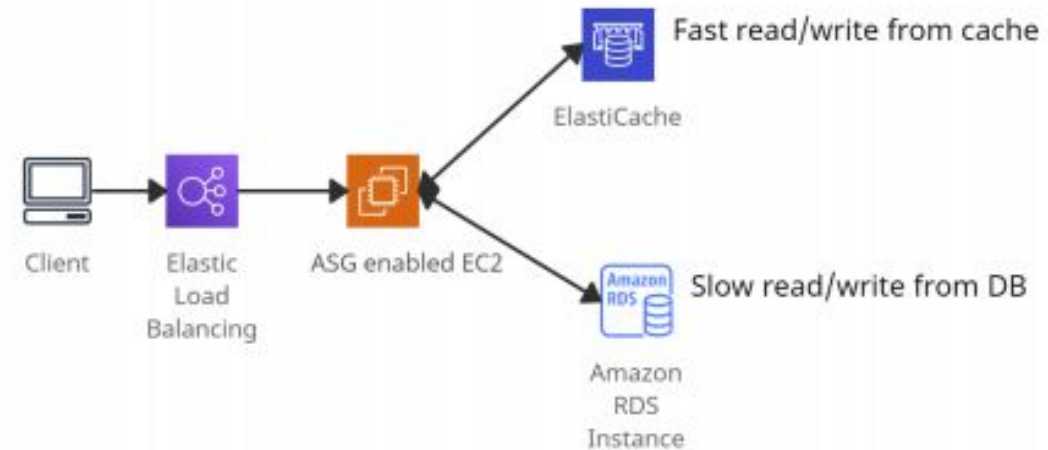
Amazon ElastiCache for Redis:

- ❖ It is useful for gaming applications, geospatial services, caching, session stores, and replication.
- ❖ Data is persistent.
- ❖ It is not multi-threaded.
- ❖ It supports Multi-AZ using read replicas.



Amazon ElastiCache for Memcached:

- ❖ It is useful for building applications that require caching layers.
- ❖ Data is not persistent.
- ❖ It supports multi-threading.
- ❖ It does not support Multi-AZ failover.
- ❖ It does not support snapshots.




Amazon ElastiCache

Amazon RDS

What is Amazon RDS?

Amazon Relational Database Service (Amazon RDS) is a service used to build and operate relational databases in the AWS Cloud



RDS provides read replicas of reading replicas and can also read replicas as a standby DB like Multi-AZ.

Read replicas feature is not available for SQL Server.

- It is best suited for structured data and Online Transaction Processing (OLTP) types of database workloads such as InnoDB.

It supports the following database engines:

- SQL Server
- PostgreSQL
- Amazon Aurora
- MySQL
- MariaDB
- Oracle

- ✓ If there is a need for unsupported RDS database engines, DB can be deployed on EC2 instances.

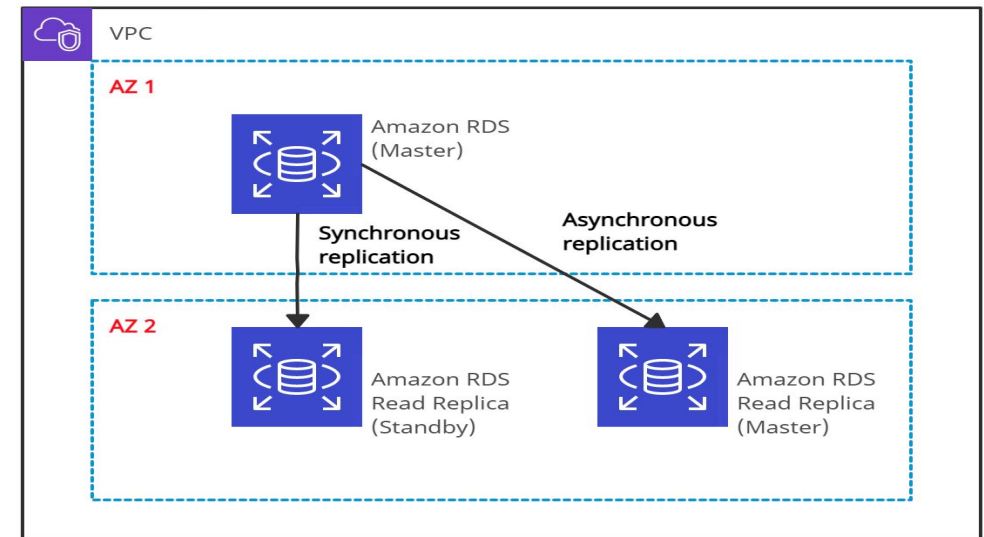
The following tasks need to be taken care of manually.

Encryption and Security

Updates and Backups

Disaster Recovery

- AWS KMS provides encryption at rest for RDS instances, DB snapshots, DB instance storage, and Read Replicas. The existing database cannot be encrypted.
- Amazon RDS only scales up for compute and storage, with no option for decreasing allocated storage
- It provides Multi-AZ and Read Replicas features for high availability, disaster recovery, and scaling.
 - **Multi-AZ Deployments** - Synchronous replication
 - **Read Replicas** - Asynchronous replication.



Amazon RDS



Migration and Transfer

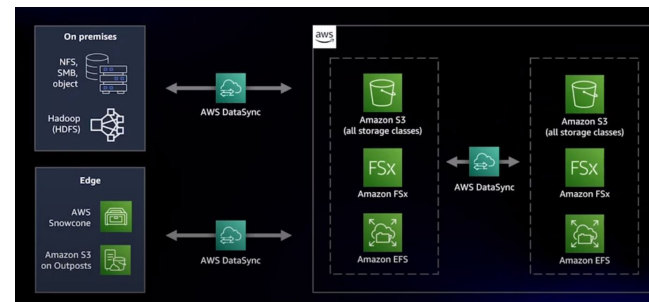
AWS DataSync

What is AWS DataSync?

AWS DataSync is a secure, reliable, managed migration Service that automates the movement of data online between storage systems. AWS DataSync provides the capability to move data between AWS storage, On-premises File Systems, Edge locations, and other Cloud Storage services like Azure. AWS DataSync helps you simplify your migration planning and reduce costs associated with the data transfer.

Use cases:

- Application Data Migration residing on On-premises storage systems like Windows Server, NAS file systems, Object storage to AWS.
- Archival of On-premises storage data to AWS to free capacity & reduce costs for continuously investing in storage infrastructure.
- Continuous replication of data present On-premises or on existing Cloud platforms for Data Protection and Disaster Recovery.



AWS Documentation

Features:

- Data movement workloads using AWS DataSync support migration scheduling, bandwidth throttling, task filtering, and logging.
- AWS DataSync provides enhanced performance using compression, and parallel transfers for transferring data at speed.
- AWS DataSync supports In-Flight encryption using TLS and encryption at rest.
- AWS DataSync provides capabilities for Data Integrity Verification ensuring that all data is transferred successfully.
- AWS DataSync integrates with AWS Management tools like CloudWatch, CloudTrail, and EventBridge.
- With DataSync, you only pay for the data you transfer without any minimum cost.
- AWS DataSync can copy data to and from Amazon S3 buckets, Amazon EFS file systems, and all Amazon FSx file system types.
- AWS DataSync supports Internet, VPN, and Direct Connect to transfer data between On-premises data centers, Cloud environments & AWS

Best Practices:

- In General, when planning a Data Migration, migration tools need to be evaluated, check for available bandwidth for online migration, and understand the source & destination migration data sources.
- For using DataSync to transfer data from On-premises storage to AWS, an Agent needs to be deployed and activated at On-premises locations. Use the Agent's local console as a tool for accessing various configurations
 - System resources
 - Network connectivity
 - Getting Agent activation key
 - View Agent ID & AWS region where the agent is activated
- A common pattern that can be used as a best practice is to use a combination of AWS DataSync & AWS Storage Gateway. DataSync can be used to archive On-premises data to AWS while Storage Gateway can be used to access commonly used data at On-premises.

Exam Tip:

- DataSync effectively manages the transfer of data between different storage devices without you having to write migration scripts or keep track of data that is transferred
- AWS DataSync can also be triggered using a Lambda function in case a migration schedule is not defined
- Data transfers between AWS services like S3 -> S3 or S3 -> EFS do not require the DataSync Agent. It is used only for data transfers from On-premises to AWS
- You pay 1.25 Cents per GigaByte of data transferred.

AWS Transfer Family

What is AWS Transfer Family?

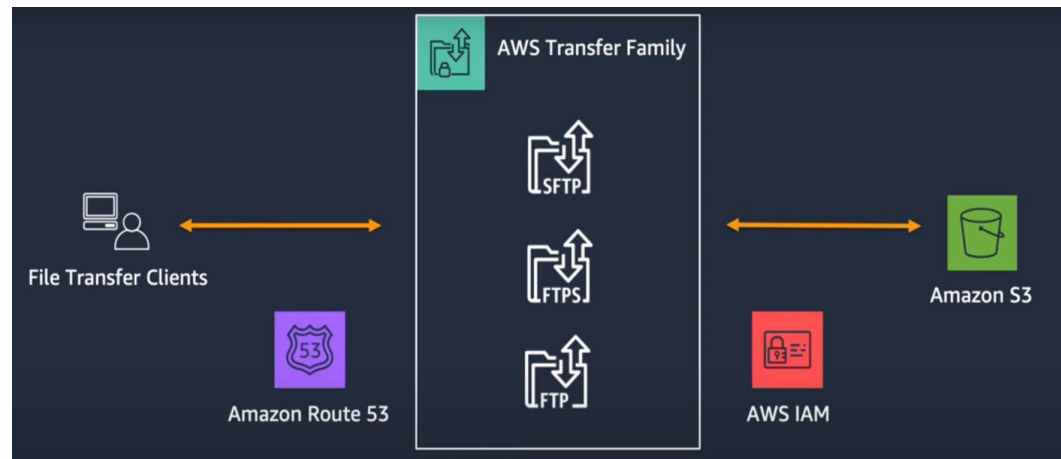
AWS Transfer Family is a fully managed & secure service that enables transfer of files using SFTP, FTPS & FTP.

The destination storage services to which files are transferred are S3, and EFS.

It helps you to seamlessly migrate File Transfer workloads to AWS without having any impact on existing application integrations or configuration

❖ Features:

- AWS Transfer Family provides a fully managed endpoint for transferring files into and out of S3, EFS.
- The Secure File Transport Protocol (SFTP) is a file transfer provided over SSH.
- File Transfer Protocol over SSL (FTPS is an FTP over a TLS-encrypted channel.
- Plain File Transfer Protocol (FTP) does not require a secure channel for transferring files.
- AWS Transfer Family exhibits high availability across the globe.
- AWS Transfer Family provides compliance with regulations within your Region.
- Using a pay-as-you-use model, the AWS Transfer Family service becomes cost-effective and is simple to use.
- AWS Transfer Family has the ability to use custom Identity Providers using AWS API Gateway & Lambda.



AWS Documentation

Exam Tip:

- IAM Roles are used to grant access to S3 buckets for file transfer clients in a secure way.
- Users can use Route 53 to migrate an existing File Transfer hostname for use in AWS.
- SFTP & FTPS protocols can be set up to be accessible from the public internet while FTP is limited for access from inside a VPC using VPC endpoints.

AWS Transfer Family

Best Practices:

- For improving Security posture while using AWS Transfer Family, the following best practices need to be adhered to.
- Use of a strong encryption mechanism for data in Transit - AWS Transfer Family provides a strong set of available ciphers for achieving this.
- Duplicate server's Host Key - This will ensure that an imposter will not impersonate your AWS Transfer Family server.
- Provide additional security on your AWS Transfer Family server to increase security posture - Use of both a password & a Key will help protect your clients if any of them are compromised.
- While using custom Identity Providers with AWS Transfer Family, set the authorizationType property of your API Gateway method to AWS_IAM - This will require the user to supply user credentials to be authenticated by the IdP.

Use cases:

- Use standard protocols (FTP) to get data into your application workflows like Data Lakes(S3) without having to rewrite applications using S3 API.
- Alleviate the challenges of a traditional MFT for managing, securing, and monitoring the MFT using AWS Transfer Family servers.
- MFT is found across industries
- Regulated industries like Financial Services and Health Care use secure document exchanges(e.g. Financial claims) where data needs to be governed.
- Supply chain where there is a trading partner relationship in its transactional data (eg purchase Orders).
- Content distribution - Software, audio, video




Networking and Content Delivery

Amazon CloudFront

What is Amazon CloudFront?

Amazon CloudFront is a content delivery network (CDN) service that securely delivers any kind of data to customers worldwide with low latency, low network, and high transfer speeds.



- It makes use of Edge locations (worldwide network of data centers) to deliver the content faster.
- Without edge locations, it retrieves data from an origin such as an Amazon S3 bucket, a Media Package channel, or an HTTP server.

CloudFront provides some security features such as:

- ❖ **Field-level encryption with HTTPS** - Data remains encrypted throughout starting from the upload of sensitive data.
- ❖ **AWS Shield Standard** - Against DDoS attacks.
- ❖ **AWS Shield Standard + AWS WAF + Amazon Route 53** - Against more complex attacks than DDoS.

CloudFront is integrated with AWS Services such as:

- Amazon S3
- Amazon EC2
- Elastic Load Balancing
- Amazon Route 53
- AWS Essential Media Services

Amazon CloudFront Access Controls:

Signed URLs:

- Use this to restrict access to individual files.

Signed Cookies:

- Use this to provide access to multiple restricted files.
- Use this if the user does not want to change current URLs.

Geo Restriction:

- Use this to **restrict** access to the data based on the geographic location of the website viewers.

Origin Access Identity (OAI):


- Outside access is restricted using signed URLs and signed cookies, but what if someone tries to access objects using Amazon S3 URL, bypassing CloudFront signed URL and signed cookies. To restrict that, OAI is used.
- Use OAI as a special CloudFront user and associate it with your CloudFront distribution to secure Amazon S3 content.

CloudFront Signed URL:

- It allows access to a path, no matter what is the origin
- It can be filtered by IP, path, date, expiration
- It leverages caching features

S3 Pre-Signed URL:


- It issues a request as the person who pre-signed the URL.



Elastic Load Balancing (ELB)

What is Elastic Load Balancing?

Elastic Load Balancing is a managed service that allows traffic to get distributed across EC2 instances, containers, and virtual appliances as target groups.



Elastic Load Balancing

Elastic Load Balancer types are as follows:

Application Load Balancer:

- Routes HTTP and HTTPS traffic at layer 7.
- Offers path-based routing, host-based routing, query-string, parameter-based routing, and source IP address-based routing.

Network Load Balancer:

- Routes TCP, UDP, and TLS traffic at layer 4.
- Suitable for high-performance and low latency applications.

Gateway Load Balancer:

- Suitable for third-party networking appliances.
- It simplifies tasks to manage, scale, and deploy virtual appliances

- ELB integrates with every AWS service throughout the applications.
- It is tightly integrated with Amazon EC2, Amazon ECS/EKS.
- ELB integrates with Amazon VPC and AWS WAF to offer extra security features to the applications.
- It helps monitor the servers' health and performance in real-time using Amazon CloudWatch metrics and request tracing.
- ELB can be placed based on the following aspects:
 - Internet-facing ELB:
 - Load Balancers have public IPs.
 - Internal only ELB:
 - Load Balancers have private IPs.
- ELB offers the functionality of Sticky sessions. It is a process to route requests to the same target from the same client.

Amazon Route 53

What is Route 53?

Route 53 is a managed DNS (Domain Name System) service where DNS is a collection of rules and records intended to help clients/users understand how to reach any server by its domain name.



- The most common records supported in Route 53 are:
- A: hostname to IPv4
 - AAAA: hostname to IPv6
 - CNAME: hostname to hostname
 - Alias: hostname to AWS resource

- **Route 53 hosted zone** is a collection of records for a specified domain that can be managed together.
- There are two types of zones:
 - **Public Hosted Zone** - Determines how traffic is routed on the Internet.
 - **Private Hosted Zone** - Determines how traffic is routed within VPC.

Route 53 CNAME	Route 53 Alias
It points a hostname to any other hostname.(app.mything.com -> abc.anything.com)	It points a hostname to an AWS Resource.(app.mything.com -> abc.amazonaws.com)
It works only for the non-root domains.(abcxyz.maindomain.com)	It works for the root domain and non-root domain. (maindomain.com)
It charges for CNAME queries.	It doesn't charge for Alias queries.
It points to any DNS record that is hosted anywhere.	It points to an ELB, CloudFront distribution, Elastic Beanstalk environment, S3 bucket as a static website, or another record in the same hosted zone.

Route 53 Routing Policies:

Simple:

- ❖ It is used when there is a need to redirect traffic to a single resource.
- ❖ It does not support health checks.

Weighted:

- ❖ It is similar to simple, but you can specify a weight associated with resources.
- ❖ It supports health checks.

Failover:

- ❖ If the primary resource is down (based on health checks), it will route to a secondary destination.
- ❖ It supports health checks.

Geo-location:

- ❖ It routes traffic to the closest geographic location you are in.

Geo-proximity:

- ❖ It routes traffic based on the location of resources to the closest region within a geographic area.

Latency based:

- ❖ It routes traffic to the destination that has the least latency.

Multi-value answer:

- ❖ It distributes DNS responses across multiple IP addresses.



Amazon Route 53

Amazon VPC

What is Amazon VPC?

Amazon Virtual Private Cloud is a service that allows users to create a virtual dedicated network for resources.

Private subnet - A subnet that does not have internet access is termed a private subnet.

Public subnet - A subnet that has internet access is termed a public subnet.

VPN only subnet - A subnet that does not have internet access but has access to the virtual private gateway for a VPN connection is termed a VPN-only subnet.

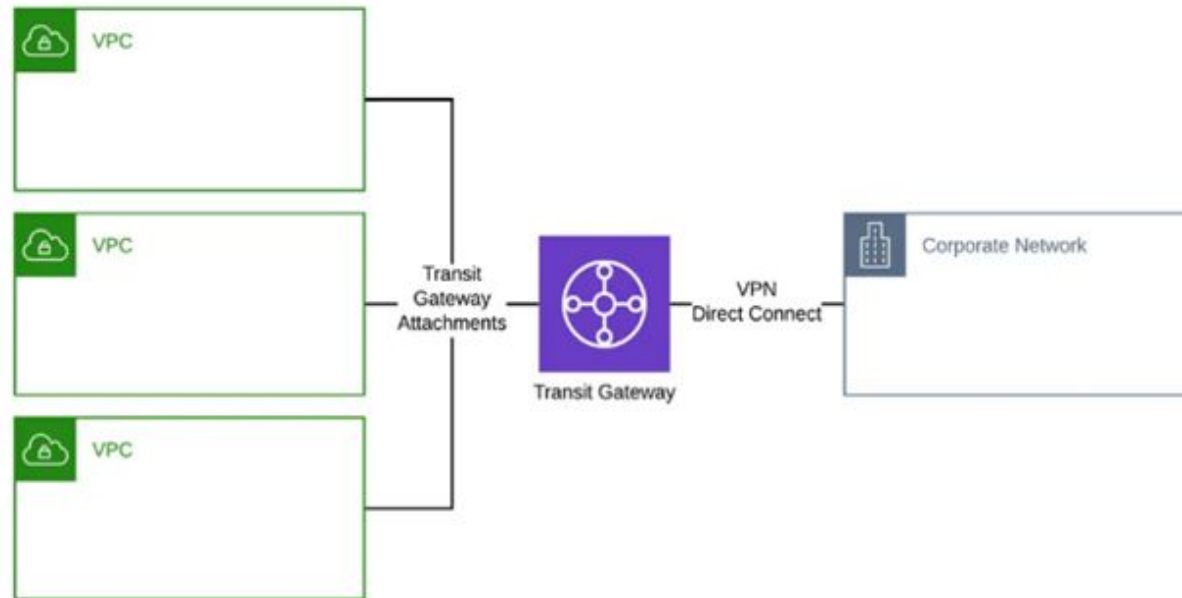
- ❑ It includes many components such as Internet gateways, VPN tools, CIDR, Subnets, Route tables, VPC endpoint, NAT instances, Bastion servers, Peering Connection, and others.
- ❑ It spans across multiple Availability Zones (AZs) within a region.
- ❑ The first four IP and last one IP addresses are reserved per subnet.
- ❑ It creates a public subnet for web servers that uses internet access and a private subnet for backend systems, such as databases or application servers.
- ❑ It can monitor resources using Amazon CloudWatch and Auto Scaling Groups.

- ❖ Every EC2 instance is launched within a default VPC with equal security and control like normal Amazon VPC. Default VPC has no private subnet.
- ❖ It uses Security Groups and NACL (Network Access Control Lists) for multi-layer security.
- ❖ Security Groups (stateful) provide instance-level security, whereas NACLs (stateless) provide subnet-level security.
- ❖ VPC sharing is a component that allows subnets to share with other AWS accounts within the same AWS Organization.

AWS Transit Gateway

What is AWS Transit Gateway?

AWS Transit Gateway is a network hub used to interconnect multiple VPCs. It can be used to attach all hybrid connectivity by controlling your organization's entire AWS routing configuration in one place.



AWS Transit Gateway

- It can be more than one per region but can not be peered within a single region.
- It helps to solve the problem of complex VPC peering connections.
- It can be connected with an AWS Direct Connect gateway from a different AWS account.
- Resource Access Manager (RAM) cannot integrate AWS Transit Gateway with Direct Connect gateway.
- To implement redundancy, Transit Gateway also allows multi-user gateway connections.
- Transit Gateway VPN attachment is a feature to create an IPsec VPN connection between your remote network and the Transit Gateway.
- Transit Gateway Network Manager is used to manage and monitor networking resources and connections to remote branch locations.
- It reduces the complexity of maintaining VPN connections with hundreds of VPCs, which become very useful for large enterprises.
- It supports attaching Amazon VPCs with IPv6 CIDRs.


Price details:

- Users will be charged for your AWS Transit Gateway on an hourly basis.

AWS Transit Gateway

Transit Gateway can be created using the following ways:

- AWS CLI
- AWS Management Console
- AWS CloudFormation



AWS Transit Gateway + VPN

Transit Gateway vs. VPC peering:

Transit Gateway	VPC peering
<ul style="list-style-type: none"> • It has an hourly charge per attachment in addition to the data transfer fees. • Multicast traffic can be routed between VPC attachments to a Transit Gateway. • It provides a Maximum bandwidth (burst) of 50 Gbps per Availability Zone per VPC connection. • Security groups feature does not currently work with Transit Gateway. 	<ul style="list-style-type: none"> • It does not charge for data transfer. • Multicast traffic cannot be routed to peering connections • It provides no aggregate bandwidth. • Security groups feature works with intra-Region VPC peering.



Security, Identity, and Compliance

AWS Certificate Manager

What is AWS Certificate Manager?

AWS Certificate Manager is a service that allows a user to protect AWS applications by storing, renewing, and deploying public and private SSL/TLS X.509 certificates.



- HTTPS transactions require server certificates X.509 that bind the public key in the certificate to provide authenticity.
- The certificates are signed by a certificate authority (CA) and contain the server's name, the validity period, the public key, the signature algorithm, and more.
- It centrally manages the certificate lifecycle and helps to automate certificate renewals.
- SSL/TLS certificates provide data-in-transit security and authorize the identity of sites and connections between browsers and applications.
- The certificates created by AWS Certificate Manager for using ACM-integrated services are free.
- With AWS Certificate Manager Private Certificate Authority, monthly charges are applied for the private CA operation and the private certificates issued.

The types of SSL certificates are:

Extended Validation Certificates (EV SSL)

Most expensive SSL certificate type

Organization Validated Certificates (OV SSL)

Validates a business' creditably.

Domain Validated Certificates (DV SSL)

Provides minimal encryption

Wildcard SSL Certificate

Secures base domain and subdomains.

Multi-Domain SSL Certificate (MDC)

Secure up to hundreds of domain and subdomains.

Unified Communications Certificate (UCC)

Single certificate secures multiple domain names.

Ways to deploy managed X.509 certificates:

AWS Certificate Manager (ACM)

Useful for customers who need a secure and public web presence.

ACM Private CA

Useful for customers that are intended for private use within an organization.

Amazon Detective

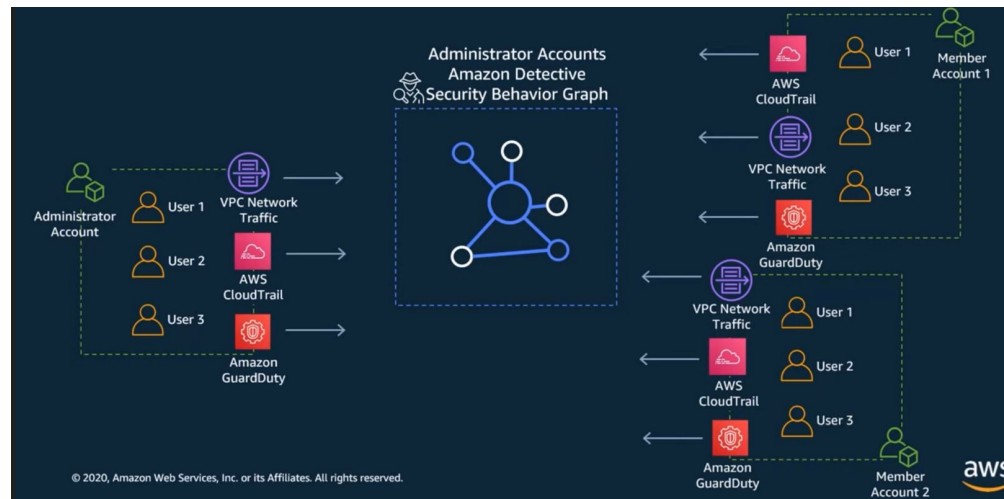
What is Amazon Detective?

Amazon Detective is a service that makes it easy to analyze, investigate & quickly find out the root cause of security findings or suspicious activities within your AWS environment using ML, Statistical analysis, and Graph theory.

It does it by automatically collecting and processing events from VPC Flow logs, CloudTrail, and Amazon GuardDuty to create a unified view.

Features:

- Amazon Detective is a multi-account service. Detective can be enabled in an AWS Organizations Management account.
- Amazon Detective creates a Security Behavior Graph for holding the log summaries & analytics of all member accounts in the Management account.
- The Management account invites member accounts. On accepting the invitations, CloudTrail management events, VPC network traffic, and GuardDuty findings for all accounts flow into the Security Behaviour Graph in the Management account. The management account can then interact with the Security Behaviour Graph in the Detective console.
- The Security Behavioural Graph created by Detective consists of security-related relationships offering contextual & behavioral insights for quickly validating, comparing & correlating data for reaching conclusions.
- Amazon Detective provides interactive visualizations leading to effective investigations using Generative AI. This makes it easy to investigate issues faster with less effort.
- Amazon Detective integrates seamlessly with AWS Security services like Amazon GuardDuty, AWS Security Hub, and Amazon Inspector. Aggregated findings from all these services help quickly investigate security issues identified in these services.
- Amazon Detective is cost-effective since there are no data sources that need to be enabled or configured for using Detective.



AWS Documentation

Amazon Detective

Best Practices:

- It is recommended to use an Administrator Account for Amazon Detective, GuardDuty & Security Hub for the following integration points to work seamlessly
 - Details of GuardDuty findings can be pivoted from the finding details to Amazon Detective's finding profile.
 - While investigating a GuardDuty finding in Amazon Detective, an option to archive the finding can be chosen.
- In order to reduce the amount of time it takes for Detective to receive updates of GuardDuty findings, it is recommended to update the Amazon CloudWatch notification frequency to 15 minutes in GuardDuty rather than its default frequency of 6 hours.

Use cases:

- **Triage Security findings/alerts** - Explore whether GuardDuty findings need to be examined further. Amazon Detective helps users to see whether a finding is a concern.
- **Incident investigation** - Since Amazon Detective allows for viewing analysis & summaries going back up to a year, it can help answer questions like how long has the security issue been there, and the resources affected because of that.
- **Threat Hunting** - Access indicators like IP addresses, users to see what interactions they would have had with the environment. Detective's Security Behaviour Graph will help here.

Exam Tip:

- For Amazon Detective to be enabled, GuardDuty should be enabled for your Account for at least 48 hours.
- For Amazon Detective to be enabled, Volume of data flowing into Amazon Detective's Security Behavior Graph for your account should be less than the maximum allowed by Detective.
- Amazon Detective is a Regional service and needs to be enabled for each Region.

AWS Directory Service

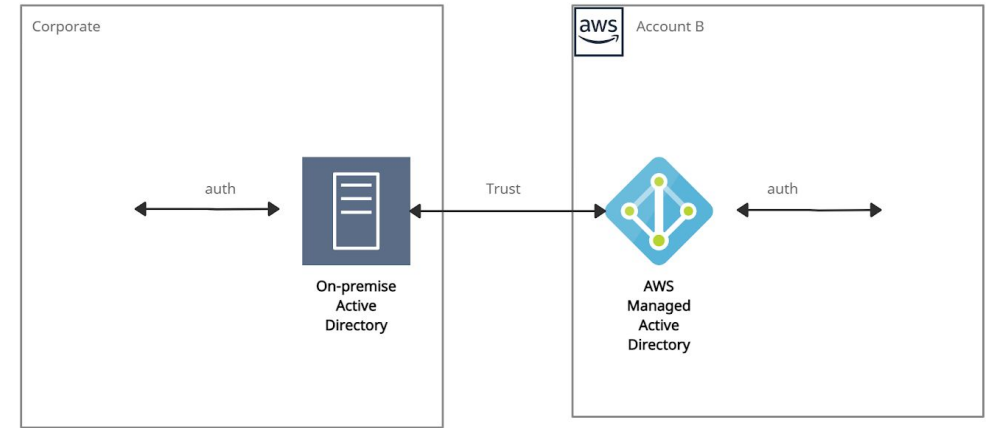
What is AWS Directory Service?

AWS Directory Service, also known as AWS Managed Microsoft Active Directory (AD), enables multiple ways to use Microsoft Active Directory (AD) with other AWS services.



- AWS Directory Service provides the following directory types to choose from:
- Simple AD
 - Amazon Cognito
 - AD Connector

- Trust relationships can be set up from on-premises Active Directories into the AWS cloud to extend authentication.
- It runs on a Windows Server, can perform schema extensions, and works with SharePoint, Microsoft SQL Server, and .Net apps.
- The directory remains available for use during the patching (updating) process for AWS Managed Microsoft AD.
- Using AWS Managed Microsoft AD, it becomes easy to migrate AD-dependent applications and Windows workloads to AWS.
- A trust relationship can be created between AWS Managed Microsoft AD and existing on-premises Microsoft Active using single sign-on (SSO).



AWS Managed AD

- Pricing:**
- Prices vary by region for the directory service.
 - Hourly charges are applied for each additional account to which a directory is shared.
 - Charges are applied per GB for the data transferred “out” to other AWS Regions where the directory is deployed.

Simple AD:

- It is an inexpensive Active Directory-compatible service driven by SAMBA 4.
- It is an isolated or self-supporting AD directory type.
- It can be used when there is a need for less than 5000 users.
- It cannot be joined with on-premise AD.
- It is not compatible with RDS SQL Server.
- It provides some features like
 - Applying Kerberos-based SSO,
 - Assigning Group policies,
 - Managing user accounts and group memberships,
 - Helping in joining a Linux domain or Windows-based EC2 instances.
- It does not support the following functionalities.
 - Multi-factor authentication (MFA),
 - Trust relationships,
 - DNS dynamic update,
 - Schema extensions,
 - Communication over LDAPS,
 - PowerShell AD cmdlets.

Amazon Cognito:

- It is a user directory type that provides sign-up and sign-in for the application using Amazon Cognito User Pools.
- It can create customized fields and store that data in the user directory.
- It helps to federate users from a SAML IdP with Amazon Cognito user pools and provide standard authentication tokens after they authenticate with a SAML IdP (identities from external identity providers).

AD Connector:

- It is like a gateway used for redirecting directory requests to the on-premise Active Directory.
- For this, there must be an existing AD, and VPC must be connected to the on-premise network via VPN or Direct Connect.
- It is compatible with Amazon WorkSpaces, Amazon WorkDocs, Amazon QuickSight, Amazon Chime, Amazon Connect, Amazon WorkMail, and Amazon EC2.
- It is also not compatible with RDS SQL Server.
- It supports multi-factor authentication (MFA) via existing RADIUS-based MFA infrastructure.

Use cases:

- It provides a Sign In option to AWS Cloud Services with AD Credentials.
- It provides Directory Services to AD-Aware Workloads.
- It enables a single-sign-on (SSO) feature to Office 365 and other Cloud applications.
- It helps to extend On-Premises AD to the AWS Cloud by using AD trusts.

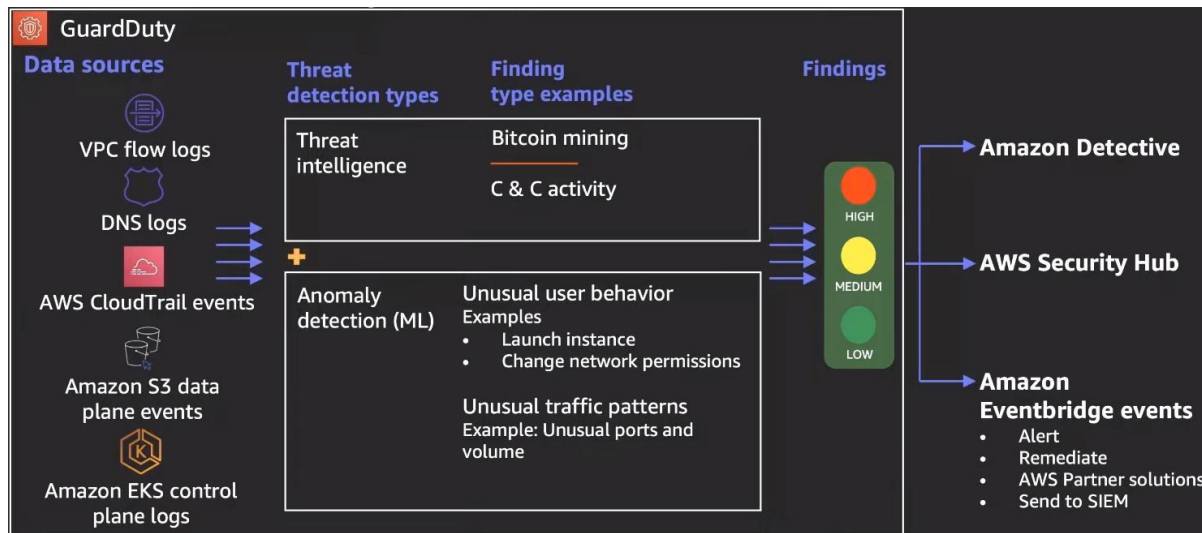
Amazon GuardDuty

What is Amazon GuardDuty?

Amazon GuardDuty is a threat detection service which uses Machine Learning, Anomaly detection, and threat intelligence for identifying and prioritizing potential threats within your AWS accounts & resources. It does it by monitoring different log sources like VPC flow logs, DNS logs, CloudTrail events, S3 data plane events, and EKS control plane logs and analyzing them.

Features:

- Amazon GuardDuty performs threat detection in the following ways
 - Using a Threat Intelligence component by accessing a database of commonly known threats that is maintained by AWS.
 - Anomaly detection using ML for unknown threats (eg Unusual behavior of a role trying to launch EC2 instances)
- Findings once detected by GuardDuty are assigned a severity of Low, Medium or High depending on the level of activity that is happening.
- GuardDuty findings are displayed in the GuardDuty console. They can also be sent to the following services for further analysis
 - AWS Detective.
 - AWS Security Hub.
 - GuardDuty findings can be sent as Amazon EventBridge events for exporting the findings to applications like Splunk, and Sumo Logic or sending event notifications(SNS).
- Multiple accounts can be managed using GuardDuty through AWS Organizations integration.
- The threat detection activity performed by GuardDuty is scalable. Detection capacity is added only when required and reduces utilization when capacity is no longer required.
- GuardDuty provides comprehensive threat protection for containerized workloads. It helps detect malicious & suspicious activity for containerized workloads running on ECS with Fargate, provides container-level context with runtime monitoring, and identifies security coverage gaps in container workloads.



Amazon GuardDuty

Best Practices:

- Ensure that GuardDuty has complete visibility over Logs for complete Detection Coverage - Eg consider enabling VPC Flow logs for all Regions and required network interfaces that are being planned to monitor for threat detection.
- GuardDuty is Region-specific and it is recommended to enable GuardDuty for all Regions for complete threat visibility.
- It is recommended to analyze GuardDuty monitoring activities with CloudTrail to ensure that users are not tampering with GuardDuty itself.
- It is recommended to integrate GuardDuty with EventBridge & Lambda for automating risk mitigation.

Use cases:

- Security analysts can be assisted to carry out investigations using the Security event findings from GuardDuty. It provides Context, Metadata, and impacted resource details using which the root cause can be detected using GuardDuty console integration with Amazon Detective.
- GuardDuty can be used to identify files containing malware - EBS can be scanned for files containing malware that creates suspicious behavior on instance, container workloads running on EC2.

Exam Tip:

When GuardDuty is enabled, the associated log sources that it accesses (VPC Flow logs, DNS Logs) need not be enabled separately. They are all enabled by default by GuardDuty and are provided access to GuardDuty.

You cannot add your own Log sources to GuardDuty other than the 5 mentioned above.

Amazon Identity and Access Management (IAM)

What is Amazon IAM ?

AWS Identity and Access Management is a free service used to define permissions and manage users to access multi-account AWS services.



Amazon Identity and Access Management

Amazon Identity and Access Management allows:

- ❖ users to analyze access and provide MFA (Multi-factor authentication) to protect the AWS environment.
- ❖ managing IAM users, IAM roles, and federated users.

IAM Policies

Policies are documents written in JSON (key-value pairs) used to define permissions.

IAM Users

User can be a person or service.

IAM Groups

Groups are collections of users, and policies are attached to them. It is used to assign permissions to users.

IAM Roles

IAM users or AWS services can assume a role to obtain temporary security credentials to make AWS API calls.

Amazon Inspector

What is Amazon Inspector?

Amazon Inspector is a vulnerability management service which continuously scans AWS resources for software vulnerabilities and network accessibility.

When activated, Amazon Inspector discovers known vulnerabilities in EC2 instances, container images/ECR, Lambda functions and provides a consolidated view of vulnerabilities across compute environments.

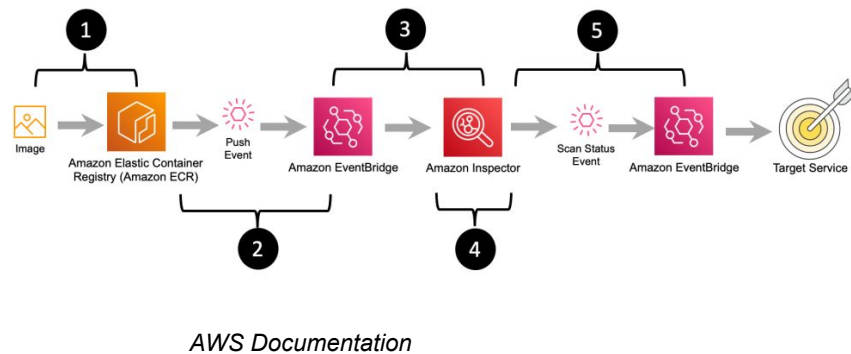
Best Practices:

- From the Management account of AWS Organizations, set up the Delegated Admin (DA) account for Amazon Inspector which will help configure member accounts & consolidate findings.
- Ensure that the AWS Systems Manager Agent is running on EC2 instances since it provides the inventory of softwares & configurations for Amazon Inspector to perform scans & detect vulnerabilities.
- Use Amazon Inspector to scan for security vulnerabilities for your code, and container images that are a part of CI/CD pipelines for build & deployment. This ensures proactive security measures & code remediation early in the software development cycle.

Use cases:

Use Common Vulnerabilities & Exposures (CVE) and network accessibility for creating contextual risk scores to Prioritize Patch remediation.

Support compliance requirements like PCI DSS, NIST CSF and other regulations by utilizing Amazon Inspector scans.



Exam Tip:

- New Amazon Inspector expands its coverage to support container images residing in ECR's in addition to EC2 instances.
- The widely adopted Systems Manager Agent used by Amazon Inspector replaces the standalone Inspector Classic Agent.


Features:

- Automation of vulnerability management - Upon activation, it automatically scans and discovers vulnerabilities in AWS resources like EC2, Lambda functions and container workloads. These vulnerabilities could compromise workloads, target resources for malicious use.
- Amazon Inspector provides multi-account support with AWS Organizations. By assigning an Inspector Delegated Administrator(DA) account for your Organization, it can seamlessly start and configure all member accounts and consolidate all findings.
- Amazon Inspector integrates with AWS Systems Manager Agent for collecting software inventory and configurations from EC2 instances. They are then used to access workloads for vulnerability.
- Findings from Amazon Inspector can be suppressed based on defined criteria. Findings that are deemed by an Organization as acceptable can be suppressed by creating suppression rules.
- A highly contextualized risk score is generated by Amazon Inspector for each finding.
- When a vulnerability has been patched or remediated, Amazon Inspector provides automatic closure of those findings.
- Amazon Inspector provides detailed monitoring of organization-wide environment coverage. It helps to avoid gaps in coverage.
- Amazon Inspector provides integration with AWS Security Hub and EventBridge for its findings. They can be used to automate workflows like Ticketing.
- Amazon Inspector scans Lambda functions for security vulnerabilities like injection flaws, and missing encryption based on AWS best practices. It uses Generative AI and automated reasoning; it provides in-context code remediations for multiple classes of vulnerability reducing the efforts required to fix them.
- Amazon Inspector integrates with CI/CD tools like Jenkins for container image assessments pushing proactive security measures early in the software development cycle.

AWS Key Management Service

What is AWS Key Management Service?


AWS Key Management Service is a global service that creates, stores, and manages encryption keys.



AWS Key Management Service

- ❑ Provides data security at rest using encryption keys and provides access control for encryption, decryption, and re-encryption.
- ❑ Offers SDKs for different languages to add digital signature capability in the application code.
- ❑ AWS KMS produces new cryptographic data for the KMS key once a year, when automatic key rotation is turned on for a KMS key.
- ❑ AWS KMS preserves all previous iterations of the cryptographic information so that you can decrypt any data that has been encrypted using that KMS key. Until the KMS key is deleted, AWS KMS does not remove any rotated key material.

Encryption using AWS KMS



Customer Managed CMKs:

The CMKs created, managed, and used by users are termed as Customer managed CMKs and support cryptographic operations.

AWS Managed CMKs:

The CMKs created, managed, and used by AWS services on the user's behalf are termed AWS-managed CMKs.

AWS Secrets Manager

What is AWS Secrets Manager?

AWS Secrets Manager is a service that prevents secret credentials from being hardcoded in the source code.



AWS Secrets Manager

AWS Secrets Manager:

- Ensures in-transit encryption of the secret between AWS and the system to retrieve the secret.
- Rotates credentials for AWS services using the Lambda function that instructs Secrets Manager to interact with the service or database.
- Stores the encrypted secret value in SecretString or SecretBinary field.
- Uses open-source client components to cache secrets and updates them when there is a need for rotation.

Secrets Manager can be accessed using the following ways:

- AWS Management Console
 - AWS Command Line Tools
 - AWS SDKs
 - HTTPS Query API
-
- It provides security and compliance facilities by rotating secrets safely without the need for code deployment.
 - It integrates with AWS CloudTrail and AWS CloudWatch to log and monitor services for centralized auditing.
 - It integrates with AWS Config and facilitates tracking of changes in Secrets Manager.

Secret rotation is supported with the below Databases:

- MySQL, PostgreSQL, Oracle, MariaDB, Microsoft SQL Server, on Amazon RDS
- Amazon Aurora on Amazon RDS
- Amazon DocumentDB
- Amazon Redshift

AWS Security Hub

What is AWS Security Hub?

AWS Security Hub is a service that provides an extensive view of the security aspects of AWS and helps to protect the environment against security industry standards and best practices.



Benefits:

- It collects data using a standard findings format and reduces the need for time-consuming data conversion efforts.
- Integrated dashboards are provided to show the current security and compliance status.

- ❖ It provides an option to aggregate, organize, and prioritize the security alerts, or findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS IAM Access Analyzer, AWS Firewall Manager, and also from AWS Partner solutions.
- ❖ It helps the Payment Card Industry Data Security Standard (PCI DSS) and the Center for Internet Security (CIS) AWS Foundations Benchmark with a set of security configuration best practices for AWS. If any problem occurs, AWS Security Hub recommends remediation steps.
- ❖ Enabling (or disabling) AWS Security Hub can be quickly done through,
 - AWS Management Console
 - AWS CLI
 - By using Infrastructure-as-Code tools -- Terraform
- ❖ If AWS architecture is divided across multiple regions, it needs to enable Security Hub within each region.
- ❖ The most powerful aspect of using Security Hub is the continuous automated compliance checks using CIS AWS Foundations Benchmark.
- ❖ The CIS AWS Foundations Benchmark consists of 43 best practice checks (such as “Ensure IAM password policy requires at least one uppercase letter” and “Ensure IAM password policy requires at least one number”).

Price details:

- Charges applied for usage of other services that Security Hub interacts with, such as AWS Config items, but not for AWS Config rules that are enabled by Security Hub security standards.
- Using the Master account’s Security Hub, the monthly cost includes the costs associated with all of the member accounts.
- Using a Member account’s Security Hub, the monthly cost is only for the member account.
- Charges are applied only for the current Region, not for all Regions in which Security Hub is enabled.

What is AWS WAF?

AWS WAF is a web application firewall that helps protect web applications from common web exploits and attacks.

It acts as a protective shield for your web applications, helping you safeguard them from threats like SQL injection, cross-site scripting (XSS), and other malicious activities.



Features:

- ❑ **Web Traffic Filtering:** AWS WAF allows you to filter and inspect web traffic coming to your applications. You can set up rules to allow, block, or monitor traffic based on various criteria, such as IP addresses, HTTP headers, request methods, and query strings.
- ❑ **Protection Against Common Attacks:** It protects a wide range of common web attacks, including SQL injection, XSS, and cross-site request forgery (CSRF).
- ❑ **Custom Rules:** You can create custom rules to address specific security requirements and business logic.
- ❑ **AWS WAF Managed Rules for AWS Organizations:** This feature allows you to centrally manage and deploy WAF rules across multiple AWS accounts within an AWS Organization.

Pricing:

AWS WAF costs depend on the quantity of web access control lists (web ACLs) you establish, the number of rules incorporated into each web ACL, and the volume of web requests you receive.

No prior commitments are required. It's important to note that AWS WAF charges are separate from Amazon CloudFront pricing, AWS Cognito pricing, Application Load Balancer (ALB) pricing, Amazon API Gateway pricing, and AWS AppSync pricing.

Best Practices:

- ❖ Combine AWS WAF with other AWS services such as AWS Shield (for DDoS protection) and Amazon CloudFront (for content delivery) to create a robust, multi-layered security strategy.
- ❖ If you're using AWS Managed Rule Sets, ensure that you keep them up to date. AWS regularly updates these rule sets to protect against emerging threats.
- ❖ Enable logging for AWS WAF to capture detailed information about web requests and potential threats. Use Amazon CloudWatch or a SIEM solution to monitor and analyze these logs.
- ❖ Implement rate-limiting rules to protect APIs from abuse and DDoS attacks. Set appropriate rate limits based on expected traffic patterns.
- ❖ Tailor your web access control lists (web ACLs) to the specific needs of your application.
- ❖ Periodically review your AWS WAF rules to make adjustments based on changing application requirements and emerging threats.



Storage

AWS Backup



What is AWS Backup?

AWS Backup is a secure service that automates and governs data backup (protection) in the AWS cloud and on-premises.

Price details:

AWS charges monthly based on the amount of backup storage used and the amount of backup data restored.

Features:

- It offers a backup console, backup APIs, and the AWS Command Line Interface (AWS CLI) to manage backups across AWS resources like instances and databases.
- It offers backup functionalities based on policies, tags, and resources.
- It provides scheduled backup plans (policies) to automate backup of AWS resources across AWS accounts and regions.
- It offers incremental backup to minimize storage costs. The first backup backs up a full copy of the data and then only the successive incremental backup changes.
- It provides backup retention plans to retain and expire backups automatically. Automated backup retention also helps to minimize storage costs for backup.
- It provides a dashboard in the AWS Backup console to monitor backup and restore activities.
- It offers an enhanced solution by providing separate encryption keys for encrypting multiple AWS resources.

AWS Backup

Features:

- It provides lifecycle policies configured to transition backups from Amazon EFS to cold storage automatically.
- It is tightly integrated with Amazon EC2 to schedule backup jobs and the storage (EBS) layer. It also simplifies recovery by restoring whole EC2 instances from a single point.
- It supports cross-account backup and restores either manually or automatically within the AWS organizations.
- It allows backups and restores to different regions, especially during any disaster, to reduce downtime and maintain business continuity.
- It integrates with Amazon CloudWatch, AWS CloudTrail, and Amazon SNS to monitor, audit API activities and notifications.




Use cases:

- It can use AWS Storage Gateway volumes for hybrid storage backup.
- AWS Storage Gateway volumes are secure and compatible with Amazon EBS, which helps restore volumes to on-premises or the AWS environment.

Amazon Elastic Block Store

What is Amazon Elastic Block Store?

Amazon Elastic Block Store is a service that provides the block-level storage drive to store persistent data.



- ❖ Multiple EBS volumes can be attached to a single EC2 instance in the same availability zone.
- ❖ A single EBS volume can not be attached to multiple EC2 instances.
- ❖ Amazon EBS Multi-Attach is a feature used to attach a single Provisioned IOPS SSD (io1 or io2) volume to multiple instances in the same Availability Zone.
- ❖ EBS volumes persist independently after getting attached to an instance, which means the data will not be erased even if it terminates.
- ❖ By default, the root EBS volume gets terminated when the instance is terminated.

By default, the non-root EBS volume does not get affected when the instance is terminated.

Amazon EBS can be attached and detached to an instance and can be reattached to other EC2 instances.

Amazon EBS easily scales up to petabytes of data storage.

Amazon EBS volumes are best suited for database servers with high reads and write and throughput-intensive workloads with continuous reads and write.

Amazon EBS uses AWS KMS service with AES-256 algorithm to support encryption.

Amazon EBS offers point-in-time snapshots for volumes to migrate to other AZs or regions.

EBS snapshots are region-specific and are incrementally stored in Amazon S3.

Amazon Elastic Block Store

EBS volumes types are as follows:

SSD (Solid-state drives)

General Purpose SSD:

- Useful for low-latency applications, development, and test environments.
- Supports volume size from 1 GiB to 16 TiB.
- Allows 16,000 as maximum IOPS per volume.
- Allows 1000 MiB/s as maximum throughput per volume.

Provisioned IOPS SSD:

- Useful for I/O-intensive database workloads and provide sub-millisecond latency.
- Supports volume size from 4 GiB to 64 TiB.
- Allows 256,000 as maximum IOPS per volume.
- Allows 4,000 MiB/s as maximum throughput per volume.
- The multi-Attach feature is supported for io1 and io2

HDD (Hard disk drives)

Throughput Optimized HDD:

- Useful for Big data and Log processing workloads.
- Supports volume size from 125 GiB to 16 TiB.
- Allows 500 as maximum IOPS per volume.
- Allows 500 MiB/s as maximum throughput per volume.

Cold HDD:

- Useful for infrequently accessed data and lowest cost workloads.
- Supports volume size from 125 GiB to 16 TiB.
- Allows 250 as maximum IOPS per volume.
- Allows 250 MiB/s as maximum throughput per volume.

Amazon Elastic File System (EFS)

What is Amazon Elastic File System?

Amazon Elastic File System is a managed service used to create and scale file storage systems for AWS and on-premises resources.



Amazon EFS

It offers the following storage classes for file storage:

- EFS Standard storage class
- EFS Infrequent Access storage class - can store less frequently accessed files.

It offers the following modes to ease the file storage system:

- ❑ Performance modes -
 - ❑ General Purpose performance mode: Useful for low-latency workloads.
 - ❑ Max I/O mode: High throughput workloads.
- ❑ Throughput modes -
 - ❑ Bursting Throughput mode: Throughput increases based on the file system storage.
 - ❑ Provisioned Throughput mode: Throughput changes are independent of the file system storage.
- ❑ It provides EFS lifecycle management policies based on the number of days ranges from 7-90 days to automatically move files from Standard storage class to EFS IA storage class.

- ❑ It spans multiple availability zones and regions.
- ❑ It uses EFS Mount Target to share a file system with multiple availability zones and VPCs.
- ❑ It is best suited for Linux-based workloads and applications.
- ❑ Multiple instances can access it at the same time leads to high throughput and low latency IOPS.
- ❑ It automatically scales storage capacity up to petabyte.
- ❑ It supports file locking and strong data consistency.
- ❑ It offers data encryption at rest and in-transit using AWS KMS and TLS, respectively.
- ❑ It uses POSIX permissions to control access to files and directories.

Amazon Simple Storage Service (S3)

What is Amazon Simple Storage Service?

Amazon S3 is a simple service used to provide key-based object storage across multiple availability zones (AZs) in a specific region.

- S3 is a global service with region-specific buckets.
- It is also termed a static website hosting service.
- It provides 99.99999999% (11 9's) of content durability.
- S3 offers strong read-after-write consistency for any object.
- Objects (files) are stored in a region-specific container known as Bucket.
- Objects that are stored can range from 0 bytes - 5TB.

- It provides 'Multipart upload' features that upload objects in parts, suitable for 100 MB or larger objects.
- It offers to choose 'Versioning' features to retain multiple versions of objects, must enable versioning at both source and destination.
- Amazon S3 Transfer Acceleration allows fast and secure transfer of objects over long distances with minimum latency using Amazon CloudFront's Edge Locations.
- Amazon S3 uses access control lists (ACL) to control access to the objects and buckets.
- Amazon S3 provides Cross-Account access to the objects and buckets by assuming a role with specified privileges.



Amazon S3

Amazon S3 uses the following ways for security:

User-based security

- IAM policies

Resource-Based

- Bucket Policies
- Bucket Access Control List (ACL)
- Object Access Control List (ACL)

Amazon S3 provides the following storage classes used to maintain the integrity of the objects:

- S3 Standard** - offers frequent data access.
- S3 Intelligent-Tiering** - automatically transfer data to other cost-effective access tiers.
- S3 Standard-IA** - offers immediate and infrequent data access.
- S3 One Zone-IA** - infrequent data access.
- S3 Glacier** - long-term archive data, cheap data retrieval.
- S3 Glacier Deep Archive** - used for long-term retention.

Amazon S3 offers to choose from the following ways to replicate objects:

- Cross-Region Replication - used to replicate objects in different AWS Regions.
- Same Region Replication - used to replicate objects in the same AWS Region.

Amazon S3 Glacier

What is Amazon S3 Glacier?

Amazon S3 Glacier is a web service with vaults that offer long-term data archiving and data backup.



It is the cheapest S3 storage class and offers 99.999999999% of data durability.

S3 Glacier provides the following data retrieval options:

Expedited retrievals -

- It retrieves data in 1-5 minutes.

Standard retrievals -

- It retrieves data between 3-5 hours.

Bulk retrievals -

- It retrieves data between 5-12 hours.

S3-Standard, S3 Standard-IA, and S3 Glacier storage classes, objects, or data are automatically stored across availability zones in a specific region.

A vault is a place for storing archives with a unique address.

Amazon S3 Glacier jobs are the select queries that execute to retrieve archived data. It uses Amazon SNS to notify when the jobs complete.

Amazon S3 Glacier does not provide real-time data retrieval of the archives.

Amazon S3 Glacier uses 'S3 Glacier Select' to query archive objects in uncompressed CSV format and store the output to the S3 bucket.

Amazon S3 Glacier Select uses common SQL statements like SELECT, FROM, and WHERE.

It offers only SSE-KMS and SSE-S3 encryption.

AWS Storage Gateway

What is the AWS Storage Gateway?

AWS Storage Gateway is a hybrid cloud storage service that allows your on-premise storage & IT infrastructure to integrate with AWS Cloud Storage Services seamlessly. It Can be AWS Provided Hardware or a Compatible Virtual Machine.



❖ Purpose of Using AWS Storage Gateway(hybrid Cloud Storage):

- To Fulfill Licencing Requirements.
- To Achieve Data-Compliance Requirements.
- To Reduce Storage & Management Cost.
- For Easy and Effective Application Storage-Lifecycle & Backup Automation.
- For Hybrid Cloud & Easy Cloud Migration.

Use Cases:

- Cost-Effective Backups and Disaster Recovery Management
- Migration to/from Cloud
- Managed Cache: Integration of Local(on-premises) Storage to Cloud Storage (Hybrid Cloud)
- To Achieve Low Latency by storing data on-premise and still leverage cloud benefits

Pricing:

- Charges are applied on what you use with the AWS Storage Gateway and based on the type and amount of storage you use.

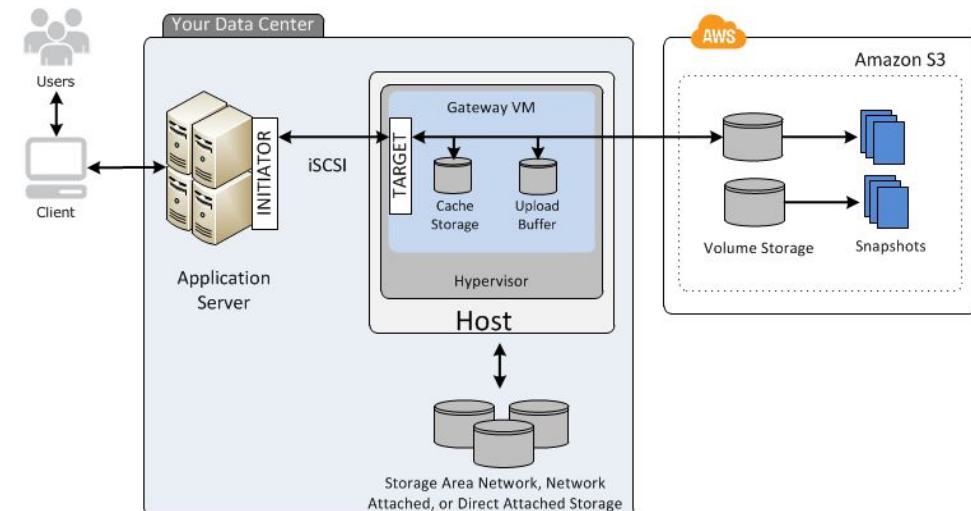
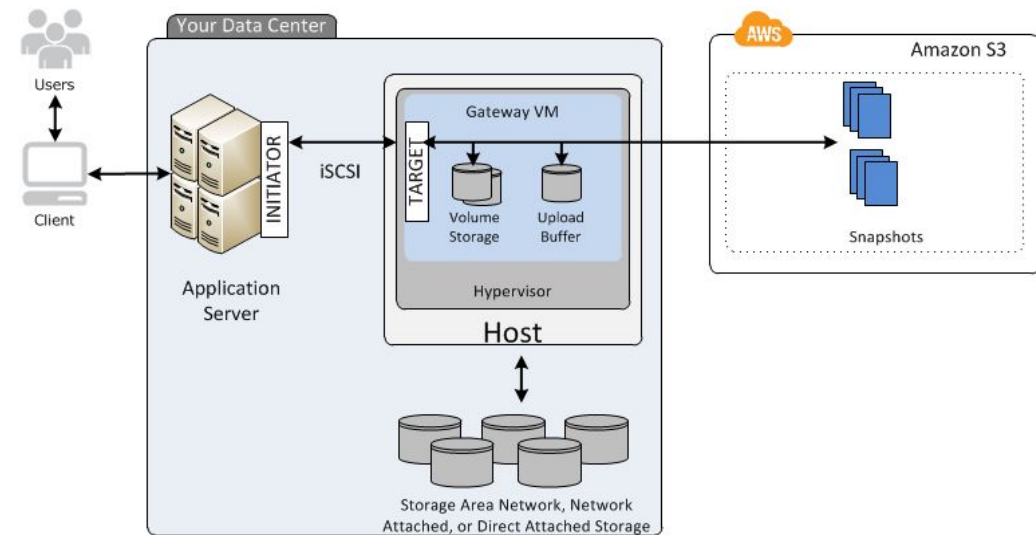
❖ Features of AWS Storage Gateway

- Cost-Effective Storage Management
- To achieve Low Latency on-premise.
- Greater Control over Data still take advantage of the cloud (Hybrid Cloud)
- Compatible and Compliance
- To meets license requirement
- Supports both hardware and software gateway
- Easy on-premise to Cloud Migrations
- Standard Protocol for storage access like NFS/SMB/iSCSI

Types of Storage Gateway

1. Volume Gateway (iSCSI)


- To Access Virtual Block-Level Storage Stored on-premise
- It can be asynchronously backed up and stored as a snapshot on AWS S3 for high reliability & durability.
 - **Storage Volume Gateway:** All Applications Data Stored on-premise and the only backup is stored on AWS S3.
 - **Cache Volume Gateway:** Only Hot Data / Cached data is Stored on-premise and all other application data is stored on AWS S3.



Types of Storage Gateway

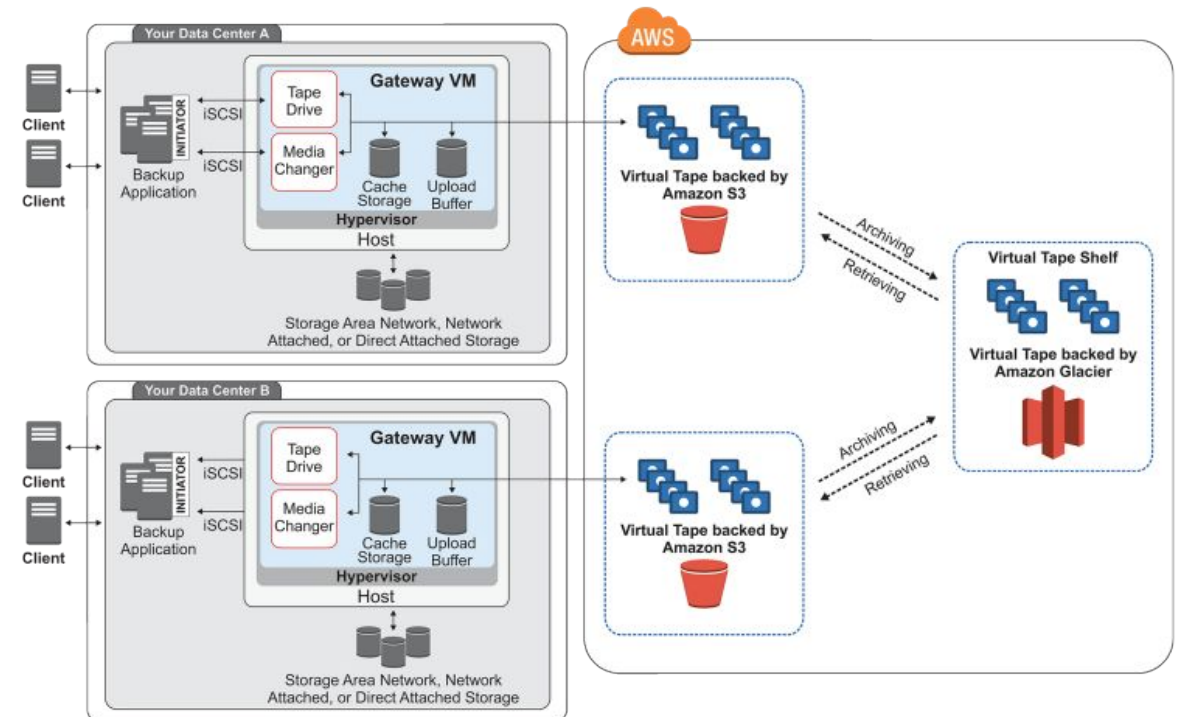
File Gateway (NFSv4 / SMB)

- To Access Object-based Storage (AWS S3 Service)
- Supports NFS Mount Point for accessing S3 Storage to the local system as Virtual Local File System
- Leverage the benefits of AWS S3 Storage Service



Tape Gateway (VTL)

- It is virtual local tape storage.
- It uses the Virtual Tape Library(VTL) by iSCSI protocol.
- It is cost-effective archive storage (AWS S3) for cloud backup.





Management and Governance

AWS CloudFormation

What is AWS CloudFormation?

AWS CloudFormation is a service that collects AWS and third-party resources and manages them throughout their lifecycles by launching them together as a stack.



AWS CloudFormation

Template:

- ❑ A **template** is used to create, update, and delete an entire stack as a single unit without managing resources individually.
- ❑ CloudFormation provides the capability to reuse the template to set the resources easily and repeatedly.

Stacks:

- ❑ **Stacks** can be created using the AWS CloudFormation console and AWS Command Line Interface (CLI).
- ❑ **Nested Stacks** are stacks created within another stack by using the 'AWS::CloudFormation::Stack' resource attribute.
- ❑ The main stack is termed as parent stack, and other belonging stacks are termed as child stack, which can be implemented by using ref variable '! Ref'.

Example: CloudFormation template for creating EC2 instance

```
EC2Instance:
  Type: AWS::EC2::Instance
  Properties:
    ImageId: 1234xyz
    KeyName: aws-keypair
    InstanceType: t2.micro
    SecurityGroups:
      - !Ref EC2SecurityGroup
    BlockDeviceMappings:
      - DeviceName: /dev/sda1
    Ebs:
      VolumeSize: 50
```

AWS does not charge for using AWS CloudFormation, and charges are applied for the CloudFormation template services.



AWS CloudTrail

What is AWS CloudTrail?

AWS CloudTrail is a service that gets enabled when the AWS account is created and is used to enable compliance and auditing of the AWS account.



- ✓ It offers to view, analyze, and respond to activity across the AWS infrastructure.
- ✓ It records actions as an event by an IAM user, role, or an AWS service.
- ✓ CloudTrail records can download Cloud Trail events in JSON or CSV file.
- ✓ **CloudWatch** monitors and manages the activity of AWS services and resources, reporting on their health and performance. Whereas **CloudTrail** resembles logs of all actions performed inside the AWS environment.
- ✓ **IAM log file** -
The below example shows that the IAM user Rohit used the AWS Management Console to call the AddUserToGroup action to add Nayan to the administrator group.

```
Records": [{
  "eventVersion": "1.0",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "PR_ID",
    "arn":
"arn:aws:iam::210123456789:user/Rohit",
    "accountId": "210123456789",
    "accessKeyId": "KEY_ID",
    "userName": "Rohit"
  },
  "eventTime": "2021-01-24T21:18:50Z",
  "eventSource": "iam.amazonaws.com",
  "eventName": "CreateUser",
  "awsRegion": "ap-south-2",
  "sourceIPAddress": "176.1.0.1",
  "userAgent": "aws-cli/1.3.2 Python/2.7.5
Windows/7",
  "requestParameters": {"userName": "Nayan"},
  "responseElements": {"user": {
    "createDate": "Jan 24, 2021 9:18:50 PM",
    "userName": "Nayan",
    "arn": "arn:aws:iam::128x:user/Nayan",
    "path": "/",
    "userId": "12xyz"
  }}
}]}
```

Amazon CloudWatch

What is Amazon CloudWatch?

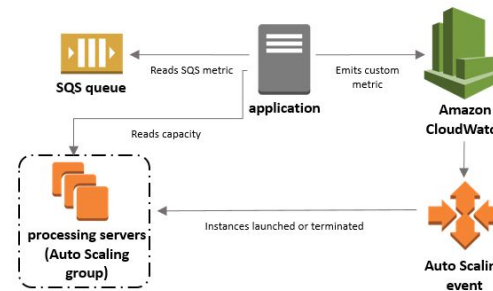
Amazon CloudWatch is a near real-time monitoring tool used to track the behavior of AWS resources and applications that run on them.

The behavior of AWS resources relates to its metrics like CPU Utilization, Memory Utilization, and the behavior of applications related to the application's logs.

On capturing different metrics and logs, actions can be taken by creating CloudWatch Alarms like AutoScaling, triggering SNS notifications for the observed behavior. Application logs can be captured by installing and running a CloudWatch Agent on AWS resources. They appear as CloudWatch Logs in the CloudWatch AWS console.

Best Practices:

- Identify monitoring needs for your AWS resources and applications. CloudWatch Metrics, Logs, Filters, Alarms, and Dashboards can then be appropriately configured based on the monitoring requirements.
- Install CloudWatch Agent for exporting Custom Metrics, and application logs to CloudWatch metrics & CloudWatch Logs respectively.
- Use CloudWatch "Metric Filters" for defining actions on CloudWatch Log streams.
- Ensure that IAM permissions are provided with AWS least privilege access. principles for AWS resources & applications that output monitoring data to CloudWatch and applications that retrieve CloudWatch metrics for performing analysis.
- Export CloudWatch Log data to S3 for performing Detailed Analytics using tools like Amazon Athena.



AWS Documentation

Features:

- The AWS resources that can be monitored using Amazon CloudWatch include EC2 instances, Amazon EBS, Elastic Load Balancers, Amazon RDS, Amazon DynamoDB, and Amazon S3.
- Default metrics like CPU Utilization, Network In-Out, and Request counts are provided automatically for AWS resources.
- Metrics like Memory Usage, Transaction volumes, and application metrics are not provided by default. They need to be provided as Custom Metrics that will also be monitored by Amazon CloudWatch.
- Using CloudWatch dashboards, statistics of CloudWatch metrics upto-the-minute can be visualized using various graphs.
- Metric filters can be created for CloudWatch Logs based on the streaming Log patterns. Alarms that perform actions like triggering a SNS notification can be created from these Metric Filters.
- CloudWatch Logs Insights allows for quickly searching and analyzing logs data using simple, powerful queries.
- Applications and services that send logs to CloudWatch need to have the appropriate IAM permissions to do so.
- Amazon CloudWatch is accessible through AWS Management Console, Amazon SDK, and command line tools.

Use cases:

- Improve Performance of AWS services: Using CloudWatch metrics, one can monitor resource utilization of different AWS services and trigger automated actions eg Auto Scaling EC2 instances for improving performance.
- Perform Analysis of CloudWatch Metrics: Products like AWS Trusted Advisor can use CloudWatch metrics and analyze them to give an idea of whether the performance or spending of your AWS resources can be optimized with respect to the usage profile.
- Error detection & troubleshooting: Using CloudWatch Logs, application errors, exceptions, and HTTP responses can be detected (eg HTTP error codes 4XX, 5XX), and appropriate actions taken like sending notifications to alert users.

Exam Tip:

- Two types of Monitoring supported by CloudWatch
 - Basic Monitoring - EC2 metrics are sent every 5 minutes by default (free).
 - Detailed Monitoring - EC2 monitoring metrics are sent every 1 minute (chargeable).
- CloudWatch supports two types of alarms
 - Metric Alarm - Performs one or more actions based on a single metric.
 - Composite Alarm - Uses rule expression that takes into account multiple alarms.
- Following are the Alarm States in CloudWatch
 - **OK** - The defined metric is within the threshold.
 - **ALARM** - Metric is outside the threshold.
 - **INSUFFICIENT_DATA** - Not enough data.

Amazon CloudWatch Logs

What is Amazon CloudWatch Logs?

Amazon CloudWatch Logs is a service provided by Amazon Web Services (AWS) that enables you to monitor, store, and access log data from various AWS resources and applications. It is designed to help you centralize and gain insights from logs generated by your AWS resources, applications, and services in a scalable and cost-effective manner.

Use Cases:

- **Application Debugging:** Developers want to troubleshoot and debug issues in a microservices-based application.
- **Cost Monitoring for EC2 Instances:** An organization wants to track and control costs associated with their Amazon EC2 instances.
- **Security and Compliance Auditing:** A company needs to monitor and audit user activities across its AWS environment to ensure compliance with security policies.

→ Features

- ❖ **Log Collection:** CloudWatch Logs allows you to collect log data from a wide range of AWS resources and services, including Amazon EC2 instances, Lambda functions, AWS CloudTrail, AWS Elastic Beanstalk, and custom applications running on AWS or on-premises.
- ❖ **Log Storage:** It provides a secure and durable repository for your log data.
- ❖ **Real-time Monitoring:** You can set up CloudWatch Alarms to monitor log data in real time and trigger notifications or automated actions when specific log events or patterns are detected.
- ❖ **Log Queries:** CloudWatch Logs Insights allows you to run ad-hoc queries on your log data to extract valuable information and troubleshoot issues. You can use a simple query language to filter and analyze logs.
- ❖ **Log Retention:** You can define retention policies for your log data, specifying how long you want to retain logs before they are automatically archived or deleted. This helps in cost management and compliance with data retention policies.
- ❖ **Log Streams:** Within a log group, log data is organized into log streams, which represent individual sources of log data. This organization makes it easy to distinguish between different sources of log data.

Amazon CloudWatch Logs

Pricing

Amazon CloudWatch operates on a pay-as-you-go model, meaning there are no initial obligations or minimum charges. You are billed based on your actual usage, with charges calculated and billed at the conclusion of each month.

CloudWatch Logs offer two distinct tiers namely Free Tier and Paid Tier.

Best Practices:

- **Log Structure:** Design your log messages with a consistent and meaningful structure. Use JSON or key-value pairs to include relevant information such as timestamps, log levels, and context.
- **Log Events Filtering:** Use log event filters to extract valuable information from log data. You can create metric filters or use CloudWatch Logs Insights for more advanced log queries.
- **Use CloudWatch Metrics:** Integrate CloudWatch Logs with CloudWatch Metrics to gain insights into log data trends and visualize log-related metrics in CloudWatch Dashboards.

→ Limitations

- ❖ **Query Execution Time:** When using CloudWatch Logs Insights to query log data, there is a maximum query execution time limit. Complex queries or queries over a large dataset may time out.
- ❖ **Data Exports:** While you can export log data to Amazon S3 or other destinations, there are limitations on the frequency of exports and the destinations you can use.
- ❖ **Data Structure:** CloudWatch Logs is designed primarily for unstructured log data. If you have structured data, you may need to parse it using CloudWatch Logs Insights to make it more accessible.
- ❖ **API Limitations:** The CloudWatch Logs API has rate limits on the number of requests you can make, and these limits vary based on your AWS account and region.
- ❖ **Log Streams:** Each log stream within a log group must have a unique name, which can make it challenging to manage log streams for resources that generate a large number of logs.
- ❖ **Log Data Size:** There are limits on the size of individual log events and log batches, which may require you to segment and compress log data if it exceeds these limits.

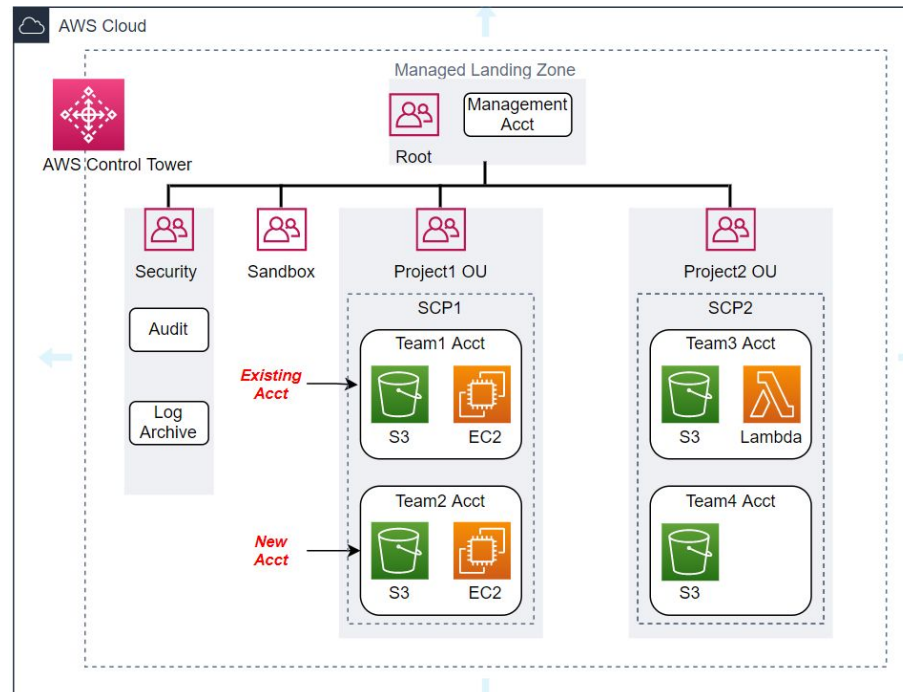
AWS Control Tower

What is AWS Control Tower?

AWS Control Tower is an extension to AWS Organizations providing additional controls.

AWS Control Tower helps create a Landing Zone which is a well-architected Multi-Account baseline based on AWS best practices.

An AWS Organization will be created if it does not already exist.



AWS Documentation

Features:

- As a part of the Landing Zone, Control Tower sets up a series of OUs - Security OU, Sandbox OU, and Production OU.
- Within the Security OU, the Control Tower creates the Audit & Log Archive accounts.
- The Sandbox & Production OUs does not contain any default accounts. Accounts related to Development and Production environments can be added to these OUs.
- Control Tower integrates with AWS Identity Center. The directory sources for SSO can be AWS Identity Center directories(default), SAML IdPs, and Microsoft AD.
- The Root user in the Management Account can perform actions that are disallowed by Guardrails similar to AWS Organizations where SCPs cannot affect the Root user in the Management Account.
- Control Tower comes with a Dashboard providing oversight into the Landing Zone and central administrative views across all Accounts, OUs, Guardrails & policies.
- Control Tower offers Account Factory which is a configurable Account Template for standardizing provisioning of new Accounts with Pre-approved Account configurations.

Exam Tip:

- AWS Control Tower provides two configuration options
 - Launch AWS Control Tower in a new AWS Organization.
 - Launch AWS Control Tower in an existing AWS Organization.
- Guardrails created by AWS Control Tower for governance & compliance fall under the following categories
 - Preventive Guardrails - These are based on SCPs that disallow certain API actions.
 - Detective Guardrails - Implemented using AWS Config & Lambda functions that monitor & govern compliance.

Use cases:

- Quick deployment of applications in a multi-account environment.
- Supporting Digital Sovereignty for Data residency, granular access restriction, encryption & resiliency by deploying Managed Controls.
- Provisioning of compliant AWS accounts for meeting business, security & compliance requirements.

Best Practices:

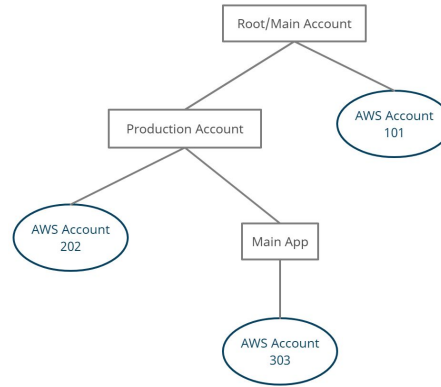
- Since Control Tower provides high levels of customizations for creating a Multi-Account environment, it is important to understand and access your Organization's OUs & workloads for correctly scoping & applying the right controls.
- AWS Control Tower uses compliance frameworks like PCI-DSS, CIS AWS Benchmark thus enabling controls for achieving specific compliance objectives for various Organizations. Organizations should align to these IT compliance frameworks which will offer a consistent & repeatable foundation for risk management and security configuration best practices in their AWS environment.
- It is advisable to test all Control Tower controls in non-production OUs for identifying & mitigating potential risks of misconfigurations early prior to actual production deployments.
- Use automation for detection & remediation of non-compliant controls by leveraging the synergy between AWS Control Tower controls and AWS Systems Manager.

AWS Organizations

What are AWS Organizations?

AWS Organizations is a global service that enables users to consolidate and manage multiple AWS accounts into an organization.

It includes account management and combined billing capabilities that help to meet the budgetary, and security needs of the business better.



AWS Organizations flow

Price details:

- AWS Organizations is free. Charges are applied to the usage of other AWS resources.
- The management account is responsible for paying charges of all resources used by the accounts in the organization.
- AWS Organizations provides consolidated billing that combines the usage of resources from all accounts, and AWS allocates each member account a portion of the overall volume discount based on the account's usage.

AWS Organizations can be accessed in the following ways:

- AWS Management Console
- AWS Command Line Tools
 - AWS Command Line Interface (AWS CLI)
 - AWS Tools for Windows PowerShell.
- AWS SDKs
- AWS Organizations HTTPS Query API

Features:


- AWS Organizations provides security boundaries using multiple member accounts.
- It makes it easy to share critical common resources across the accounts.
- It organizes accounts into organizational units (OUs), which are groups of accounts that serve specified applications.
- Service Control Policies (SCPs) can be created to provide governance boundaries for the OUs. SCPs ensure that users in the accounts only perform actions that meet security requirements.
- Cost allocation tags can be used in individual AWS accounts to categorize and track the AWS costs.
- It integrates with the following services:
 - AWS CloudTrail - Manages auditing and logs all events from accounts.
 - AWS Backup - Monitor backup requirements.
 - AWS Control Tower - to establish cross-account security audits and view policies applied across accounts.
 - Amazon GuardDuty - Managed security services, such as detecting threats.
 - AWS Resource Access Manager (RAM) - Can reduce resource duplication by sharing critical resources within the organization.
- **Steps to be followed for migrating a member account:**
 - Remove the member account from the old Organization.
 - Send an invitation to the member account from the new Organization.
 - Accept the invitation to the new Organization from the member account.

AWS Systems Manager

What is AWS Systems Manager?

AWS Systems Manager (SSM) is a service that allows users to centralize or group operational data using multiple services and automate operations across AWS infrastructure.

- ✓ It simplifies maintenance and identifies issues in the resources that may impact the applications.
- ✓ It displays the operational data, system and application configurations, software installations, and other details on a single dashboard known as AWS Systems Manager Explorer.
- ✓ It manages secrets and configuration data and separates them from code using a centralized store known as Parameter Store.
- ✓ It helps to communicate with the Systems Manager agent installed on AWS servers and in an on-premises environment. Agents are installed to manage resources on servers using different operating systems.

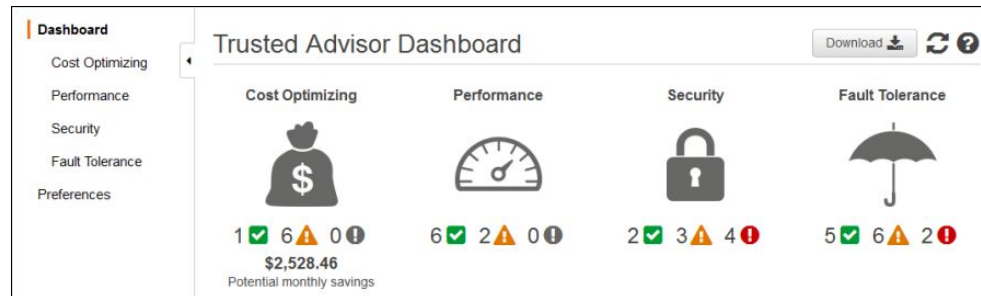


- It helps to manage servers without actually logging into the server using a web console known as Session Manager.
- It helps to automate repetitive operations and management tasks using predefined playbooks.
- It connects with Jira Service Desk and ServiceNow to allow ITSM platform users to manage AWS resources.
- Systems Manager Distributor helps to distribute software packages on hosts along with versioning.

What is AWS Trusted Advisor?

AWS Trusted Advisor is a service that provides real-time guidance to organizations using AWS Cloud to provision resources that follow AWS best practices.

As a Global service, Trusted Advisor inspects your AWS infrastructure comparing it with AWS best practices, and provides recommended actions for optimizing your infrastructure.



AWS Documentation

Features:

- Trusted Advisor performs specific checks based on best practices identified by experts in each AWS service and learnings from serving customers.
- The Trusted Advisor Dashboard provides a category-level summary of check results.
- For each category, the summary includes an aggregation of Check Status sorted by OK (Green), Warning (Yellow), and Error (Red) across all Regions. Each category can be further drilled down or downloaded as an Excel sheet for viewing Alerts and recommended actions.
- Trusted Advisor adds new checks, and new features, updates existing checks, and expands to new Regions on an ongoing basis.
- Drill down results for specific resources, check statuses can be obtained by using a filter tag. Eg a specific Application, Business Unit.
- Trusted Advisor comes with the following types of plans
 - Basic & Developer plans.
 - Business and Enterprise support plans.

Exam Tip:

- The Basic & Developer support plans have Limited access to Trusted Advisor checks(7 checks) accessible through AWS Console only.
- The Business & Enterprise support plans have access to all Trusted Advisor checks through the console, access checks via API calls, and setup of EventBridge to consume check results that can be sent as notifications to interested parties.

Best Practices:

- Trusted Advisor itself provides checks based on Best Practices in the Cost Optimization, Security, Fault Tolerance, and Performance improvement categories.
 - **Cost Optimization** - Provides recommendations to Organizations for saving money on their AWS infrastructure by terminating unused & idle resources, using Reserved capacity for continuous usage.
 - **Security** - Provides recommendations to Organizations for improving the security of their applications by Restricting access using SG/NACL, Checking permissions on S3 Buckets, and enabling various security features.
 - **Fault Tolerance** - Increasing availability and redundancy of applications using Auto Scaling, Performing Health checks, configuring Multi-AZ environments, and taking backups.
 - **Performance** - Provides recommendations to Organizations for improving the performance of applications by taking advantage of provisioned throughput, and monitoring of over-utilized instances
 - **Service Limits** - Notifies Organizations when their resource usage is more than 80%.

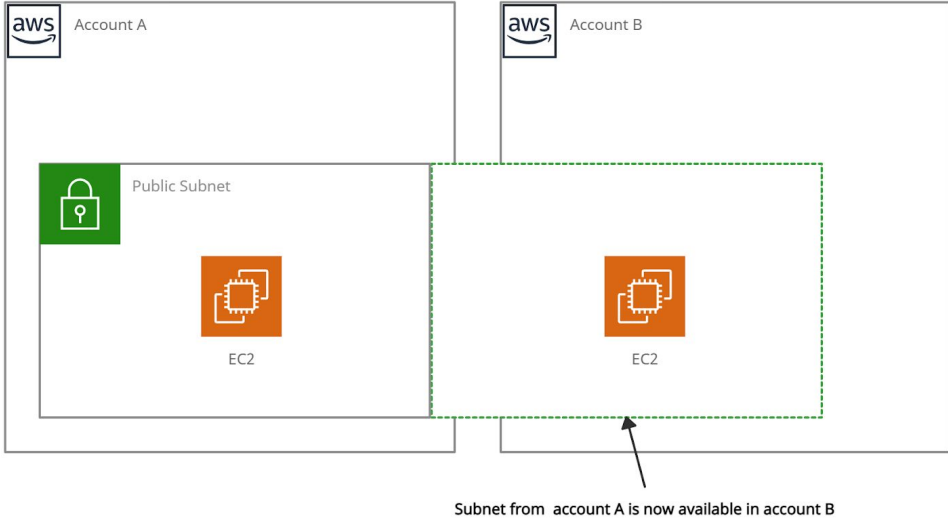
Use cases:

- **Optimization of cost & efficiency** - Trusted Advisor helps identify resources that are not used to capacity or idle resources and provides recommendations to lower costs.
- **Address Security Gaps** - Trusted Advisor performs Security checks of your AWS environment based on security best practices. It flags off errors or warnings depending on the severity of the security threat e.g. Open SG/NACL ports for unrestricted external user access, and open access permissions for S3 buckets in Accounts.
- **Performance Improvement** - Trusted Advisor checks for usage & configuration of your AWS resources and provides recommendations that can improve performance e.g. it can check for Provisioned IOPS EBS volumes on EC2 instances that are not EBS-optimized.

AWS Resource Access Manager

What is AWS Resource Access Manager?

AWS Resource Access Manager (RAM) is a service that permits users to share their resources across AWS accounts or within their AWS Organization.



AWS Resource Access Manager

Resources that can be integrated with AWS RAM are:

- AWS App Mesh
- Amazon Aurora
- AWS Certificate Manager Private Certificate Authority
- AWS CodeBuild
- EC2 Image Builder
- AWS Glue
- AWS License Manager
- AWS Network Firewall
- AWS Outposts
- AWS Resource Groups

- Benefits:**
- The resource sharing feature of AWS RAM reduces customers' need to create duplicate resources in each of their accounts.
 - It controls the consumption of shared resources using existing policies and permissions.
 - It can be integrated with Amazon CloudWatch and AWS CloudTrail to provide detailed visibility into shared resources and accounts.
 - Access control policies in AWS Identity & Access Management (IAM) and Service Control Policies in AWS Organizations provide security and governance controls to AWS Resource Access Manager (RAM).

Price details:

The charges only differ based on the resource type. No charges are applied for creating resource shares and sharing your resources across accounts.