AWS Architecting Associate Exam Tips

Updated September 2019

IAM

Know that Deny always overrides Allow, and that permissions have to be explicity allowed.

Know that it's a global service.

Know the different types of credentials eg a user needs a Username and Password in order to access the console, and Access and Secret keys in order to use CLI or make API calls using the SDKs.

Know what access permissions an IAM user has by default. Hint: None.

Know the use case for roles eg allows a resource to access a service eg EC2 to access S3/Dynamo, Cross Account access, SAML and Web Identify Federation.

Know that you can assign a role to a running instance which does not currently have a role.

S3/Glacier

Know that you can store an "unlimited" amount of data in S3.

The maximum size of an individual object is 5TB.

Know that S3 objects are accessible via http/https and a URL.

Know the differences between the storage classes. Ie S3 Standard, IA and One Zone IA

Know how to make an S3 object public. This can be done with an ACL or a bucket policy.

Know the consistency model for puts, overwrite puts, deletes. Puts are strongly consistent. Overwrite puts and deletes are eventually consistent.

Know the options for encrypting data at rest eg server side with AMZ managed keys, server-side with KMS, server-side with customer provided keys, or encrypt the data on the client side before ingesting into S3

Know that signed URLs can provide a user application with time limited access to S3 objects, without requiring public read access. It prevents other sites hot linking to your objects.

Know which features help prevent or recover from accidental data loss. Eg versioning and a conditional deletion based on MFA.

Know the benefits of multipart upload which includes improved throughput and quick recovery from failed uploads due to a network error.

Know that an alternative to Cloud Trail to track access to S3 is server access logging on a bucket

Know that Glacier natively encrypts data at rest

Know that Glacier now supports expedited, standard and bulk retrieval.

EC2

Know the purpose of cluster placement groups – puts a set of instances in the same AZ for lowest latency between them.

Know the purpose of spread placement groups – it puts instances on separate physical servers for redundancy purposes.

Know how to find metadata from within an instance eg the pubic IP address Eg curl http://169.254.169.254/latest/meta-data/

Know some use cases for tagging EC2 resources eg keeping track of projects, billing.

Know about the choices of instance tenancy which are default (shared), dedicated instance and dedicated host. Dedicated instances guarantees your instances will not share hardware with other accounts. Dedicated host does the same but also provides you with a unique host ie, in other words it is a specific host.

Know that an instance must be assigned at least one security group.

Can you change the security group attached to an instance after it is launched? Yes!

Know that for EMR and Elastic Beanstalk the customer has full admin privilege of the instance, they can logon to the OS. For RDS you cannot login to the instance.

Know there are three choices of Reserved Instance ie Standard RIs, Convertible RIs, Scheduled RIs. For all of them, they can apply to usage across all AZs, or, if you specify an AZ, it is also a capacity reservation. Attributes of an RI are instance type, platform (AMZ EC2, SUSE, Redhat, Windows, MS SQL), tenancy (dedicated vs default), AZ (optional).

Know the different pricing models eg On-demand, RIs, Dedicated, Spot. What is spot good for? Eg Data processing where the application can recover gracefully from failures.

Know the use case for user data ie bootstrapping an instance

Know that you only need to retain the most recent snapshot for the snapshot to be complete allowing the ability to fully restore the data

Know that a requirement for restoring individual files may imply the use of a 3rd party backup solution.

EBS

Know how to encrypt an EBS volume (a tick box if using the console, allowing choice of KMS key)

Can a volume be used while a snapshot is in progress? Yes.

Know that snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3)

Know the difference between EBS and Instance store. EBS is a service accessed over the network. Instance Store is local storage in the physical server and the storage is lost when the instance is stopped. So instance store is only for scratch data, temporary data, buffers, swap files.

Know that to fully benefit from SSD PIOPS or Enhanced Throughput HDD it is recommended to use an "EBS Optimized2 instance which gives a dedicated network path between the instance and EBS reducing contention with other network traffic.

VPC

Know that you can allocate a permanent private address to an instance.

Know the detail of NACLs and Security groups in terms of where they are applied and the default rules.

What is the purpose of an internet gateway?

Know when you are charged for an EIP. You are not charged while it is attached to a running instance.

Know the rules of VPC peering. Eg no overlapping CIDR ranges, non transitive, traffic stays on the AWS network. Can be cross region, cross account.

Know the requirements to set up a web server accessible using a custom DNS name – ie Route 53, EIP, Public Subnet, Internet Gateway, appropriate security group and NACLs.

Know the use cases for VPN and Direct Connect. Know how to make a single DC connection more highly available eg by having vpn backup or a second direct connect connection.

Can you detach the primary (eth0) interface from an instance? No.

What are some possible reasons for an instance not being able to reach the internet, when this is desired? Hint, IGW. NAT if instance is in a private subnet.

What are some possible reasons for an instance not being reachable from the internet, when this is desired. Hint, IGW, Public Subnet, Public address or EIP, Approprate SGs and NACLs.

What are some additional issues if the instance is behind a load balancer? Eg the load balancer must be in a public subnet and security groups or network ACLs could be the problem.

Know how a Bastion Host/Jump box can provide access to instances in a private subnet in terms of the what instances need a public address and what security groups would be needed.

Database

Know when to use RDS v NoSQL. For an application requiring complex queries and table joins, use RDS. For applications that need extreme read/write performance use DynamoDB.

Know whether you have OS access to RDS instances. Hint: No.

Know which DB engines are supported with RDS. MS SQL, Postgres, MySQL, MariaDB, Oracle, Aurora.

Know about RDS read replicas. Know the use case, eg to offload the performance impact of a read ony/reporting application from the Master DB.

Can you initiate a forced failover an RDS instance? Hint: Yes, there is a reboot with failover option.

Know which RDS databases support read replicas (MySQL, MariaDB and PostgreSQL). You can replicate across regions.

Autoscaling

Know the best practice is to use ASGs across AZs. Can an ALB route to instances across AZs? Hint: Yes

What are the implications for user state during scale-in? Servers need to be stateless, which means storing user session state elsewhere, eg in a database The DB may refer to larger objects in S3.

Know how you would change the instance type for instances that are part of an ASG. Hint: Change the Launch Configuration.

Load Balancing

Know some use cases. Know the 3 choices of LB and what each supports ie HTTP/HTTPS for the ALB, any TCP or UDP port with or without TLS for the NLB. The NLB gets a static address whereas the ALB gets a DNS name.

Know that to capture client connection information in detail you can enable access logs.

Route 53

Know how to create ALIAS and CNAME record sets and what they can point to. Eg an ALIAS record can point to the zone apex (xyz.com) , but a CNAME cannot.

Know the the routing policies eg Geolocation and others.

Cloud Front

Know its use case. It is used in many of the course scenarios

Know that if you are sending application logs to an S3 bucket, and then you deploy Cloud Front, you would want to enable Cloud Front to also deliver logs to S3

Cloudwatch

Know the difference between Cloudwatch and Cloudtrail

Know that you would need a custom metric in order to keep track of memory utilization within an EC2 instance.

Know the default intervals when monitoring EC2 instances with and without “Detailed Monitoring”. Hint 5 mins and 1 min.

Know that Cloud Watch can integrate with SNS for notifications when an Alarm changes state.

SNS

Know the endpoints supported ie email, SMS and others.

SQS

Know the use cases covered in the course.

Cloud Trail

Know that you can choose to enable it for all regions or only some regions (all regions is recommended). You cannot choose which services are included. Logs can be delivered to an s3 bucket.

Configuration Management

OpsWorks – a high level view. Supports Chef and Puppet

Chef uses terms such as chef, recipes and cookbooks.

Security

Know the use case for SAML eg SSO from a SAML IDP.

Know the difference between SAML and Web Identity

Know the difference between SAML and Web Identify Federation eg Web Identity allows an app using an OICD compatible identify provider to gain access to eg S3, using STS)

Know the difference between encryption at rest and in transit and some of the supported features

Which security aspects are the customers responsibility?

Know that multiple accounts can provide autonomy and control to divisions with an organisation, and that IAM cross account access can enable corporate IT to maintain governance, and that Consolidated Billing can provide cost oversight.

Directory service

Know that one AD connector allows, amongst other things, a user to authenticate to an on premices AD and gain access to the AWS console. A user can login to an application in the cloud using his or her AD credentials. An EC2 instance can join the on prem domain.

Big Data

Have a high level knowledge of eg Kinesis for streaming large amounts of data into the cloud

Kinesis

Know that Kinesis streams allows for re l time data processing, to collect data as it generated andreact to critical information. For example, to analyse log data in real time to react to business oroperational critical information.

Know that Kinesis Streams stores data for 24 hours, configurable up to 7 days. This would allow a consumer application to re-analyse the data.

Know the difference between the 3 (there is a 4th now) services

Kinesis Streams

Kinesis Firehose

Kinesis Analytics

Redshift

Know some use cases. it's a data warehousing solution.

EMR

Know some use cases. It's a managed Hadoop service.

Support options

Know the names of the 4 support options and the maximum ticket response times

Misc

Know which services are global, regional and per AZ

ECS

Read the FAQ. Here are some highlights.

A managed EC2 container service. Allows you to manage Docker containers on a cluster of EC2 instances

Containers allow you to run an app and its dependencies in a resource isolated process

Containers are created from a read-only template called an image

Images are stored in a registry such as DockerHub or AWS Elastic Container Registry which is a managed Docker registry service.

A task definition – required to run docker containers. Json files.

ECS service allows you to run a desired count of a task definition. Similar to an autoscaling group.

ECS cluster – a group of instances

ECS agent connects EC2 instances to your ECS cluster. Linux only not windows

Security groups operate at the instance level.

Elasticache

Know the use cases

Elastic Transcoder

Converts media files into formats for Smartphones, Tablets, PCs. It has knowledge of popular output formats. Pay per minute and resolution

Storage gateway

A high level knowledge of the use cases for Storage Gateway

Know the different types of Storage Gateway