

Top 100 Amazon Web Service Interview Questions with Answers

1. What is AWS ?

AWS attains as Amazon Web Service; this is a gathering of remote computing settings also identified as cloud computing policies. This unique realm of cloud computing is also recognized as IaaS or Infrastructure as a Service.

2. What are the key components of AWS ?

The fundamental elements of AWS are

Route 53: A DNS web service

- **Easy E-mail Service:** It permits addressing e-mail utilizing RESTFUL API request or through normal SMTP
- **Identity and Access Management:** It gives heightened protection and identity control for your AWS account
- **Simple Storage Device or (S3):** It is a warehouse equipment and the well-known widely utilized AWS service
- **Elastic Compute Cloud (EC2):** It affords on-demand computing sources for hosting purposes. It is extremely valuable in trouble of variable workloads
- **Elastic Block Store (EBS):** It presents persistent storage masses that connect to EC2 to enable you to endure data beyond the lifespan of a particular EC2
- **CloudWatch:** To observe AWS sources, It permits managers to inspect and obtain key Additionally, one can produce a notification alert in the state of crisis.

3. what is S3 ?

S3 holds for Simple Storage Service. You can utilize S3 interface to save and recover the unspecified volume of data, at any time and from everywhere on the web. For S3, the payment type is “pay as you go”.

4.What Is The Importance Of Buffer In Amazon Web Services?

An Elastic Load Balancer ensures that the incoming traffic is distributed optimally across various AWS instances. A buffer will synchronize different

components and makes the arrangement additional elastic to a burst of load or traffic. The components are prone to work in an unstable way of receiving and processing the requests. The buffer creates the equilibrium linking various apparatus and crafts them effort at the identical rate to supply more rapid services.

5.What does an AMI include ?

- An AMI comprises the following elements.
- A template to the source quantity concerning the instance.
- Launch authorities determine which AWS accounts can avail the AMI to drive instances
- A base design mapping that defines the amounts to join to the instance while it is originated.

6.How can you send request to Amazon S3 ?

Amazon S3 is a REST service, you can transmit the appeal by applying the REST API or the AWS SDK wrapper archives that envelop the underlying Amazon S3 REST API.

7. How many buckets can you create in AWS by default ?

In each of your AWS accounts, by default, You can produce up to 100 buckets.

8.What Is The Importance Of Buffer In Amazon Web Services?

An Elastic Load Balancer ensures that the incoming traffic is distributed optimally across various AWS instances. A buffer will synchronize different components and makes the arrangement additional elastic to a burst of load or traffic. The components are prone to work in an unstable way of receiving and processing the requests. The buffer creates the equilibrium linking various apparatus and crafts them effort at the identical rate to supply more rapid services.

9. What Is The Way To Secure Data For Carrying In The Cloud?

One thing must be ensured that no one should seize the information in the cloud while data is moving from point one to another and also there should not be any leakage with the security key from several storerooms in the cloud. Segregation of information from additional companies' information and then encrypting it by means of approved methods is one of the options.

10. Name The Several Layers Of Cloud Computing?

Here is the list of layers of the cloud computing

PaaS – Platform as a Service

IaaS – Infrastructure as a Service

SaaS – Software as a Service

11.Explain Can You Vertically Scale An Amazon Instance ? How?

Surely, you can vertically estimate on Amazon instance. During that Twist up a fresh massive instance than the one you are currently governing. Delay that instance and separate the source webs mass of server and dispatch. Next, quit your existing instance and separate its source quantity.

Note the different machine ID and connect that source mass to your fresh server

Also, begin it repeatedly Study AWS Training Online From Real Time Experts

12. What Are The Components Involved In Amazon Web Services?

There are 4 components involved and are as below. Amazon S3: with this, one can retrieve the key information which are occupied in creating cloud structural design and amount of produced information also can be stored in this component that is the consequence of the key specified. Amazon EC2 instance: helpful to run a large distributed system on the Hadoop cluster. Automatic parallelization and job scheduling can be achieved by this component.

Amazon SQS: this component acts as a mediator between different controllers. Also worn for cushioning requirements those are obtained by the manager of Amazon.

Amazon SimpleDB: helps in storing the transitional position log and the errands executed by the consumers.

13. What Is Lambda@edge In Aws?

In AWS, we can use Lambda@Edge utility to solve the problem of low network latency for end users.

In Lambda@Edge there is no need to provision or manage servers. We can just upload our Node.js code to AWS Lambda and create functions that will be triggered on CloudFront requests.

When a request for content is received by CloudFront edge location, the Lambda code is ready to execute.

This is a very good option for scaling up the operations in CloudFront without managing servers.

14. Distinguish Between Scalability And Flexibility?

The aptitude of any scheme to enhance the tasks on hand on its present hardware resources to grip inconsistency in command is known as scalability. The capability of a scheme to augment the tasks on hand on its present and supplementary hardware property is recognized as flexibility, hence enabling the industry to convene command devoid of putting in the infrastructure at all. AWS has several configuration management solutions for AWS scalability, flexibility, availability and management.

15. Name The Various Layers Of The Cloud Architecture?

There are 5 layers and are listed below

CC- Cluster Controller

SC- Storage Controller

CLC- Cloud Controller

Walrus

NC- Node Controller

16.Explain can you vertically scale an Amazon instance ? How ?

Surely, you can vertically estimate on Amazon instance. During that

Twist up a fresh massive instance than the one you are currently governing

Delay that instance and separate the source webs mass of server and dispatch

Next, quit your existing instance and separate its source quantity

Note the different machine ID and connect that source mass to your fresh server

Also, begin it repeatedly Study AWS Training Online From Real Time Experts

17. Explain what is T2 instances ?

T2 instances are outlined to present average baseline execution and the ability to explode to powerful execution as needed by the workload.

18. In VPC with private and public subnets, database servers should ideally be launched into which subnet ?

Among private and public subnets in VPC, database servers should ideally originate toward separate subnets.

19. Explain how the buffer is used in Amazon web services ?

The buffer is utilized to deliver the system further robust to handle traffic or load by synchronizing different component. Usually, elements sustain and process the demands in an unreliable mode, With the aid of buffer, the elements will be (sap training) equivalent and will operate at the similar speed to accommodate high-speed services.

20. While connecting to your instance what are the possible connection issues one might face ?

The feasible connection failures one might battle while correlating instances are

- Consolidation timed out
- User key not acknowledged by the server
- Host key not detected, license denied
- Unguarded private key file
- Server rejected our key or No sustained authentication program available
- Error handling Mind Term on Safari Browser
- Error utilizing Mac OS X RDP Client

21. Explain Elastic Block Storage ? What type of performance can you expect ? How do you back it up? How do you improve performance ?

That indicates it is RAID warehouse to begin with, so it's irrelevant and faults tolerant. If disks expire in the RAID you don't miss data. Excellent! It is more virtualized, therefore you can provision and designate warehouse, and connect it

to your server with multiple API appeals. No calling the storage specialist and asking him or her to operate specific requests from the hardware vendor.

Execution on EBS can manifest variability. Such signifies that can run above the SLA enforcement level, suddenly descend under it. The SLA gives you among a medium disk I/O speed you can foresee. That can prevent any groups particularly performance specialists who suspect stable and compatible disk throughput on a server. Common physically entertained servers perform that direction. Pragmatic AWS cases do not.

Backup EBS masses by utilizing the snap convenience through API proposal or by a GUI interface same elasticfox.

Progress execution by practicing Linux software invasion and striping over four extents.

21. What Are The Different Types Of Events Triggered By Amazon Cloud Front?

Different types of events triggered by Amazon CloudFront are as follows:

- **Viewer Request:** When an end user or a client program makes an HTTP/HTTPS request to CloudFront, this event is triggered at the Edge Location closer to the end user.
- **Viewer Response:** When a CloudFront server is ready to respond to a request, this event is triggered.
- **Origin Request:** When CloudFront server does not have the requested object in its cache, the request is forwarded to Origin server. At this time this event is triggered.
- **Origin Response:** When CloudFront server at an Edge location receives the response from Origin server, this event is triggered.

22. Which Automation Gears Can Help With Spinup Services?

The API tools can be used for spinup services and also for the written scripts. Those scripts could be coded in Perl, bash or other languages of your preference. There is one more option that is patterned administration and stipulating tools such as a dummy or improved descendant. A tool called Scalr

can also be used and finally we can go with a controlled explanation like a Rightscale.

23. What Is An Ami ? How Do I Build One?

AMI holds for Amazon Machine Image. It is efficiently a snap of the source filesystem. Products appliance servers have a bio that shows the master drive report of the initial slice on a disk. A disk form though can lie anyplace physically on a disc, so Linux can boot from an absolute position on the EBS warehouse interface.

Create a unique AMI at beginning rotating up and instance from a granted AMI. Later uniting combinations and components as needed. Comprise wary of setting delicate data over an AMI (learn salesforce online). For instance, your way credentials should be joined to an instance later spinup. Among a database, mount an external volume that carries your MySQL data next spinup actually enough.

24. What Are The Main Features Of Amazon Cloud Front?

Some of the main features of Amazon CloudFront are as follows: Device Detection Protocol Detection Geo Targeting Cache Behavior Cross Origin Resource Sharing Multiple Origin Servers HTTP Cookies Query String Parameters Custom SSL.

25. What Is The Relation Between An Instance And Ami?

AMI can be elaborated as Amazon Machine Image, basically, a template consisting software configuration part. For example an OS, applications, application server. If you start an instance, a duplicate of the AMI in a row as an unspoken attendant in the cloud.

26. What Is Amazon Ec2 Service?

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable (scalable) computing capacity in the cloud. You can use Amazon EC2 to launch as many virtual servers you need. In Amazon EC2 you can configure security and networking as well as manage storage. Amazon EC2 service also helps in obtaining and configuring capacity using minimal friction.

27. What Are The Features Of The Amazon Ec2 Service?

As the Amazon EC2 service is a cloud service so it has all the cloud features. Amazon EC2 provides the following features:

- Virtual computing environment (known as instances)
- re-configured templates for your instances (known as Amazon Machine Images – AMIs)
- Amazon Machine Images (AMIs) is a complete package that you need for your server (including the operating system and additional software)
- Amazon EC2 provides various configurations of CPU, memory, storage and networking capacity for your instances (known as instance type)
- Secure login information for your instances using key pairs (AWS stores the public key and you can store the private key in a secure place)
- Storage volumes of temporary data is deleted when you stop or terminate your instance (known as instance store volumes)
- Amazon EC2 provides persistent storage volumes (using Amazon Elastic Block Store – EBS)
- A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances using security groups
- Static IP addresses for dynamic cloud computing (known as Elastic IP address)
- Amazon EC2 provides metadata (known as tags)
- Amazon EC2 provides virtual networks that are logically isolated from the rest of the AWS cloud, and that you can optionally connect to your own network (known as virtual private clouds – VPCs)

28. What Is Amazon Machine Image And What Is The Relation Between Instance And Ami?

- Amazon Web Services provides several ways to access Amazon EC2, like web-based interface, AWS Command Line Interface (CLI) and Amazon Tools for Windows Powershell. First, you need to sign up for an AWS account and you can access Amazon EC2.

- Amazon EC2 provides a Query API. These requests are HTTP or HTTPS requests that use the HTTP verbs GET or POST and a Query parameter named Action.

29. What Is Amazon Machine Image (ami)?

An Amazon Machine Image (AMI) is a template that contains a software configuration (for example, an operating system, an application server, and applications). From an AMI, we launch an instance, which is a copy of the AMI running as a virtual server in the cloud. We can even launch multiple instances of an AMI.

30. What Is The Relation Between Instance And Ami?

We can launch different types of instances from a single AMI. An instance type essentially determines the hardware of the host computer used for your instance. Each instance type offers different compute and memory capabilities.

After we launch an instance, it looks like a traditional host, and we can interact with it as we would do with any computer. We have complete control of our instances; we can use sudo to run commands that require root privileges.

31. Explain Storage For Amazon Ec2 Instance.?

Amazon EC2 provides many data storage options for your instances. Each option has a unique combination of performance and durability. These storages can be used independently or in combination to suit your requirements.

There are mainly four types of storages provided by AWS:

- Amazon EBS: Its durable, block-level storage volumes can be attached to a running Amazon EC2 instance. The Amazon EBS volume persists independently from the running life of an Amazon EC2 instance. After an EBS volume is attached to an instance, you can use it like any other physical hard drive. Amazon EBS encryption feature supports encryption feature.
- Amazon EC2 Instance Store: Storage disk that is attached to the host computer is referred to as instance store. The instance storage provides temporary block-level storage for Amazon EC2 instances. The data on an

instance store volume persists only during the life of the associated Amazon EC2 instance; if you stop or terminate an instance, any data on instance store volumes is lost.

- Amazon S3: Amazon S3 provides access to reliable and inexpensive data storage infrastructure. It is designed to make web-scale computing easier by enabling you to store and retrieve any amount of data, at any time, from within Amazon EC2 or anywhere on the web.
- Adding Storage: Every time you launch an instance from an AMI, a root storage device is created for that instance. The root storage device contains all the information necessary to boot the instance. You can specify storage volumes in addition to the root device volume when you create an AMI or launch an instance using block device mapping.

32. What Are The Security Best Practices For Amazon Ec2?

There are several best practices for secure Amazon EC2. Following are few of them.

- Use AWS Identity and Access Management (IAM) to control access to your AWS resources.
- Restrict access by only allowing trusted hosts or networks to access ports on your instance.
- Review the rules in your security groups regularly, and ensure that you apply the principle of least
- Privilege — only open up permissions that you require.
- Disable password-based logins for instances launched from your AMI. Passwords can be found or cracked, and are a security risk.

33. Explain Stopping, Starting, And Terminating An Amazon Ec2 Instance?

Stopping and Starting an instance: When an instance is stopped, the instance performs a normal shutdown and then transitions to a stopped state. All of its Amazon EBS volumes remain attached, and you can start the instance again at a later time. You are not charged for additional instance hours while the instance is in a stopped state.

Terminating an instance: When an instance is terminated, the instance performs a normal shutdown, then the attached Amazon EBS volumes are deleted unless the volume's deleteOnTermination attribute is set to false. The instance itself is also deleted, and you can't start the instance again at a later time.

34 .What is S3 ? What is it used for ? Should encryption be used ?

S3 implies for Simple Storage Service. You can believe it similar ftp warehouse, wherever you can transfer records to and from beyond, merely not uprise it similar to a filesystem. AWS automatically places your snaps there, at the same time AMIs there. sensitive data is treated with Encryption, as S3 is an exclusive technology promoted by Amazon themselves, and as still unproven vis-a-vis a protection viewpoint.

35. What is an AMI ? How do I build one ?

AMI holds for Amazon Machine Image. It is efficiently a snap of the source filesystem. Products appliance servers have a bio that shows the master drive report of the initial slice on a disk. A disk form though can lie anyplace physically on a disc, so Linux can boot from an absolute position on the EBS warehouse interface.

Create a unique AMI at beginning rotating up and instance from a granted AMI. Later uniting combinations and components as needed. Comprise wary of setting delicate data over an AMI (learn salesforce online). For instance, your way credentials should be joined to an instance later spinup. Among a database, mount an external volume that carries your MySQL data next spinup actually enough.

36.Can I Vertically Scale An Amazon Instance? How?

Yes.This is an incredible feature of AWS and cloud virtualization. Spin up a new larger instance than the one you are currently running. Pause that instance and detach the root ebs volume from this server and discard. Then stop your live instance, detach its root volume. Note down the unique device ID and attach that root volume to your new server. And then start it again. Voila, you have scaled vertically in-place!!

37. Define Auto Scaling ?

Answer: Auto-scaling is one of the conspicuous characteristics feature of AWS anywhere it authorizes you to systematize and robotically obligation and twist up new models externally that necessary for your entanglement. This can be accomplished by initiating brims and metrics to view.If these proposals are demolished, the latest model of your preference will be configured, wrapped up and cloned into the weight administrator panel.

38. Which automation gears can help with spinup services ?

For the written scripts we can use spinup services with the help of API tools.These scripts could be coded in bash, Perl, or any another language of your choice.There is one more alternative that is patterned control and stipulating devices before-mentioned as a dummy or advanced descendant. A machine termed as Scalar can likewise be utilized and ultimately we can proceed with a constrained expression like a RightScale.

39. Is it possible to scale an Amazon instance vertically ? How ?

Yes, it is possible to scale an Amazon instance vertically because of an unbelievable characteristic of cloud virtualization and AWS. Spinup is a huge case while correlated to the one which you are working with. Let up the case and distribute the source EBS bulk of this server and eliminate. Subsequent, end your existing instance, exclude its root volume. Enter down the peculiar device ID and join source volume to your fresh server and begin it repeatedly. This is the way to scaling vertically in position.

40. How the processes start, stop and terminate works ?

Starting and stopping of an instance: If an instance goes arrested or died, the instance performs a normal power cut and then transfer over to a sealed area. You can build the case then for all the EBS masses of Amazon persist and associated. If an instance is in ending state, suddenly you will not get charged to the additional instance

Finishing the instance: If an instance goes stopped it serves to perform a standard blackout, therefore the EBS capacities which are connected will get

excluded save the volume's delete On Termination feature is fixed to zero. In such instances, the instance will get eliminated and cannot set it up afterward.

41. Explain in detail the function of Amazon Machine Image (AMI) ?

An Amazon Machine Image AMI is a pattern that comprises a software conformation (for instance, an operative system, a request server, and applications). From an AMI, we present an example, which is a duplicate of the AMI successively as a virtual server in the cloud. We can even offer plentiful examples of an AMI.

42. If I'm expending Amazon Cloud Front, can I custom Direct Connect to handover objects from my own data centre ?

Certainly. Amazon Cloud Front stipulations culture rises computing sources of separate AWS. By AWS Direct Connect, you will be accelerating with the appropriate information substitution rates. AWS Training Free Demo

43.If my AWS Direct Connect flops, will I lose my connection ?

If a gridlock AWS Direct connects has been transposed, in the event of a let-down, it will convert over to the next one. It is voluntary to allow Bidirectional Forwarding Detection (BFD) while systematizing your rules to safeguard quicker identification and failover. Proceeding the opposite hand, if you have built a backup IPsec VPN connecting as an option, all VPC transactions will failover to the backup VPN association routinely.

44. What is AWS Certificate Manager ?

AWS Certificate Manager (ACM) manages the complexity of extending, provisioning, and regulating certificates granted over ACM (ACM Certificates) to your AWS-based websites and forms. You work ACM to petition and maintain the certificate and later practice other AWS services to provision the ACM Certificate for your website or purpose. As designated in the subsequent instance, ACM Certificates are currently ready for performance with only Elastic Load Balancing and Amazon CloudFront. You cannot handle ACM Certificates outside of AWS.

45. Explain What is Redshift ?

The executes it easy and cost-effective to efficiently investigate all your data employing your current marketing intelligence devices which is a completely controlled, high-speed, it is petabyte-scale data repository service known as Redshift.

46. Mention what are the differences between Amazon S3 and EC2 ?

S3: Amazon S3 is simply a storage aid, typically applied to save huge binary records. Amazon too has additional warehouse and database settings, same as RDS to relational databases and DynamoDB concerning NoSQL.

EC2: An EC2 instance is similar to a foreign computer working Linux or Windows and on which you can install whatever software you need, including a Network server operating PHP code and a database server.

47.Explain what is C4 instances ?

C4 instances are absolute for compute-bound purposes that serve from powerful-performance processors. AWS Interview Questions and Answers

48. Explain what is DynamoDB in AWS ?

Amazon DynamoDB is a completely controlled NoSQL database aid that renders quick and anticipated execution with seamless scalability. You can perform Amazon DynamoDB to formulate a database table that can save and reclaim any quantity of data, and help any level of application transactions. Amazon DynamoDB automatically increases the data and transactions for the table above an adequate number of servers to supervise the inquiry function designated by the customer and the volume of data saved, while keeping constant and quick execution.

49. Explain what is ElastiCache ?

A web service that executes it comfortable to set up, maintain, and scale classified in-memory cache settings in the cloud is known as ElastiCache.

50. What is the AWS Key Management Service ?

A managed service that makes it easy for you to create and control the encryption keys used to encrypt your data is known as the AWS Key Management Service (AWS KMS).

51. What is AWS WAF ? What are the potential benefits of using WAF ?

AWS WAF is a web application firewall that lets you monitor the HTTP and HTTPS applications that are promoted to Amazon CloudFront and gives you regulate path to your content. Based on circumstances that you stipulate, such as the IP addresses that grants originate from or the consequences of query series, CloudFront returns to applications either with the petitioned content or with an HTTP 403 situation code (Forbidden). You can further configure CloudFront to restore a pattern failure page when an application is obstructed.

Advantages of utilizing WAF:

- Further security versus web initiatives relating circumstances that you designate. You can describe situations by managing characteristics of web inquiries such as the IP address that the applications originate from, the rates in headers, chains that rise in the applications, and the presence of hateful SQL code in the call, which is recognized as SQL injection.
- Rules that you can reuse for various network appeals
- Real-time metrics and examined web demands
- Computerized command practicing the AWS WAF API

52. What is Amazon EMR ?

Amazon Elastic MapReduce (Amazon EMR) is a survived cluster stage that interprets working big data structures, before-mentioned as Apache Spark and Apache Hadoop, on AWS to treat and investigate enormous volumes of data. By adopting these structures and relevant open-source designs, such as Apache Pig and Apache Hive, you can prepare data for analytics goals and marketing intellect workloads. Additionally, you can use Amazon EMR to convert and migrate vast masses of information into and of other AWS data repositories and databases, such as Amazon DynamoDB and Amazon Simple Storage Service (Amazon S3).

53. What is AWS Data Pipeline ? and what are the components of AWS Data Pipeline ?

A web service that you can implement to automate the journey and exchange of data are called AWS Data Pipeline. Beside AWS you can define data-driven workflows so that companies can be reliant on the favorable execution of initial jobs.

The succeeding components of AWS Data Pipeline work collectively to get your data:

- A pipeline key indicates the business appraised of your data administration. For additional data, observe Pipeline Definition File Syntax.
- A pipeline registers and tracks responsibilities. You upload your pipeline accuracy to the pipeline and when excite the pipeline. You can control the pipeline variety for a working pipeline and stimulate the pipeline regularly for it to receive the issue. You can deactivate the pipeline, replace a data storage, and before initiate the pipeline newly. If you are terminated with your pipeline, you can cancel it.
- Task Runner studies for services and then performs those duties. For instance, Task Runner could replicate log records to Amazon S3 and push Amazon EMR organizations. Task Runner is uns automatically on devices designed by your pipeline keys. You can create a custom task runner application, or you can make the Task Runner form that is offered by AWS Data Pipeline.

54. What is Amazon Kinesis Firehose ?

A fully managed service for delivering real-time streaming data to destinations such as Amazon Simple Storage Service (Amazon S3) and Amazon Redshift is known as Amazon Kinesis Firehose.

55. What Is Amazon CloudSearch and its features ?

A thoroughly managed service in the cloud that creates it simple to set up, maintain, and estimate a search solution for your website or application is called Amazon CloudSearch.

we can use Amazon CloudSearch to catalog and explore both plain text and structured data. Amazon CloudSearch characteristics:

Entire text search with language-specific text processing

- Range searches
- Prefix searches
- Boolean search
- FacetingTerm boosting
- Highlighting
- Autocomplete Advices

56. Explain what is Regions and Endpoints in AWS ?

An endpoint is a URL that is the entry point for a web service. To decrease data latency in your forms, most Amazon Web Services results enable you to choose a sectional endpoint to make your applications.

Some services, before-mentioned as Amazon EC2, let you define an endpoint that does not cover a particular area. IAM, do not sustain regions; their endpoints, consequently, do not incorporate a region proposed by Amazon Web Services Tutorials Some services. Amazon Web Services Tutorials

57. What are the different types of cloud services ?

Infrastructure as a Service (IaaS), Software as a Service (SaaS), Platform as a Service (PaaS), and Data as a Service (DaaS).

58. What is SimpleDB ?

A structured records or data repository that encourages indexing and data doubts to both EC2 and S3 is known as SimpleDB.

59. What is the type of architecture, where half of the workload is on the public load while at the same time half of it is on the local storage ?

Hybrid cloud architecture.

60. Should encryption be used for S3 ?

Encryption should be examined for delicate information or data as S3 is a proprietary technology.

61. What are the various AMI design options ?

Fully Baked AMI, JeOS (just enough operating system) AMI, and Hybrid AMI.

62. What is Geo Restriction in CloudFront ?

Geo restriction, also known as geoblocking, is used to prevent users in specific geographic locations from accessing content that you're distributing through a CloudFront web distribution. Amazon Web Services Training

63.Can S3 be used with EC2 instances, how?

It can be used for instances with root devices backed by local instance storage. By using Amazon S3, developers have access to the same highly scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites. In order to execute systems in the Amazon EC2 environment, developers use the tools provided to load their Amazon Machine Images (AMIs) into Amazon S3 and to move them between Amazon S3 and Amazon EC2.

Another use case could be for websites hosted on EC2 to load their static content from S3.

64.Can I connect my corporate datacenter to the Amazon Cloud?

Yes, you can do this by establishing a VPN(Virtual Private Network) connection between your company's network and your VPC (Virtual Private Cloud), this will allow you to interact with your EC2 instances as if they were within your existing network.

65.Is it possible to change the private IP addresses of an EC2 while it is running/stopped in a VPC?

Primary private IP address is attached with the instance throughout its lifetime and cannot be changed, however secondary private addresses can be unassigned, assigned or moved between interfaces or instances at any point.

66.If I'm using Amazon CloudFront, can I use Direct Connect to transfer objects from my own data center?

Yes. Amazon CloudFront supports custom origins including origins from outside of AWS. With AWS Direct Connect, you will be charged with the respective data transfer rates.

67. If my AWS Direct Connect fails, will I lose my connectivity?

If a backup AWS Direct connect has been configured, in the event of a failure it will switch over to the second one. It is recommended to enable Bidirectional Forwarding Detection (BFD) when configuring your connections to ensure faster detection and failover. On the other hand, if you have configured a backup IPsec VPN connection instead, all VPC traffic will failover to the backup VPN connection automatically. Traffic to/from public resources such as Amazon S3 will be routed over the Internet. If you do not have a backup AWS Direct Connect link or a IPsec VPN link, then Amazon VPC traffic will be dropped in the event of a failure.

68.What is the difference between Scalability and Elasticity?

Scalability is the ability of a system to increase its hardware resources to handle the increase in demand. It can be done by increasing the hardware specifications or increasing the processing nodes.

Elasticity is the ability of a system to handle increase in the workload by adding additional hardware resources when the demand increases(same as scaling) but also rolling back the scaled resources, when the resources are no longer needed. This is particularly helpful in Cloud environments, where a pay per use model is followed.

69.How will you change the instance type for instances which are running in your application tier and are using Auto Scaling. Where will you change it ?

In Auto Scaling launch configuration, Auto scaling tags configuration, is used to attach metadata to your instances, to change the instance type you have to use auto scaling launch configuration.

70.Suppose you have an application where you have to render images and also do some general computing. From the following services which service will best fit your need?

Classic Load Balancer and Application Load Balancer. You will choose an application load balancer, since it supports path based routing, which means it can take decisions based on the URL, therefore if your task needs image rendering it will route it to a different instance, and for general computing it will route it to a different instance.

71. You have a content management system running on an Amazon EC2 instance that is approaching 100% CPU utilization. How to reduce load on the Amazon EC2 instance?

Create a load balancer, and register the Amazon EC2 instance with it.

- Creating alone an autoscaling group will not solve the issue, until you attach a load balancer to it. Once you attach a load balancer to an autoscaling group, it will efficiently distribute the load among all the instances. Option B – CloudFront is a CDN, it is a data transfer tool therefore will not help reduce load on the EC2 instance. Similarly the other option – Launch configuration is a template for configuration which has no connection with reducing loads.

72. When should I use a Classic Load Balancer and when should I use an Application load balancer?

A Classic Load Balancer is ideal for simple load balancing of traffic across multiple EC2 instances, while an Application Load Balancer is ideal for microservices or container-based architectures where there is a need to route traffic to multiple services or load balance across multiple ports on the same EC2 instance.

73. What does Connection draining do?

- A. Terminates instances which are not in use.
- B. Re-routes traffic from instances which are to be updated or failed a health check.
- C. Re-routes traffic from instances which have more workload to instances which have less workload.
- D. Drains all the connections from an instance, with one click.

Answer B.

Connection draining is a service under ELB which constantly monitors the health of the instances. If any instance fails a health check or if any instance has to be patched with a software update, it pulls all the traffic from that instance and re routes them to other instances.

74. When an instance is unhealthy, it is terminated and replaced with a new one, which of the following services does that?

- A. Sticky Sessions
- B. Fault Tolerance
- C. Connection Draining
- D. Monitoring

Answer B.

When ELB detects that an instance is unhealthy, it starts routing incoming traffic to other healthy instances in the region. If all the instances in a region becomes unhealthy, and if you have instances in some other availability zone/region, your traffic is directed to them. Once your instances become healthy again, they are re routed back to the original instances.

75. What are lifecycle hooks used for in AutoScaling?

They are used to put an additional wait time to a scale in or scale out event. Lifecycle hooks are used for putting wait time before any lifecycle action i.e launching or terminating an instance happens. The purpose of this wait time, can be anything from extracting log files before terminating an instance or installing the necessary softwares in an instance before launching it.

76. A user has setup an Auto Scaling group. Due to some issue the group has failed to launch a single instance for more than 24 hours. What will happen to Auto Scaling in this condition?

- A. Auto Scaling will keep trying to launch the instance for 72 hours
- B. Auto Scaling will suspend the scaling process
- C. Auto Scaling will start an instance in a separate region
- D. The Auto Scaling group will be terminated automatically

Answer B.

Auto Scaling allows you to suspend and then resume one or more of the Auto Scaling processes in your Auto Scaling group. This can be very useful when you want to investigate a configuration problem or other issue with your

web application, and then make changes to your application, without triggering the Auto Scaling process.

77. Suppose you have an application where you have to render images and also do some general computing. which service will best fit your need?

Application Load Balancer, since it supports path based routing, which means it can take decisions based on the URL, therefore if your task needs image rendering it will route it to a different instance, and for general computing it will route it to a different instance.

78. What is the difference between Scalability and Elasticity?

Scalability is the ability of a system to increase its hardware resources to handle the increase in demand. It can be done by increasing the hardware specifications or increasing the processing nodes.

Elasticity is the ability of a system to handle increase in the workload by adding additional hardware resources when the demand increases (same as scaling) but also rolling back the scaled resources, when the resources are no longer needed. This is particularly helpful in Cloud environments, where a pay per use model is followed.

79. How will you change the instance type for instances which are running in your application tier and are using Auto Scaling. Where will you change it from the following areas?

Auto Scaling launch configuration

Auto scaling tags configuration, is used to attach metadata to your instances, to change the instance type you have to use auto scaling launch configuration.

80. You have a content management system running on an Amazon EC2 instance that is approaching 100% CPU utilization. Which option will reduce load on the Amazon EC2 instance?

Create a load balancer, and register the Amazon EC2 instance with it. Creating alone an autoscaling group will not solve the issue, until you attach a load balancer to it. Once you attach a load balancer to an autoscaling group, it will

efficiently distribute the load among all the instances. Option B – CloudFront is a CDN, it is a data transfer tool therefore will not help reduce load on the EC2 instance. Similarly the other option – Launch configuration is a template for configuration which has no connection with reducing loads.

81. When should I use a Classic Load Balancer and when should I use an Application load balancer?

A Classic Load Balancer is ideal for simple load balancing of traffic across multiple EC2 instances, while an Application Load Balancer is ideal for microservices or container-based architectures where there is a need to route traffic to multiple services or load balance across multiple ports on the same EC2 instance.

82. What does Connection draining do?

Re-routes traffic from instances which are to be updated or failed a health check. Connection draining is a service under ELB which constantly monitors the health of the instances. If any instance fails a health check or if any instance has to be patched with a software update, it pulls all the traffic from that instance and re routes them to other instances.

83. When an instance is unhealthy, it is terminated and replaced with a new one, which of the following services does that?

Fault Tolerance. When ELB detects that an instance is unhealthy, it starts routing incoming traffic to other healthy instances in the region. If all the instances in a region becomes unhealthy, and if you have instances in some other availability zone/region, your traffic is directed to them. Once your instances become healthy again, they are re routed back to the original instances.

84. What are lifecycle hooks used for in AutoScaling?

- A. They are used to do health checks on instances
- B. They are used to put an additional wait time to a scale in or scale out event.
- C. They are used to shorten the wait time to a scale in or scale out event

Answer B.

Lifecycle hooks are used for putting wait time before any lifecycle action i.e launching or terminating an instance happens. The purpose of this wait time, can be anything from extracting log files before terminating an instance or installing the necessary softwares in an instance before launching it.

85. A user has setup an Auto Scaling group. Due to some issue the group has failed to launch a single instance for more than 24 hours. What will happen to Auto Scaling in this condition?

A.Auto Scaling will keep trying to launch the instance for 72 hours

B.Auto Scaling will suspend the scaling process

C.Auto Scaling will start an instance in a separate region

D.The Auto Scaling group will be terminated automatically

Answer B.

Auto Scaling allows you to suspend and then resume one or more of the Auto Scaling processes in your Auto Scaling group. This can be very useful when you want to investigate a configuration problem or other issue with your web application, and then make changes to your application, without triggering the Auto Scaling process.

86. Which services you would not use to deploy an app?

Lambda is used for running server-less applications. It can be used to deploy functions triggered by events. When we say serverless, we mean without you worrying about the computing resources running in the background. It is not designed for creating applications which are publicly accessed.

87. How does Elastic Beanstalk apply updates?

By having a duplicate ready with updates before swapping. Elastic Beanstalk prepares a duplicate copy of the instance, before updating the original instance, and routes your traffic to the duplicate instance, so that, incase your updated application fails, it will switch back to the original instance, and there will be no downtime experienced by the users who are using your application.

88. How is AWS Elastic Beanstalk different than AWS OpsWorks?

AWS Elastic Beanstalk is an application management platform while OpsWorks is a configuration management platform. BeanStalk is an easy to use service which is used for deploying and scaling web applications developed with Java, .Net, PHP, Node.js, Python, Ruby, Go and Docker. Customers upload their code and Elastic Beanstalk automatically handles the deployment. The application will be ready to use without any infrastructure or resource configuration.

In contrast, AWS Opsworks is an integrated configuration management platform for IT administrators or DevOps engineers who want a high degree of customization and control over operations.

89. What happens if my application stops responding to requests in beanstalk?

AWS Beanstalk applications have a system in place for avoiding failures in the underlying infrastructure. If an Amazon EC2 instance fails for any reason, Beanstalk will use Auto Scaling to automatically launch a new instance. Beanstalk can also detect if your application is not responding on the custom link, even though the infrastructure appears healthy, it will be logged as an environmental event(e.g a bad version was deployed) so you can take an appropriate action.

90.How is AWS OpsWorks different than AWS CloudFormation?

OpsWorks and CloudFormation both support application modelling, deployment, configuration, management and related activities. Both support a wide variety of architectural patterns, from simple web applications to highly complex applications. AWS OpsWorks and AWS CloudFormation differ in abstraction level and areas of focus.

AWS CloudFormation is a building block service which enables customer to manage almost any AWS resource via JSON-based domain specific language. It provides foundational capabilities for the full breadth of AWS, without prescribing a particular model for development and operations. Customers define templates and use them to provision and manage AWS resources, operating systems and application code.

In contrast, AWS OpsWorks is a higher level service that focuses on providing highly productive and reliable DevOps experiences for IT administrators and ops-minded developers. To do this, AWS OpsWorks employs a configuration management model based on concepts such as stacks and layers, and provides integrated experiences for key activities like deployment, monitoring, auto-scaling, and automation. Compared to AWS CloudFormation, AWS OpsWorks supports a narrower range of application-oriented AWS resource types including Amazon EC2 instances, Amazon EBS volumes, Elastic IPs, and Amazon CloudWatch metrics.

91. I created a key in Oregon region to encrypt my data in North Virginia region for security purposes. I added two users to the key and an external AWS account. I wanted to encrypt an object in S3, so when I tried, the key that I just created was not listed. What could be the reason?

- A.External aws accounts are not supported.
- B.AWS S3 cannot be integrated KMS.
- C.The Key should be in the same region.
- D.New keys take some time to reflect in the list.

Answer C.

The key created and the data to be encrypted should be in the same region. Hence the approach taken here to secure the data is incorrect.

92. A company needs to monitor the read and write IOPS for their AWS MySQL RDS instance and send real-time alerts to their operations team. Which AWS services can accomplish this?

- A.Amazon Simple Email Service
- B.Amazon CloudWatch
- C.Amazon Simple Queue Service
- D.Amazon Route 53

Answer B.

Amazon CloudWatch is a cloud monitoring tool and hence this is the right service for the mentioned use case. The other options listed here are used for other purposes for example route 53 is used for DNS services, therefore CloudWatch will be the apt choice.

93. What happens when one of the resources in a stack cannot be created successfully in AWS OpsWorks?

When an event like this occurs, the “automatic rollback on error” feature is enabled, which causes all the AWS resources which were created successfully till the point where the error occurred to be deleted. This is helpful since it does not leave behind any erroneous data, it ensures the fact that stacks are either created fully or not created at all. It is useful in events where you may accidentally exceed your limit of the no. of Elastic IP addresses or maybe you may not have access to an EC2 AMI that you are trying to run etc.

94. What automation tools can you use to spinup servers?

Any of the following tools can be used:

Roll-your-own scripts, and use the AWS API tools. Such scripts could be written in bash, perl or other language of your choice.

Use a configuration management and provisioning tool like puppet or its successor Opscode Chef. You can also use a tool like Scalr.

Use a managed solution such as Rightscale.

95. Which AWS services will you use to collect and process e-commerce data for near real-time analysis?

- A. Amazon ElastiCache
- B. Amazon DynamoDB
- C. Amazon Redshift
- D. Amazon Elastic MapReduce

Answer B,C.

DynamoDB is a fully managed NoSQL database service. DynamoDB, therefore can be fed any type of unstructured data, which can be data from e-commerce websites as well, and later, an analysis can be done on them using Amazon Redshift. We are not using Elastic MapReduce, since a near real time analyses is needed.

96. Can I retrieve only a specific element of the data, if I have a nested JSON data in DynamoDB?

Yes. When using the GetItem, BatchGetItem, Query or Scan APIs, you can define a Projection Expression to determine which attributes should be retrieved from the table. Those attributes can include scalars, sets, or elements of a JSON document.

97.What happens to my backups and DB Snapshots if I delete my DB Instance?

When you delete a DB instance, you have an option of creating a final DB snapshot, if you do that you can restore your database from that snapshot. RDS retains this user-created DB snapshot along with all other manually created DB snapshots after the instance is deleted, also automated backups are deleted and only manually created DB Snapshots are retained.

98.How can I load my data to Amazon Redshift from different data sources like Amazon RDS, Amazon DynamoDB and Amazon EC2?

You can load the data in the following two ways:

You can use the COPY command to load data in parallel directly to Amazon Redshift from Amazon EMR, Amazon DynamoDB, or any SSH-enabled host.

AWS Data Pipeline provides a high performance, reliable, fault tolerant solution to load data from a variety of AWS data sources. You can use AWS Data Pipeline to specify the data source, desired data transformations, and then execute a pre-written import script to load your data into Amazon Redshift.

99.If my AWS Direct Connect fails, will I lose my connectivity?

If a backup AWS Direct connect has been configured, in the event of a failure it will switch over to the second one. It is recommended to enable Bidirectional Forwarding Detection (BFD) when configuring your connections to ensure faster detection and failover. On the other hand, if you have configured a

backup IPsec VPN connection instead, all VPC traffic will failover to the backup VPN connection automatically. Traffic to/from public resources such as Amazon S3 will be routed over the Internet. If you do not have a backup AWS Direct Connect link or a IPsec VPN link, then Amazon VPC traffic will be dropped in the event of a failure.

100.What are the best practices for Security in Amazon EC2?

There are several best practices to secure Amazon EC2. A few of them are given below:

- Use AWS Identity and Access Management (IAM) to control access to your AWS resources.
- Restrict access by only allowing trusted hosts or networks to access ports on your instance.
- Review the rules in your security groups regularly, and ensure that you apply the principle of least
- Privilege – only open up permissions that you require.
- Disable password-based logins for instances launched from your AMI. Passwords can be found or cracked, and are a security risk.

CREDO SYSTEMZ