

Most 70 AWS Cloud Interview Questions and answers

Topics : [AWS](#)

Written on [December 11, 2023](#)

Basics and General AWS Questions

1. What is AWS?

AWS (Amazon Web Services) is a cloud computing platform provided by Amazon that offers a wide range of services like computing power, storage, databases, machine learning, analytics, and more over the internet.

2. What are the key components of AWS?

Key components include Compute, Storage, Database, Networking, Security & Identity, Analytics, Machine Learning, and Internet of Things (IoT).

3. What is the difference between EC2 and S3?

EC2 (Elastic Compute Cloud) is a virtual server for running applications, while S3 (Simple Storage Service) is an object storage service used to store and retrieve data.

4. Explain what Elasticity and Scalability mean in AWS.

Elasticity is the ability to automatically scale resources up or down based on demand. Scalability is the ability to handle varying workloads.

5. What is the AWS Free Tier?

The AWS Free Tier provides limited free access to a wide range of AWS services for 12 months, allowing users to explore and test AWS offerings.

EC2 (Elastic Compute Cloud) Questions

6. What is an EC2 instance?

An EC2 instance is a virtual server in the cloud that can run applications. It provides scalable compute capacity.

7. What is an Amazon Machine Image (AMI)?

An AMI is a pre-configured template used to create EC2 instances. It includes an operating system, application server, and applications.

8. What is an EC2 instance type?

EC2 instance types define the virtual hardware of the host computer used for an EC2 instance. They vary in terms of CPU, memory, storage, and networking capacity.

9. Explain the difference between On-Demand and Reserved Instances.

On-Demand Instances are pay-as-you-go with no upfront cost. Reserved Instances involve a one-time upfront payment for a significant discount over On-Demand pricing.

10. What is an Auto Scaling Group?

An Auto Scaling Group is a group of EC2 instances that automatically adjusts the number of instances based on specified conditions.

S3 (Simple Storage Service) Questions

11. What is S3 and what are its main features?

S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance.

12. What is the maximum size for a single S3 object?

As of my last update, the maximum size for a single S3 object is 5 terabytes.

13. Explain the S3 storage classes.

S3 offers storage classes like STANDARD, INTELLIGENT_TIERING, ONEZONE_IA, GLACIER, and DEEP_ARCHIVE, each with varying durability, availability, and cost.

14. What is a bucket policy in S3?

A bucket policy is a JSON-based access policy that specifies what actions are allowed or denied on the objects within a bucket.

15. What is S3 Transfer Acceleration?

S3 Transfer Acceleration uses the CloudFront globally distributed edge locations to accelerate uploads to S3.

RDS (Relational Database Service) Questions

16. What is Amazon RDS?

Amazon RDS is a managed relational database service that makes it easier to set up, operate, and scale a relational database in the cloud.

17. Which databases engines does Amazon RDS support?

RDS supports various database engines, including MySQL, PostgreSQL, Oracle, SQL Server, and MariaDB.

18. What is Multi-AZ deployment in RDS?

Multi-AZ deployment in RDS involves replicating the database in a different Availability Zone to enhance availability and fault tolerance.

19. Explain the Read Replica in RDS.

Read Replicas are copies of the primary database instance that can be used to offload read traffic from the primary database.

20. What is the purpose of the AWS Database Migration Service (DMS)?

AWS DMS helps migrate databases to and from the cloud seamlessly.

Networking Questions

21. What is Amazon VPC?

Amazon Virtual Private Cloud (VPC) enables users to launch Amazon Web Services resources into a virtual network.

22. What is a subnet in Amazon VPC?

A subnet is a range of IP addresses in your VPC, and it can be public or private.

23. What is an Internet Gateway in Amazon VPC?

An Internet Gateway enables communication between instances in your VPC and the internet.

24. Explain Elastic Load Balancing (ELB).

ELB automatically distributes incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones.

25. What is AWS Direct Connect?

AWS Direct Connect provides dedicated network connections from your on-premises data centers to AWS.

Identity and Access Management (IAM) Questions

26. What is IAM?

Identity and Access Management (IAM) is a service that helps manage access to AWS resources securely.

27. What is an IAM role?

An IAM role is an AWS identity with permission policies that determine what the identity can and cannot do in AWS.

28. What is the difference between IAM users and IAM roles?

IAM users represent a person or service, while IAM roles are identities that you can create and grant permissions to.

29. How do you secure access keys in AWS?

Access keys should be securely stored, and best practices include rotating them regularly and avoiding the use of root account access keys.

30. What is AWS Security Token Service (STS)?

AWS STS enables you to request temporary, limited-privilege credentials for IAM users or federated users.

AWS Lambda Questions

31. What is AWS Lambda?

AWS Lambda is a serverless compute service that runs your code in response to events and automatically manages the computing resources.

32. What languages are supported by AWS Lambda?

AWS Lambda supports multiple programming languages, including Node.js, Python, Ruby, Java, Go, .NET Core, and custom runtimes.

33. What is an event source in AWS Lambda?

An event source is an AWS service or developer-created application that produces events that trigger an AWS Lambda function to run.

34. Explain the concept of cold start in AWS Lambda.

A cold start in AWS Lambda occurs when a function is invoked for the first time or after changes that invalidate the existing execution environment.

35. What is the maximum execution time for an AWS Lambda function?

The maximum execution time for an AWS Lambda function is 15 minutes.

DynamoDB Questions

36. What is Amazon DynamoDB?

DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability.

37. What is the primary key in DynamoDB?

DynamoDB uses a primary key, which can be a single attribute called the partition key or a combination of partition key and sort key.

38. What is the difference between a scan and query operation in DynamoDB?

Scan reads every item in a table, while query retrieves items based on the values of primary key attributes.

39. What is Provisioned Throughput in DynamoDB?

Provisioned Throughput is the capacity reserved for the read and write capacity units for a DynamoDB table.

40. Explain DynamoDB Accelerator (DAX).

DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers fast read performance.

CloudWatch and CloudTrail Questions

41. What is Amazon CloudWatch?

Amazon CloudWatch is a monitoring and observability service that provides data and actionable insights for AWS resources.

42. What is AWS CloudTrail?

AWS CloudTrail is a service that records API calls made on your account and delivers log files to your Amazon S3 bucket.

43. What is the difference between Amazon CloudWatch and AWS CloudTrail?

CloudWatch provides monitoring and observability, while CloudTrail provides audit history of API calls.

44. What are CloudWatch Alarms?

CloudWatch Alarms watch a single metric over a specified time period and perform one or more actions based on the value of the metric relative to a threshold over time.

45. How can you monitor EC2 instance memory usage with CloudWatch?

EC2 instances need to have the CloudWatch agent installed to collect memory usage metrics.

AWS Elastic Beanstalk Questions

46. What is AWS Elastic Beanstalk?

AWS Elastic Beanstalk is a fully managed service that makes it easy to deploy and run applications in multiple languages.

47. What types of applications can be deployed using Elastic Beanstalk?

Elastic Beanstalk supports a variety of applications, including web applications, APIs, and backend services.

48. How does AWS Elastic Beanstalk differ from AWS Lambda?

Elastic Beanstalk is for more traditional applications with server instances, while Lambda is for serverless computing.

49. What is the role of AWS Elastic Beanstalk environment tiers?

Elastic Beanstalk environment tiers include Web Server, Worker, and Multicontainer Docker, providing different configurations based on your application's needs.

50. Can you customize the underlying AWS resources in an Elastic Beanstalk environment?

Yes, you can customize the underlying resources, but Elastic Beanstalk abstracts the infrastructure details for ease of use.

Security and Compliance Questions

51. What is the AWS Shared Responsibility Model?

The Shared Responsibility Model defines the security responsibilities between AWS and the customer. AWS is responsible for the security of the cloud infrastructure, while customers are responsible for securing their data and applications in the cloud.

52. How do you encrypt data at rest in AWS?

AWS offers services like Amazon S3 and Amazon EBS that support encryption of data at rest using AWS Key Management Service (KMS) keys.

53. What is AWS Key Management Service (KMS)?

AWS KMS is a managed service that makes it easy for you to create and control encryption keys used to encrypt your data.

54. What is AWS WAF?

AWS WAF (Web Application Firewall) helps protect web applications from common web exploits by allowing you to configure rules that control access to content.

55. What is AWS Organizations?

AWS Organizations helps consolidate multiple AWS accounts into an organization that you create and centrally manage.

AWS Cost Management Questions

56. What is AWS Pricing Calculator?

AWS Pricing Calculator is a tool that allows you to estimate your monthly bill using pricing information for AWS services.

57. How can you control costs in AWS?

You can control costs in AWS by using AWS Budgets, tagging resources, utilizing reserved instances, and regularly reviewing your architecture for cost optimization.

58. What are AWS Budgets?

AWS Budgets allow you to set custom cost and usage budgets that alert you when you exceed your thresholds.

59. Explain the concept of Reserved Instances.

Reserved Instances involve making a one-time payment to reserve compute capacity in an AWS data center for a specified term and receive a significant discount compared to On-Demand pricing.

60. What is AWS Marketplace?

AWS Marketplace is a digital catalog with thousands of software listings from independent software vendors that make it easy to find, test, buy, and deploy software that runs on AWS.

DevOps and Automation Questions

61. What is AWS CloudFormation?

AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so you can spend less time managing resources and more time focusing on your applications that run in AWS.

62. What is AWS CodeDeploy?

AWS CodeDeploy is a deployment service that automates the process of deploying applications to instances.

63. What is AWS CodePipeline?

AWS CodePipeline is a continuous integration and continuous delivery (CI/CD) service that automates the build, test, and deployment phases of your release process.

64. What is AWS OpsWorks?

AWS OpsWorks is a configuration management service that uses Chef and Puppet to automate how servers are configured, deployed, and managed across your EC2 instances.

65. Explain the use of AWS Elastic Beanstalk in a DevOps environment.

Elastic Beanstalk simplifies the deployment and management of applications in a DevOps environment, providing an easy way to deploy, scale, and update applications.

Machine Learning and AI Questions

66. What is Amazon SageMaker?

Amazon SageMaker is a fully managed service that enables developers and data scientists to

quickly build, train, and deploy machine learning models at scale.

67. What is Amazon Polly?

Amazon Polly is a service that turns text into lifelike speech using deep learning.

68. What is Amazon Rekognition?

Amazon Rekognition is a service for image and video analysis that can identify objects, people, text, scenes, and activities.

69. What is AWS DeepLens?

AWS DeepLens is a deep learning-enabled video camera for developers, allowing them to experiment with deep learning models and build applications with computer vision.

70. Explain the concept of Amazon Machine Learning (AML).

Amazon Machine Learning is a cloud service that enables developers to easily build predictive applications, including fraud detection, demand forecasting, and personalized user experiences.