

# Amazon Web Services Interview Questions

---

## Amazon Web Services Interview Questions

### Q. Explain what is AWS?

AWS stands for Amazon Web Service; it is a collection of remote computing services also known as cloud computing platform. This new realm of cloud computing is also known as IaaS or Infrastructure as a Service.

### Q. Mention what are the key components of AWS?

The key components of AWS are as follows:

- **Route 53:** A DNS web service
- **Simple E-mail Service:** It allows sending e-mail using RESTFUL API call or via regular SMTP
- **Identity and Access Management:** It provides enhanced security and identity management for your AWS account
- **Simple Storage Device or (S3):** It is a storage device and the most widely used AWS service
- **Elastic Compute Cloud (EC2):** It provides on-demand computing resources for hosting applications. It is very useful in case of unpredictable workloads
- **Elastic Block Store (EBS):** It provides persistent storage volumes that attach to EC2 to allow you to persist data past the lifespan of a single EC2

- **CloudWatch:** To monitor AWS resources, It allows administrators to view and collect key Also, one can set a notification alarm in case of trouble.

**Q. Explain what is S3?**

S3 stands for Simple Storage Service. You can use S3 interface to store and retrieve any amount of data, at any time and from anywhere on the web. For S3, the payment model is “pay as you go”.

**Q. Explain what is AMI?**

AMI stands for Amazon Machine Image. It's a template that provides the information (an operating system, an application server, and applications) required to launch an instance, which is a copy of the AMI running as a virtual server in the cloud. You can launch instances from as many different AMIs as you need.

**Q. Mention what is the relation between an instance and AMI?**

From a single AMI, you can launch multiple types of instances. An instance type defines the hardware of the host computer used for your instance. Each instance type provides different compute and memory capabilities. Once you launch an instance, it looks like a traditional host, and we can interact with it as we would with any computer.

**Q. What does an AMI include?**

An AMI includes the following things

- A template for the root volume for the instance
- Launch permissions decide which AWS accounts can avail the AMI to launch instances
- A block device mapping that determines the volumes to attach to the instance when it is launched

**Q. How can you send a request to Amazon S3?**

Amazon S3 is a REST service, you can send a request by using the REST API or the AWS SDK wrapper libraries that wrap the underlying Amazon S3 REST API.

**Q. What is the 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.

### **Q. 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:
- The virtual computing environment (known as instances)
- Pre-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 are 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)

### **Q. 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.

### **Q. What is the 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.

### **Q. 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.

#### **Q. 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 storage provided by AWS.

**Amazon EBS:** Its durable, block-level storage volumes can attach to 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 the 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.

#### **Q. What are the Security Best Practices for Amazon EC2?**

There are several best practices for secure Amazon EC2. The following are a 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.

**Q. 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.

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

EBS is a virtualized SAN or storage area network. That means it is RAID storage to start with, so it's redundant and faults tolerant. If disks die in that RAID you don't lose data. Great! It is also virtualized, so you can provision and allocate storage, and attach it to your server with various API calls. No calling the storage expert and asking him or her to run specialized commands from the hardware vendor.

Performance on EBS can exhibit variability. That is, it can go above the SLA performance level, then drop below it. The SLA provides you with an average disk I/O rate you can expect. This can frustrate some folks, especially performance experts who expect reliable and consistent disk throughput on a server. Traditional physically hosted servers behave that way. Virtual AWS instances do not.

Backup EBS volumes by using the snapshot facility via API call or via a GUI interface like elasticfox.

Improve performance by using Linux software raid and striping across four volumes.

**Q. What is S3? What is it used for? Should encryption be used?**

S3 stands for Simple Storage Service. You can think of it like FTP storage, where you can move files to and from there, but not mount it like a filesystem. AWS automatically puts your snapshots there, as well as AMIs there. Encryption should be considered for sensitive data, as S3 is a proprietary technology developed by Amazon themselves, and as yet unproven vis-a-vis a security standpoint.

**Q. What is an AMI? How do I build one?**

>> AMI stands for Amazon Machine Image. It is effectively a snapshot of the root filesystem. Commodity hardware, servers have bios that point the

master boot record of the first block on a disk. A disk image, though can sit anywhere physically on a disk, so Linux can boot from an arbitrary location on the EBS storage network.

>> Build a new AMI by first spinning up an instance from a trusted AMI. Then adding packages and components as required. Be wary of putting sensitive data onto an AMI. For instance, your access credentials should be added to an instance after spinup with a database, mount an outside volume that holds your MySQL data after spinup as well.

**Q. 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!!

**Q. What is auto-scaling? How does it work?**

Autoscaling is a feature of AWS which allows you to configure and automatically provision and spin up new instances without the need for your intervention. You do this by setting thresholds and metrics to monitor. When those thresholds are crossed, a new instance of your choosing will be spun up, configured, and rolled into the load balancer pool. Voila, you've scaled horizontally without any operator intervention!

**Q. What automation tools can I use to spin up servers?**

The most obvious way is to roll-your-own scripts and use the AWS API tools. Such scripts could be written in bash, Perl or another language or your choice. The next option is to use configuration management and provisioning tools like puppet or better it's successor Opscode Chef. You might also look towards a tool like Scalr. Lastly, you can go with a managed solution such as RightScale.

**Q. What is configuration management? Why would I want to use it with cloud provisioning of resources?**

>> Configuration management has been around for a long time in web operations and systems administration. Yet the cultural popularity of it has been limited. Most systems administrators configure machines as the software was developed before version control – that is manually making changes on servers. Each server can then and usually is slightly different. Troubleshooting though is straightforward as you login to the box and operate on it directly. Configuration management brings a large automation tool in the picture, managing servers like strings of a puppet. This forces standardization, best practices, and reproducibility as all configs are

versioned and managed. It also introduces a new way of working which is the biggest hurdle to its adoption.

>> Enter the cloud, then configuration management becomes even more critical. That's because virtual servers such as Amazon EC2 instances are much less reliable than physical ones. You absolutely need a mechanism to rebuild them as-is at any moment. This pushes best practices like automation, reproducibility and disaster recovery into center stage.

**Q. Explain how you would simulate perimeter security using the Amazon Web Services model?**

Traditional perimeter security that we're already familiar with using firewalls and so forth is not supported in the Amazon EC2 world. AWS supports security groups. One can create a security group for a jump box with ssh access – only port 22 open. From there a webserver group and database group are created. The webserver group allows 80 and 443 from the world, but port 22 \*only\* from the jump box group. Further, the database group allows port 3306 from the webserver group and port 22 from the jump box group. Add any machines to the webserver group and they can all hit the database. No one from the world can, and no one can directly ssh to any of your boxes.

**Q. How to use Amazon SQS?**

Amazon SQS (Simple Queue Service) is a message passing mechanism that is used for communication between different connectors that are connected with each other. It also acts as a communicator between various components of Amazon. It keeps all the different functional components together. This functionality helps different components to be loosely coupled and provides an architecture that is a more failure resilient system.

These are the top Amazon Web Services Interview Questions.

---