

[By OnlineInterviewQuestions.com](http://OnlineInterviewQuestions.com)

[Amazon Redshift Interview Questions](#)

Read Basic Amazon Redshift Interview Questions

We have listed below the very important **Amazon Redshift Interview Questions** that you may face as Amazon Redshift Interview. The interviewer will try their best to judge whether you are well equipped with appropriate knowledge and skills. So, it's your responsibility to ensure that you do the best from your side by preparing these **Amazon Redshift Interview Questions and Answers thoroughly**. These Interview Questions are very helpful for the freshers & Experienced candidates. apart from this, you can also download below the **Amazon Redshift Interview Questions PDF** completely free.

Q1. [What is Amazon Redshift? How it work?](#)

Amazon Redshift is a petabyte-scale data warehouse service in the cloud that is fully managed. This enables you to use your data to acquire new insights for your business and customers.

Q2. [What are the limitations of Amazon Redshift?](#)

some limitations of Amazon Redshift are as follows: Amazon Redshift imposes a limit on the number of tables that you can create in a cluster by node type. An Amazon Redshift table cannot have more than 1,600 columns.

Q3. [Which query language is used by Amazon Redshift?](#)

SQL (Structured Query Language) is used by Amazon Redshift.

Q4. [Explain the architecture of Amazon Redshift?](#)

An Amazon Redshift data repository is a business-class relational database query and administration system. It provides connection of clients with a great number of applications including reporting, business intelligent (BI) and analytics tools.

Amazon Redshift has great storage and excellent query performance with an aggregation of column data storage, massively parallel processing, targeted data compression encoding schemes. It is all about the architecture of Redshift system architecture

Q5. List some Pros and Cons of Amazon Redshift?

Pros of Amazon Redshift

- It offers network isolation.
- It offers result caching.
- It integrates with third-party tools.
- It offers a consistent backup for your data.

Cons of Amazon Redshift

- It does not work as a live app database.
- 2. It is a little behind the times with its Postgre setup.
- 3. Your performance levels decrease as the clusters increase.
- 4. There are no stored procedures available to you in Amazon Redshift.

Q6. What is redshift spectrum?

Redshift Spectrum is one of the popular features of Amazon web services. It allows you to do complex analysis of data that is stored in AWS cloud faster. The redshift spectrum is perfect for a data analyst who is performing on SQL queries in the bucket of Amazon S3. For the purpose of analysis, one can easily load data to the cluster of Amazon redshift from Amazon S3. The two major reasons for using the redshift spectrum are, it is budget-friendly as the pricing is dependable on the size of the cluster and the efforts in loading data are a bit less than any other source. If we talk about the spectrum of redshift it is a kind of bridge for providing an interface between redshift and S3 data.

Q7. What are key differences between SQL Server and Amazon Redshift?

The key difference between the SQL Server the Amazon Redshift is as follows.

Name	Amazon Redshift	SQL Server
Description	Large scale data warehouse service for use with business intelligence tools	Microsofts relational DBMS
Primary Database Model	Relational DBMS	Relational DBMS
Secondary Database Model		Document store Graph DBMS

Q8. What problems have you faced while working with Amazon Redshift?

- Majority of the people facing the problem of the queries which are very slow and take a lot of time

answering it.

- Another problem that is seemed is on the dashboard. The dashboard is too slow.
- Another problem in the Amazon Redshift is “black box”. It is very difficult to observe ‘what’s going on’.

Q9. How do we execute sql file on Redshift?

You can be done this job by using a simple Python script running on an EC2 to set up a JDBC connection to Redshift. When it is done, execute the queries in the your.SQL file

Q10. What are cluster in Redshift? How do I create and delete a cluster in AWS redshift ?

Computing resources in Amazon Redshift data warehouse are called nodes which are further arranged in a group known as a cluster.

This cluster contains at least one database and it works on Amazon Redshift engine.

To create a Cluster, you have to follow these steps: -

- The very first step to create a cluster is open the Amazon ECS console by using this link <https://console.aws.amazon.com/ecs/>.
- After that, you need to select the region to use which you can find from the navigation bar.
- When it is done, select cluster in the navigation panel.
- Then, select Create Cluster can be seen on the Cluster page.
- At last, you should select the selection compatibility which might be networking, EC2 Linux+ networking or EC2 window + networking.

To delete a cluster in AWS, follow these steps: -

The very first step to delete a cluster is to need you to open the Amazon Redshift console by using this link <https://console.aws.amazon.com/redshift/>.

- After that, select the Cluster which you want to remove from the navigation panel
- When it is done, on the Configuration tab of the cluster details page and then select Cluster, and after that select Delete option.
- Before going through the end, you need to do some final steps one of the following in the Delete Cluster dialog box.
- You must choose YES to remove the cluster in creating a snapshot and then take the last snapshot. And then you give the name to that snapshot. And finally, select the delete option.
- Or you can do choose NO to delete in creating a snapshot without the taking final snapshot and then select the delete option.

Q11. What is Amazon Redshift ODBC?

The Amazon Redshift ODBC Driver allows you to connect with live Amazon Redshift data, directly from

applications that support ODBC connectivity. It is also helpful to read, write, and update Amazon Redshift data through a standard ODBC Driver interface.

Q12. How to connect a private Redshift cluster?

By selecting the option NO, you access for your private IP address within the VPC. Bu doing this, you execute the public IP address. Now, the way of its accessing is through the VPC.

One more method most of the people use to connect to a private database is by using the port forwarding by a Bastion server.

Q13. How to Stop/Start Redshift cluster?

You can Start the Redshift cluster by using the following steps:

- In the Redshift Snapshots, select the snapshot of the cluster that you want to restore.
- Select the Restore option on the Dropdown “Actions” Snapshot menu.
- Complete the configuration details, then click the “Restore” button at the bottom right.

You can Stop the Redshift cluster by using the following steps:

- Select the cluster you want to stop from the AWS Console.
- Select the “Delete” option on the Dropdown “Cluster” menu.
- Enter the Snapshot name.
- Click on Stop.

Please Visit OnlineInterviewquestions.com to download more pdfs