

Service Summary Cards (SSC)

Reference:

[FAQs](#)

Category:

Database



Amazon Redshift

What?

- Amazon Redshift is a fully managed, petabyte-scale data warehouse, to analyze data and run complex queries.
- Amazon Redshift uses SQL to analyze large amounts of structured and semi-structured data across data warehouses, operational databases, and data lakes.

Why?

- Focus on getting insights from data in seconds without managing your data warehouse infrastructure.
- Gain up to 3x better price performance than other cloud data warehouses with automation to improve query speed.
- It offers both provisioned and serverless options, without having to manage your data warehouse.

When?

- You want to get real-time insights and predictive analytics on all your data across your operational databases, data lake, data warehouse, and third-party datasets.
- You want benefit of Massively parallel processing (MPP), columnar data storage, and columnar data compression.

Where?

- Amazon Redshift is a regional service, but the Amazon Redshift cluster is created in a Single AZ which consists of a leader node and one or more compute nodes.
- Clusters can also be relocated to alternative AZs without any data loss or application changes.

Who?

- Amazon Redshift service manages all of the work of setting up, operating, scaling a data warehouse, provisioning capacity, monitoring, backing up the cluster, applying patches and upgrades to the Amazon Redshift engine.
- You can use the default database to load data and run queries on your data or can create additional databases as needed.

How?

- As a first step you launch a set of nodes, called an Amazon Redshift cluster. Each cluster runs an Amazon Redshift engine and contains one or more databases.
- Afterwards you can upload your data set and then perform data analysis queries using the same SQL-based and BI tools.

How much?

- Simply pay an hourly rate based on the chosen instance type (Reserved or On-Demand) and number of nodes in your cluster.
- When you choose on-demand pricing, you can use the pause and resume feature to suspend on-demand billing when a cluster is not in use.

More SSCs:

[Click Here](#)

Complete Book

[Click Here](#)

Created by:

[Ashish Prajapati](#)

