

Course Outline of Data Warehousing on AWS Training

WHAT YOU'LL LEARN

- Evaluate the relationship between Amazon Redshift and other Big Data systems
- Evaluate use cases for data warehousing workloads and review real-world implementation of AWS data and analytic services as part of a data warehousing solution
- Choose an appropriate Amazon Redshift node type and size for your data needs
- Understand which security features are appropriate for Amazon Redshift, such as encryption, IAM permissions, and database permissions
- Launch an Amazon Redshift cluster and use the components, features, and functionality to implement a data warehouse in the cloud
- Use other AWS data and analytic services, such as Amazon DynamoDB, Amazon EMR, Amazon Kinesis Firehose, and Amazon S3, to contribute to the data warehousing solution
- Evaluate approaches and methodologies for designing data warehouses
- Identify data sources and assess requirements that affect the data warehouse design
- Design the data warehouse to make effective use of compression, data distribution, and sort methods
- Load and unload data and perform data maintenance tasks
- Write queries and evaluate query plans to optimize query performance
- Configure the database to allocate resources such as memory to query queues and define criteria to route certain types of queries to your configured query queues for improved processing

- Audit, monitor, and receive event notifications about activities in the data warehouse by using features and services such as Amazon Redshift database audit logging, Amazon CloudTrail, Amazon CloudWatch, and Amazon Simple Notification Service (Amazon SNS)
- Prepare for operational tasks such as resizing Amazon Redshift clusters and using snapshots to back up and restore clusters
- Use a BI application to perform data analysis and visualization tasks against your data

Day 1

- Introduction to Data Warehousing
- Introduction to Amazon Redshift
- Launching Clusters

Day 2

- Designing the Database Schema
- Identifying Data Sources
- Loading Data

Day 3

- Writing Queries and Tuning Performance
- Amazon Redshift Spectrum
- Maintaining Clusters
- Analyzing and Visualizing Data