

Data Warehousing on AWS

Course 3582 – 24 Hours

Overview

Data Warehousing on AWS introduces you to concepts, strategies, and best practices for designing a cloud-based data warehousing solution using Amazon Redshift, the petabyte-scale data warehouse in AWS. This course demonstrates how to collect, store, and prepare data for the data warehouse by using other AWS services such as Amazon DynamoDB, Amazon EMR, Amazon Kinesis, and Amazon S3. Additionally, this course demonstrates how to use Amazon QuickSight to perform analysis on your data.

On Completion, Delegates will be able to

- Discuss the core concepts of data warehousing.
- Discuss the intersection between data warehousing and big data solutions.
- 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, and Amazon S3, to contribute to the data warehousing solution.
- Evaluate approaches and methodologies for designing data warehouses.
- Identify data sources and determine requirements for accessing the data.
- Architect the data warehouse.
- Use important commands, such as COPY, UNLOAD, and VACUUM, to manage the data in the data warehouse.
- Identify performance issues, optimize queries, and tune the database for better performance
- Use Amazon Redshift Spectrum to analyze data directly from an Amazon S3 bucket.
- Use features and services, such as Amazon Redshift database auditing, Amazon CloudWatch, Amazon CloudTrail, and Amazon Simple Notification Service (Amazon SNS), to monitor and audit the data warehouse.
- Use Amazon QuickSight to perform data analysis and visualization tasks against the data warehouse.

Who Should Attend

- Database architects
- Database administrators
- Database developers
- Data analysts and scientists

Prerequisites

- Courses taken: 3550 - AWS Technical Essentials (or equivalent experience with AWS)
- Familiarity with relational databases and database design concepts

Course Contents

- Course Introduction
- Introduction to Data Warehousing
- Introduction to Amazon Redshift
- Understanding Amazon Redshift Components and Resources
- Launching an Amazon Redshift Cluster
- Choosing a Data Warehousing Approach
- Identifying Data Sources and Requirements
- Architecting the Data Warehouse
- Loading Data into the Data Warehouse
- Optimizing Queries and Tuning Performance
- Monitoring and Auditing the Data Warehouse
- Maintaining the Data Warehouse
- Analyzing and Visualizing Data