+33 (0)1 48 10 09 46
formation@ambient-it.net
140 Avenue Jean Lolive 93500 Pantin - France

Updated on 23/07/2025

Sign up

# Redshift Data Warehousing Training

3 days (21 hours)

## Overview

Master Amazon Redshift in its expert dimension with this advanced training course, designed for demanding data architects and engineers. You'll learn how to exploit the full potential of Redshift data warehousing, optimizing performance, industrializing your pipelines, and guaranteeing security, governance and scalability in high-level cloud architectures.

The training begins with a deep dive into Redshift's MPP architecture, node types, distribution, sorting and compression mechanisms, and best practices for boosting your SQL queries and exploiting execution plans.

You'll learn how to build efficient ingestion pipelines, combining COPY, S3, Kinesis, Glue and dbt, while applying advanced security strategies (IAM, RLS, CLS) and auditing via system views.

Industrialization is at the heart of the program: CI/CD with Terraform, modeling with dbt, promotion between environments, CloudWatch integration and performance diagnostics become your everyday tools.

As with all our training courses, this one will be presented with the latest updates from redshift data.

## Objectives

- Understand Redshift's advanced architecture and its specific features MPP, columnar storage and scaling
- Design, optimize and audit high-performance data models in Redshift
- Master strategies for ingesting massive data via COPY, Kinesis, Glue or Spectrum
- Be able to secure, govern and trace access and processing on critical Redshift environments

- Industrialize Redshift deployments with Terraform and automate processing via dbt and CI/CD
- Understand and exploit Redshift Data Sharing, Lakehouse mode and hybrid
  hybrid analytic architectures

## Target audience

- Data Engineers
- Redshift administrators

## Prerequisites

- Master the fundamentals of the AWS cloud

## Redshift Data Warehousing training program Reminders and advanced Redshift developments

- Redshift MPP architecture: deep dive
- Redshift RA3 vs DS2: impact on performance
- Redshift Spectrum, AQUA, Concurrency Scaling, Data Sharing
- Integration with S3, Glue, Athena

## Advanced storage management

- Columnar storage and compression
- Optimal sorting (SORTKEY) and distribution (DISTKEY)
- Vacuum and Analyze management
- Storage monitoring

## Performance optimization

- Execution plan analysis
- SQL tuning and anti-patterning
- View materialization: strategies and use cases

- Logical partitioning via UNLOAD/EXTERNAL TABLES
- Practical workshop: Optimizing an existing cluster, refactoring slow queries

## Mass loading and data pipelines

- Optimized COPY from S3, Kinesis or Glue
- Error and recovery management
- Redshift Streaming Ingestion
- ELT with Redshift + dbt

## Governance, security and access

- Row- and column-level security (RLS/CLS)
- IAM vs. native authentication
- Audit with STL/ SVL / SVV tables
- Integration with Lake Formation and AWS KMS

## Monitoring, diagnostics and auditing

- Use of STL_ALERT_EVENT_LOG, SVL_S3LOG
- Integration with CloudWatch, EventBridge, Lambda
- Audit connections, slow requests, contention
- Workshop: Creating a secure ingestion pipeline + monitoring dashboards

## Redshift Data Sharing and multi-cluster architecture

- Redshift Data Sharing between accounts
- Multi-team, multi-environment use cases
- Dataset security and isolation

## CI/CD and industrialization

- Redshift deployment with Terraform / CloudFormation
- dbt for modeling, documentation and testing
- Versioning, code review, promotion between environments

## Advanced use cases and hybrid architecture

- Lakehouse with Redshift + Iceberg
- Redshift + Apache Hudi / Delta Lake
- Query federation with external sources
- Final workshop: Implementation of a complete project Data Lake + Redshift + CI/CD + Monitoring

## Companies concerned

This course is aimed at both individuals and companies, large or small, wishing to train their teams in a new advanced IT technology, or to acquire specific business knowledge or modern methods.

## Positioning on entry to training

Positioning at the start of training complies with Qualiopi quality criteria. As soon as registration is finalized, the learner receives a self-assessment questionnaire which enables us to assess his or her estimated level of proficiency in different types of technology, as well as his or her expectations and personal objectives for the training to come, within the limits imposed by the selected format. This questionnaire also enables us to anticipate any connection or security difficulties within the company (intra-company or virtual classroom) which could be problematic for the follow-up and smooth running of the training session.

## Teaching methods

Practical training: 60% hands-on, 40% theory. Training material distributed in digital format to all participants.

## Organization

The course alternates theoretical input from the trainer, supported by examples, with brainstorming sessions and group work.

## Validation

At the end of the session, a multiple-choice questionnaire verifies the correct acquisition of skills.

## Certification

A certificate will be awarded to each trainee who completes the training course.