



Architecture Diagrams

Automated Archival for Amazon Redshift



Automated Archival for Amazon Redshift: Architecture Diagrams

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Automated Archival **i**

Automated Archival for Amazon Redshift 1

Further reading 2


Diagram history 2

Automated Archival for Amazon Redshift

Publication date: July 12, 2021 ([Diagram history](#))

This architecture shows how to automate the periodic data archival process for an [Amazon Redshift](#) database. The cold data remains available instantly, and you can join it with existing datasets in the Amazon Redshift cluster using Amazon Redshift Spectrum.

Automated Archival for Amazon Redshift



The following steps describe the architecture:

1. Ingest data into the Amazon Redshift cluster at various frequencies.
2. After every ingestion load, create a queue of metadata about tables populated into Amazon Redshift tables in Amazon RDS. Data engineers can also create the archival queue manually.
3. Use [Amazon EventBridge](#) to trigger an [AWS Lambda](#) function periodically. The function reads the queue from the RDS table and creates an [Amazon Simple Queue Service](#) message for every period due for archival.
4. A proxy Lambda function dequeues the Amazon SQS messages and invokes [AWS Step Functions](#) for every message.
5. Step Functions unloads the data from the Amazon Redshift cluster into an [Amazon Simple Storage Service](#) bucket for the given table and period.

6. Amazon S3 Lifecycle configuration moves data in buckets from S3 Standard storage class to S3 Glacier storage class after 90 days.
7. Amazon S3 inventory tool generates manifest files from the Amazon S3 bucket dedicated for cold data on a daily basis and stores them in an Amazon S3 bucket for manifest files.
8. Every time an inventory manifest file is created in a manifest Amazon S3 bucket, a Lambda function triggers through an Amazon S3 event notification.
9. A Lambda function normalizes the manifest file for easy consumption in the event of restore.
10. Query the data stored in the Amazon S3 bucket for cold data using Amazon Redshift Spectrum.

Further reading

For additional information, refer to the following resources:

- [AWS Architecture Icons](#)
- [AWS Architecture Center](#)
- [AWS Well-Architected](#)

Diagram history

To be notified about updates to this reference architecture diagram, subscribe to the RSS feed.

Change	Description	Date
Initial publication	Reference architecture diagram first published.	July 12, 2021

Note

To subscribe to RSS updates, you must have an RSS plugin enabled for the browser you are using.