How-To Library

Informatica

# Configuring AWS IAM Authentication for Amazon Redshift and Amazon Redshift V2 Connectors

# Abstract

You can use AWS Identity and Access Management (IAM) to control individual and group access to Amazon Redshift resources. You can configure AWS IAM to run tasks on the Secure Agent that is installed on the EC2 system. This article describes the guidelines to configure IAM Authentication for Amazon Redshift and Amazon Redshift V2 Connectors.

# Supported Versions

- Informatica Cloud® Data Integration Amazon Redshift Spring 2018
- Informatica Cloud® Data Integration Amazon Redshift V2 Spring 2018

# Table of Contents



# Overview

To control the access of Amazon Redshift resources, you can define permissions to the users by configuring Amazon Identity and Access Management (IAM). The Amazon IAM service provides enhanced security.

Perform the following steps to configure Amazon IAM authentication:

1. Create a minimal Amazon IAM policy.
2. Create the Amazon EC2 Role.
3. Create the Amazon Redshift Role.
4. Associate the Amazon Redshift Role with the Redshift cluster.
5. Create an Amazon Redshift or Amazon Redshift V2 connection.
6. Create a mapping.

# Create a Minimal Amazon IAM Policy

You can configure the minimal Amazon IAM policy through the AWS console.

You can use the following minimum required actions when you use Amazon Redshift Connector and Amazon Redshift V2 Connector to successfully read data from and write data to Amazon Redshift resources:

- PutObject
- GetObject

- DeleteObject
- ListBucket
- GetBucketPolicy

  **Note:** Do not add the GetBucketPolicy permission in the Amazon IAM policy when you use Amazon Redshift V2 Connector. Amazon Redshift V2 Connector does not support the GetBucketPolicy permission.

The following snippet shows a sample Amazon IAM policy for Amazon Redshift Connector:

```
{
"Version": "2012-10-17", "Statement": [
{ "Effect": "Allow", "Action": [ "s3:PutObject", "s3:GetObject","s3:DeleteObject",
"s3:ListBucket", "s3:GetBucketPolicy" ],
"Resource": [ "arn:aws:s3:::<bucket_name>/*", "arn:aws:s3:::<bucket_name>" ] }
]
}
```

The following snippet shows a sample Amazon IAM policy for Amazon Redshift V2 Connector:

```
{
"Version": "2012-10-17", "Statement": [
{ "Effect": "Allow", "Action": [ "s3:PutObject", "s3:GetObject", "s3:DeleteObject",
"s3:ListBucket" ],
"Resource": [ "arn:aws:s3:::<bucket_name>/*", "arn:aws:s3:::<bucket_name>" ] }
]
}
```

## Create the Amazon EC2 Role

You must create an Amazon EC2 Role to provide users access to the Redshift resources. You can use the Amazon EC2 Role when you create an EC2 system.

1. Log in to the **AWS Console**.
2. Click **Dashboard** from the left panel.

   The **AWS Service** dashboard page appears.
3. Click **IAM**.

   The **Welcome to Identity and Access Management** page appears.
4. Click **Policies** from the left panel.

   The **Policy** page appears.
5. Click **Create Policy** or select the required existing Amazon S3 Policy.

   You can edit or review the policy.
6. Select **Role** from the left panel and click **Create role**.

   The **Create role** page appears.
7. Select **EC2** under the **Choose the service that will use this role** section.
8. Select the required **Amazon EC2** role type under the **Select your use case** section.
9. Click **Next: Permission**.
10. Select the required **Amazon S3 Policy** in the **Attach Permission Policies** page.
11. Click **Next: Review**.
12. Specify the name of the role in the **Create role** review page.
13. Click **Create Role**.
14. Review the Role ARN, Instance Profile ARNs, and Policy values in the **Summary** page.

After you create the Amazon EC2 Role, create an EC2 instance. Assign the Amazon EC2 Role to the EC2 instance.

For more information about creating an EC2 instance and assigning an Amazon EC2 Role to the Amazon EC2 instance, see the Amazon Redshift documentation.

## Create the Amazon Redshift Role

You must create the Amazon Redshift Role Resource Name (ARN) for secure access to Amazon Redshift resources.

1. Log in to the **AWS Console**.
2. Click **Dashboard** from the left panel.

   The **AWS Service** dashboard page appears.
3. Click **IAM**.

   The **Welcome to Identity and Access Management** page appears.
4. Click **Policies** from the left panel.

   The **Policy** page appears.
5. Click **Create Policy** or select the required existing Amazon S3 Policy.

   You can edit or review the policy.
6. Select **Role** from the left panel and click **Create role**.

   The **Create role** page appears.
7. Select **Redshift** under the **Choose the service that will use this role** section.
8. Select the required Amazon Redshift cluster under the **Select your use case** section.
9. Click **Next: Permission**.
10. Select the required **Amazon S3 Policy** in the **Attach Permission Policies** page.
11. Click **Next: Review**.
12. Specify the name of the role in the **Create role** review page.
13. Click **Create Role**.
14. Review the Role ARN, Instance Profile ARNs, and Policy values in the **Summary** page.

The following example shows a sample Amazon Redshift Role Resource Name (ARN):

```
arn:aws:iam::123123456789:role/redshift_write
```

You must assign the Amazon Redshift Role that you created to the Amazon Redshift cluster to successfully perform the read and write operations.

## Associate the Amazon Redshift Role with the Redshift Cluster

You must associate the Amazon Redshift Role Resource Name (ARN) with an Amazon Redshift cluster to read data from Amazon Redshift and write data to the Amazon S3 bucket.

1. Log in to the **AWS Console**.
2. Click **Dashboard** from the left panel.

   The **AWS Service** dashboard page appears.
3. Click **Amazon Redshift**.

   The **Redshift** dashboard page appears.
4. Click **Clusters** under **Resources**.
5. Select the required Amazon Redshift cluster.

6. Click **Manage IAM roles**.

   The **Manage IAM roles** dialog box displays.

7. Select the required Amazon Redshift Role. For example, `arn:aws:iam::123123456789:role/redshift_write`

8. Click **Apply changes**.

After you associate the Amazon Redshift Role Resource Name (ARN) with the Amazon Redshift cluster, install the Secure Agent on the EC2 instance.

For more information about installing a Secure Agent, see *Informatica Cloud Data Integration online help*.

# Create a Connection

You can create an Amazon Redshift or Amazon Redshift V2 connection. Specify the connection properties for configuring the IAM authentication to control secure access of Amazon Redshift resources.

## *Create an Amazon Redshift Connection*

When you set up an Amazon Redshift connection, you must configure the connection properties.

To run a mapping on Secure Agent installed on an EC2 system, you must not provide the **Access Key ID** and **Secret Access Key** when you create an Amazon Redshift connection.
The following image shows the sample values in the Amazon Redshift connection properties:

**Connection Details**

| | |
|---|---|
| Connection Name: | AmazonRedshift_IAM_Connection |
| Description: | Connection to write to an Amazon Redshift target. |
| Type: | AmazonRedshift New (Informatica Cloud) |
| Created On: | May 3, 2018 4:06:35 AM |
| Updated On: | May 3, 2018 4:07:01 AM |
| Created By: | |
| Updated By: | |

**AmazonRedshift New Connection Properties**

| | |
|---|---|
| Runtime Environment: | WIN-PH8J5A6ECRK |
| Username: | Sample_User_Name |
| Password: | ******** |
| Schema: | public |
| AWS Access Key ID: | |
| AWS Secret Access Key: | |
| Master Symmetric Key: | |
| Customer Master Key ID: | |
| Jdbc URL: | jdbc:redshift://infa-rs-qa-cluster.ca8dsxvs46sw.us-west- |
| Number of bytes needed to support multibytes for varchar: | 1 |

The Secure Agent uses the **Username**, **Password**, and **JDBC URL** properties to validate the connection.

Provide the Amazon Redshift Role Resource Name (ARN) in the **AWS_IAM_ROLE** option in the **UnloadOptions Property File** and **CopyOptions Property File** properties when you create a task. The Secure Agent uses the Amazon Redshift Role Resource Name (ARN) associated with the IAM Role to access the data from the Amazon Redshift target. When you run the mapping task, the Secure Agent validates the IAM policy.

### *Create an Amazon Redshift V2 Connection*

When you set up an Amazon Redshift V2 connection, you must configure the connection properties.

To run a mapping on Secure Agent installed on an EC2 system, you must not provide the **Access Key ID** and **Secret Access Key** when you create an Amazon Redshift V2 connection.

The following image shows the sample values in the Amazon Redshift V2 connection properties:

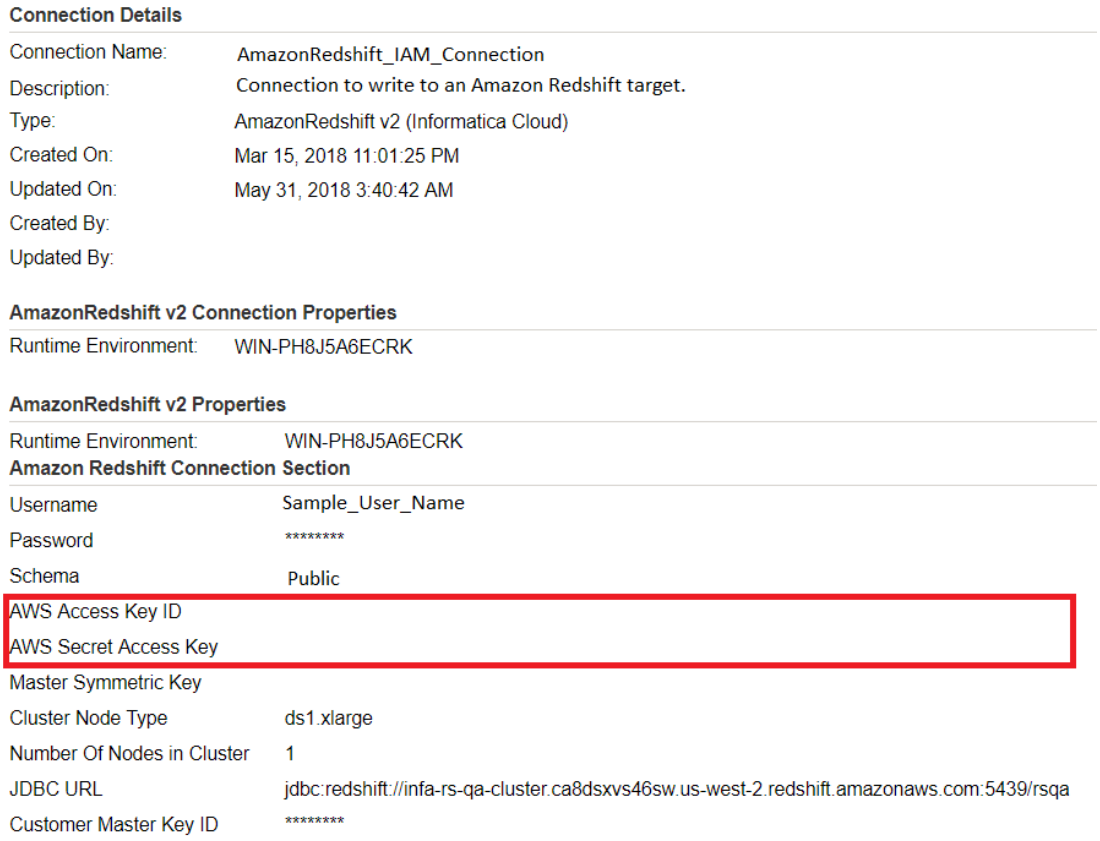

| **Connection Details** | |
|---|---|
| Connection Name: | AmazonRedshift_IAM_Connection |
| Description: | Connection to write to an Amazon Redshift target. |
| Type: | AmazonRedshift v2 (Informatica Cloud) |
| Created On: | Mar 15, 2018 11:01:25 PM |
| Updated On: | May 31, 2018 3:40:42 AM |
| Created By: | |
| Updated By: | |
| **AmazonRedshift v2 Connection Properties** | |
| Runtime Environment: | WIN-PH8J5A6ECRK |
| **AmazonRedshift v2 Properties** | |
| Runtime Environment: | WIN-PH8J5A6ECRK |
| **Amazon Redshift Connection Section** | |
| Username | Sample_User_Name |
| Password | ******** |
| Schema | Public |
| AWS Access Key ID | |
| AWS Secret Access Key | |
| Master Symmetric Key | |
| Cluster Node Type | ds1.xlarge |
| Number Of Nodes in Cluster | 1 |
| JDBC URL | jdbc:redshift://infa-rs-qa-cluster.ca8dsxvs46sw.us-west-2.redshift.amazonaws.com:5439/rsqa |
| Customer Master Key ID | ******** |

The Secure Agent uses the **Username**, **Password**, and **JDBC URL** properties to validate the connection.

Provide the Amazon Redshift Role Resource Name (ARN) in the **AWS_IAM_ROLE** option in the **UnloadOptions Property File** and **CopyOptions Property File** properties when you create a task. The Secure Agent uses the Amazon Redshift Role Resource Name (ARN) associated with the IAM Role to access the data from the Amazon Redshift target. When you run the mapping task, the Secure Agent validates the IAM policy.

## Create a Mapping

Create a mapping to read data from an on-premise MySQL database and write data to an Amazon Redshift target for analysis. Configure AWS IAM authentication for secure and controlled access to Amazon Redshift resources when you run the mapping.

1. In Data Integration, click **New** > **Mappings** > **Create**.

   The **New Mapping** dialog box appears.
2. Enter a name, location, and description for the mapping.
3. On the Source transformation, specify a name and description in the general properties.

4. On the **Source** tab, perform the following steps to provide the source details to read data from the MySQL source:
    1. In the **Connection** field, select the MySQL source connection.
    2. In the **Source Type** field, select the type of the source.
    3. In the **Object** field, select the required object.
    4. In the **Advanced Properties** section, provide the appropriate values.
5. On the **Fields** tab, map the MySQL source fields to the target fields.
6. On the Target transformation, specify a name and description in the general properties.
7. On the **Target** tab, perform the following steps to provide the target details to write data to the Amazon Redshift target:
    1. In the **Connection** field, select the required target connection.
    2. In the **Target Type** field, select the type of the target.
    3. In the **Object** field, select the required object.
    4. In the **Operation** field, select the required operation.
    5. In the **Advanced Properties** section, provide appropriate values for the advanced target properties and ensure that you specify the AWS IAM ROLE that you created in the **CopyOptions Property File** field. The Secure Agent writes the data to the required target when you specify the Amazon Redshift Role Resource Name (ARN) in the **CopyOptions Property File** field.

       **Note:** Verify that the Amazon S3 bucket and Amazon Redshift cluster reside in the same region.
8. Map the MySQL source and the Amazon Redshift target.

   The following image shows a sample mapping:

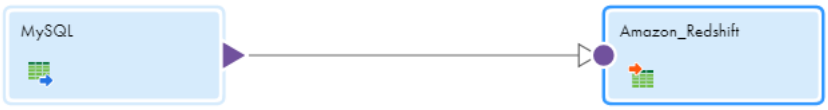


9. Click **Save** > **Run** to validate the mapping.

   In **Monitor**, you can monitor the status of the logs after you run the task.
10. Click **Action** > **New Mapping Task** on the left corner of the task wizard.

    The **Mapping Task** page appears.
11. Provide a name of the mapping task and select the runtime environment.

    The mapping that you created is selected automatically.
12. Click **Save** > **Run** to run the mapping task.

# Authors

**Salam Subhashree**

**Chanchal Das**

**Shivaprasad Yallappagoudar**