Configuring AWS S3 Integration for Union Pacific Data Sharing

This document outlines the necessary steps to configure your AWS S3 cloud data storage to enable Union Pacific's secure data sharing service via Snowflake. By following these instructions, you will grant Union Pacific's Snowflake environment the required permissions to deliver data within your designated AWS cloud S3 bucket.

Contents

Configuring Secure Access 2

Amazon S3 Bucket Integration Steps 2

Step 1: Create an S3 Bucket (or use an existing one) 2

Step 2: Create an IAM Policy for Snowflake Access 2

Step 3: Create an IAM Role for Snowflake access 4

Step 4: Provide Union Pacific with Your S3 Bucket Details 5

Step 5: Union Pacific will provide you with an IAM Role ARN 5

Step 6: Grant the IAM User Permissions to Access Bucket 5

Step 7: Union Pacific Confirms Access 8

References: 8

Configuring Secure Access

To securely write to an S3 bucket, it is crucial to ensure that the security and access management policies on the bucket are configured to grant Union Pacific's Snowflake the necessary permissions. This involves setting up a storage integration object in Union Pacific's Snowflake, which delegates authentication responsibilities to a Snowflake-specific identity and access management (IAM) entity. By implementing this method, we are eliminating the need to provide AWS IAM credentials directly when creating stages or loading data, thereby reducing the risk of credential exposure. Additionally, you can apply fine-grained access controls and monitor access logs to enhance security further.

Amazon S3 Bucket Integration Steps

This section describes the steps to configure your Amazon S3 bucket to allow Union Pacific's Snowflake to securely deliver data via Snowflake's Storage Integration feature.

Step 1: Create an S3 Bucket (or use an existing one)

Choose an existing Amazon S3 bucket where you would like Union Pacific to deliver the shared data. If you don't have a bucket, create a new one.

Make note of the **Region** where your chosen S3 bucket is located. You will need to provide this information to Union Pacific.

Step 2: Create an IAM Policy for Snowflake Access

These steps describe how to configure access permissions for Union Pacific's Snowflake in your AWS Management Console so that an S3 bucket can be configured to load and unload data

- Navigate to the AWS IAM service in your AWS Management Console.
- Under Account settings, select Security Token Service (STS) in the Endpoints list, find the Snowflake region where your account is located. If the STS status is inactive, move the toggle to Active.
- Click on **Policies** in navigation pane and then click **Create policy**.
- Select the **JSON** tab and paste the following policy document into the editor. Add a policy document that will allow Snowflake to access the S3 bucket and folder. Be sure to replace <your-bucket-name> with the actual name of your S3 bucket. Select Create policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {  "Effect": "Allow",
       "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:GetObjectVersion",
        "s3:DeleteObject",
        "s3:DeleteObjectVersion"
       ],
       "Resource": "arn:aws:s3:::<bucket>/<prefix>/*"
    },
    {
       "Effect": "Allow",
       "Action": [
          "s3:ListBucket",
          "s3:GetBucketLocation"
       ],
       "Resource": "arn:aws:s3:::<bucket>",
       "Condition": {
          "StringLike": {
             "s3:prefix": [
                "<prefix>/*"
             ]
          }
       }
    }
  ]
}
```

Step 3: Create an IAM Role for Snowflake access

To configure access permissions for Snowflake, please follow the instructions:

- In the IAM service, click on **Roles** and then click **Create role**.
- For the "Select the type of trusted entity", choose **AWS account**. And “Another AWS account”
- In the Account ID field, enter your own AWS account ID temporarily. Later, you modify the trust relationship and grant access to Snowflake.
- Select the Require external ID option. Enter a placeholder ID such as 0000. In a later step, you will modify the trust relationship for your IAM role and specify the external ID for your storage integration.
- Attach the IAM policy you created earlier.
- Enter a name and description for your role and click **Create role**.

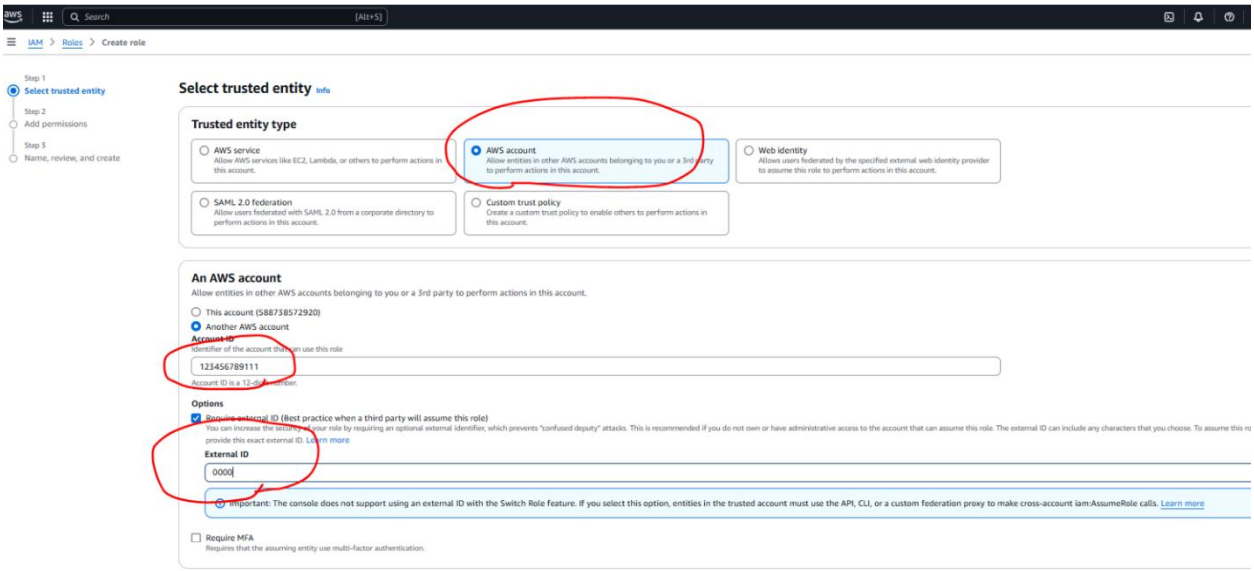

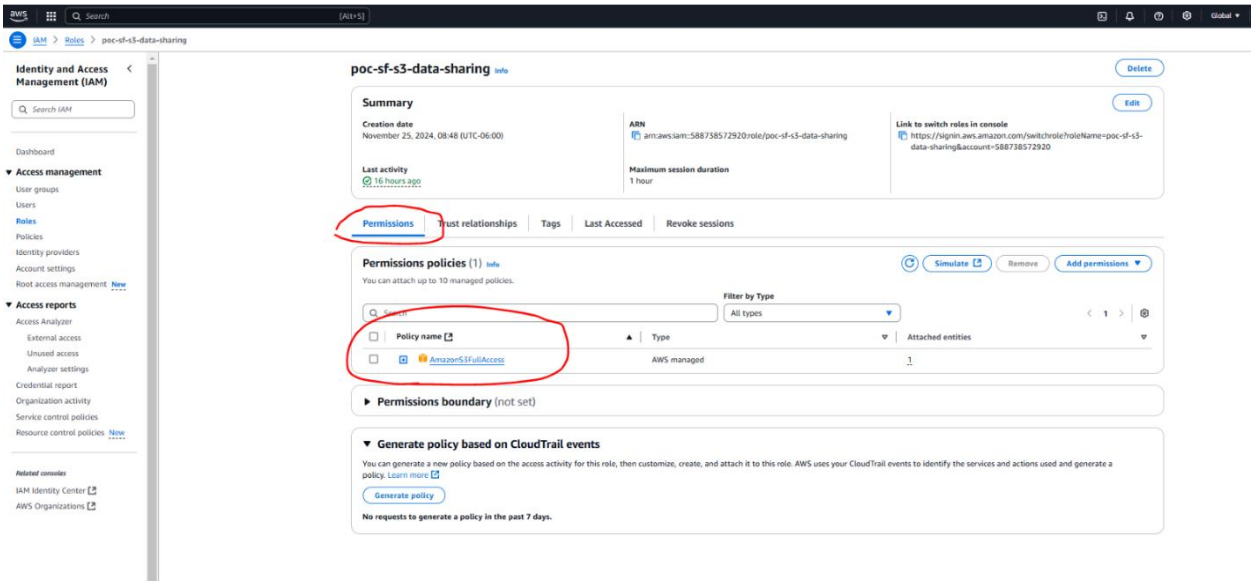

- You have now created an IAM policy for a bucket, created an IAM role, and attached the policy to the role.
- On the role summary page, locate and record the Role ARN value.

Step 4: Provide Union Pacific with Your S3 Bucket Details

You will need to contact and inform Union Pacific of the following:

- The name of the S3 bucket where you want to receive the data.
- The AWS region where your bucket is located (e.g. us-east-1).
- Amazon Resource Name (ARN) of the IAM role you created

Step 5: Union Pacific will provide you with an IAM Role ARN

- Union Pacific will create a Snowflake Storage Integration using the information you provided in their Snowflake account.
- As part of this process, Union Pacific will now contact you to provide following details:
    - snowflake_user_arn
    - snowflake_external_id

Step 6: Grant the IAM User Permissions to Access Bucket

Now configure IAM access permissions for Snowflake in your AWS Management Console so that Union Pacific can use S3 bucket to load data.

- Go to IAM Role and select Roles
- Select the Trust relationships tab

- Select Edit trust policy.
- Modify the policy document with the values you received from Union Pacific

Policy document for IAM role

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "",
      "Effect": "Allow",
      "Principal": {
        "AWS": "<snowflake_user_arn>"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId": "<snowflake_external_id>"
        }
      }
    }
  ]
}
```

- Replace following values you received from Union Pacific in **Step 5** and select update policy
    - snowflake_user_arn
    - snowflake_external_id

Step 7: Union Pacific Confirms Access

Union Pacific will perform tests to verify that Snowflake environment can successfully access your S3 bucket using the configured IAM role and External ID.

References:

https://docs.snowflake.com/en/user-guide/data-load-s3-config-storage-integration

https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html