Set up resources with AWS CloudFormation

We provide three AWS CloudFormation templates in this post: for the producer account, central account, and consumer account. Deploy the CloudFormation templates in the order of producer, central, and consumer, because there are dependencies between the templates.

The CloudFormation template for the central account generates the following resources:

- Two IAM users:
  - `DataMeshOwner`
  - `ProducerSteward`
- Grant `DataMeshOwner` as the LakeFormation Admin
- One IAM role:
  - `LFRegisterLocationServiceRole`
- Two IAM policies:
  - `ProducerStewardPolicy`
  - `S3DataLakePolicy`
- Create databases "`credit-card`" for `ProducerSteward` to manage Data Catalog
- Share the data location permission for producer account to manage Data Catalog

The CloudFormation template for the producer account generates the following resources:

- Two Amazon Simple Storage Service (Amazon S3) buckets:
  - `credit-card`, which holds one table:
    - `Credit_Card`
- Allow Amazon S3 bucket access for the central account Lake Formation service role.
- One AWS Glue crawler
- One AWS Glue crawler service role
- Grant permissions on the S3 bucket locations `credit-card-lf-<ProducerAccountID>-<aws-region>` to the AWS Glue crawler role
- One producer steward IAM user

The CloudFormation template for the consumer account generates the following resources:

- One S3 bucket:
  - `<AWS Account ID>-<aws-region>-athena-logs`
- One Athena workgroup:
  - `consumer-workgroup`
- Three IAM users:
  - `ConsumerAdmin`

Launch the CloudFormation stack in the central account

To create resources in the central account, complete the following steps:

1. Sign in to the central account's AWS CloudFormation console in the target Region.
2. Choose **Launch Stack:**
   Launch Stack
3. Choose **Next**.
4. For **Stack name**, enter `stack-central`.

5. For **DataMeshOwnerUserPassword**, enter the password you want for the data lake admin IAM user in the central account.
6. For **ProducerStewardUserPassword**, enter the password you want for the producer steward IAM user in the producer account.
7. For **ProducerAWSAccount**, enter the AWS `<ProducerAccountID>`.
8. Choose **Next**.
9. On the next page, choose **Next**.
10. Review the details on the final page and select **I acknowledge that AWS CloudFormation might create IAM resources**.
11. Choose **Create stack**.
12. Collect the value for `LFRegisterLocationServiceRole` on the stack's Outputs tab.

Launch the CloudFormation stack in the producer account

To set up resources in the producer account, complete the following steps:

1. Sign in to the producer account's AWS CloudFormation console in the target Region.
2. Choose **Launch Stack**:
   Launch Stack
3. Choose **Next**.
4. For **Stack name**, enter `stack-producer`.
5. For **CentralAccountID**, copy and paste the value of the `<CentralAccountID>` .
6. For **CentralAccountLFServiceRole**, copy and paste the value of the `LFRegisterLocationServiceRole` collected from the stack-central.
7. For **LFDatabaseName**, keep the default value of the `lf-ml` database name.
8. For **ProducerStewardUserPassword**, enter the password you want for the data lake admin IAM user on the producer account.
9. Choose **Next**.
10. On the next page, choose **Next**.
11. Review the details on the final page and select **I acknowledge that AWS CloudFormation might create IAM resources**.
12. Choose **Create stack.**

Launch the CloudFormation stack in the consumer account

To create resources in the consumer account, complete the following steps:

1. Sign in to the consumer account's AWS CloudFormation console in the target Region.
2. Choose **Launch Stack**:
   Launch Stack
3. Choose **Next**.
4. For **Stack name**, enter `stack-consumer`.
5. For **ConsumerAdminUserName** and **ConsumerAdminUserPassword**, enter the user name and password you want for the data lake admin IAM user.
6. For **ConsumerAnalyst1UserName** and **ConsumerAnalyst1UserPassword**, enter the user name and password you want for the `consumeranalyst1` IAM user.
7. For **ConsumerAnalyst2UserName** and **ConsumerAnalyst2UserPassword**, enter the user name and password you want for the `consumeranalyst2` IAM user.

8. Choose **Next**.
9. On the next page, choose **Next**.
10. Review the details on the final page and select **I acknowledge that AWS CloudFormation might create IAM resources**.
11. Choose **Create stack**.

Configure Lake Formation cross-account sharing

After you create your resources with AWS CloudFormation, you perform the following steps in the producer and central account to set up Lake Formation cross-account sharing.

Central governance account

In the central account, complete the following steps:

1. Sign in to the Lake Formation console as admin.
2. In the navigation pane, choose **Permissions**, then choose **Administrative roles and tasks**.

The CloudFormation template added the data mesh owner as the data lake administrator.

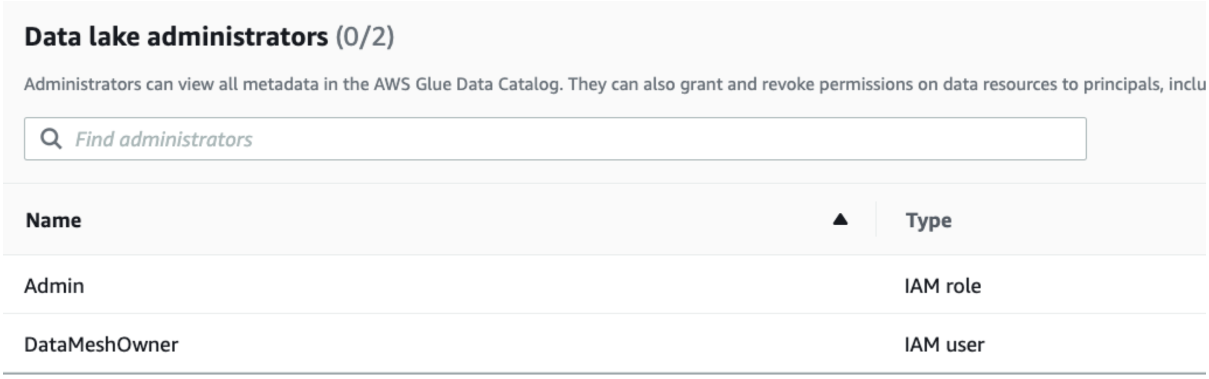

Data lake administrators (0/2)

Administrators can view all metadata in the AWS Glue Data Catalog. They can also grant and revoke permissions on data resources to principals, inclu

Name	Type
Admin	IAM role
DataMeshOwner	IAM user

Next, we update the Data Catalog settings to use Lake Formation permissions to control catalog resources instead of IAM-based access control.

3. In the navigation pane, under **Data catalog**, choose **Settings**.
4. Uncheck **Use only IAM access control for new databases**.
5. Uncheck **Use only IAM access control for new tables in new databases**.

6. Choose **Save**.

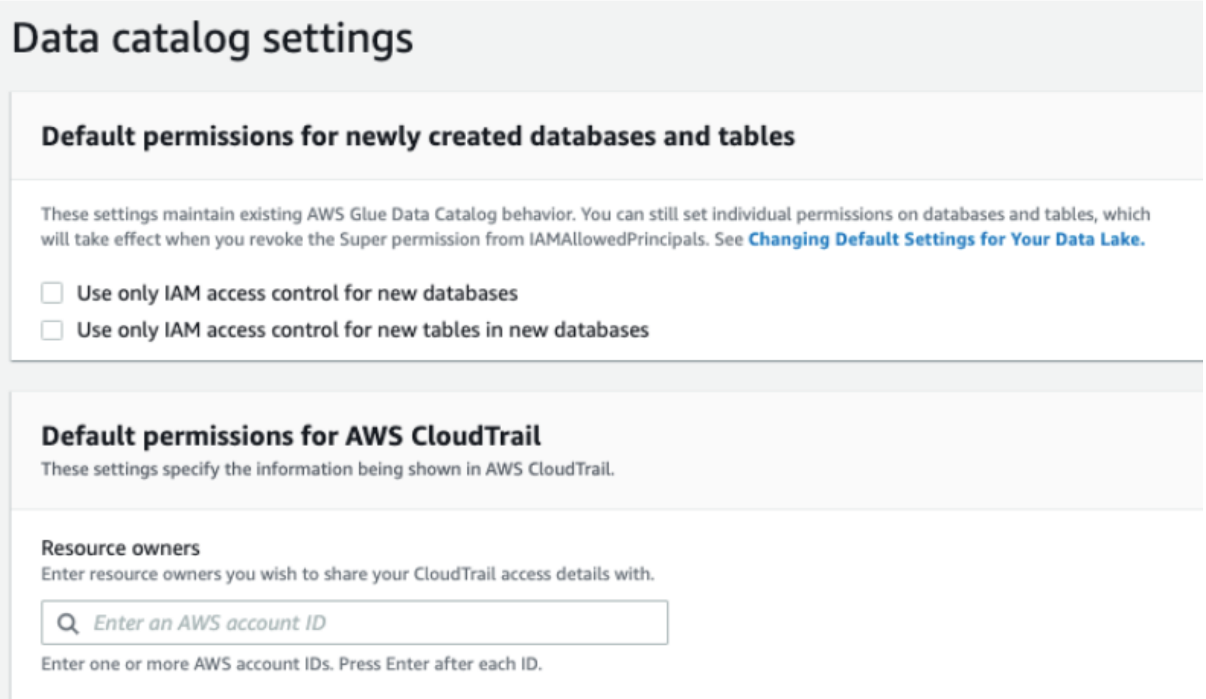

Data catalog settings

Default permissions for newly created databases and tables

These settings maintain existing AWS Glue Data Catalog behavior. You can still set individual permissions on databases and tables, which will take effect when you revoke the Super permission from IAMAllowedPrincipals. See **Changing Default Settings for Your Data Lake.**

□ Use only IAM access control for new databases

□ Use only IAM access control for new tables in new databases

Default permissions for AWS CloudTrail

These settings specify the information being shown in AWS CloudTrail.

Resource owners

Enter resource owners you wish to share your CloudTrail access details with.

Enter an AWS account ID

Enter one or more AWS account IDs. Press Enter after each ID.

Next, we need to set up the AWS Glue Data Catalog resource policy to grant cross-account access to Data Catalog resources.

7. Use the following policy, and replace the account number and Region with your own values:

```
8. {
9.    "PolicyInJson": "{\"Version\" : \"2012-10-17\",\"Statement\" : [
   {\"Effect\" : \"Allow\",\"Principal\" : {\"AWS\" :
   [\"arn:aws:iam::<ProducerAccountID>:root\",\"arn:aws:iam::<ConsumerAc
   countID>:root\"]},\"Action\" : \"glue:*\",\"Resource\" : [
   \"arn:aws:glue:<aws-region>:<CentralAccountID>:table/*\",
   \"arn:aws:glue:<aws-region>:<CentralAccountID>:database/*\",
   \"arn:aws:glue:<aws-region>:<CentralAccountID>:catalog\"
   ],\"Condition\" : {\"Bool\" : {\"glue:EvaluatedByLakeFormationTags\"
   : \"true\"}}}, {\"Effect\" : \"Allow\",\"Principal\" : {\"Service\" :
   \"ram.amazonaws.com\"},\"Action\" :
   \"glue:ShareResource\",\"Resource\" : [ \"arn:aws:glue:<aws-
   region>:<CentralAccountID>:table/*\", \"arn:aws:glue:<aws-
   region>:<CentralAccountID>:database/*\", \"arn:aws:glue:<aws-
   region>:<CentralAccountID>:catalog\" ]} ]}",
10.     "EnableHybrid": "TRUE"
   }
```

Replace the <aws-region>, <ProducerAccountID>, <ConsumerAccountID> and <CentralAccountID> values in the above policy as appropriate and save it in a file called `policy.json`.

9. Next, run the following AWS Command Line Interface (AWS CLI) command on AWS CloudShell.

```
aws glue put-resource-policy --region <aws-region> --cli-input-json
file://policy.json
```

For more information about this policy, see put-resource-policy.

10. Next, we verify the source data S3 bucket isregistered as data lake location in the central account. This is completed by the CloudFormation template.
11. Under **Register and ingest** in the navigation pane, choose **Data lake locations.**

You should see the S3 bucket registered under the data lake locations.

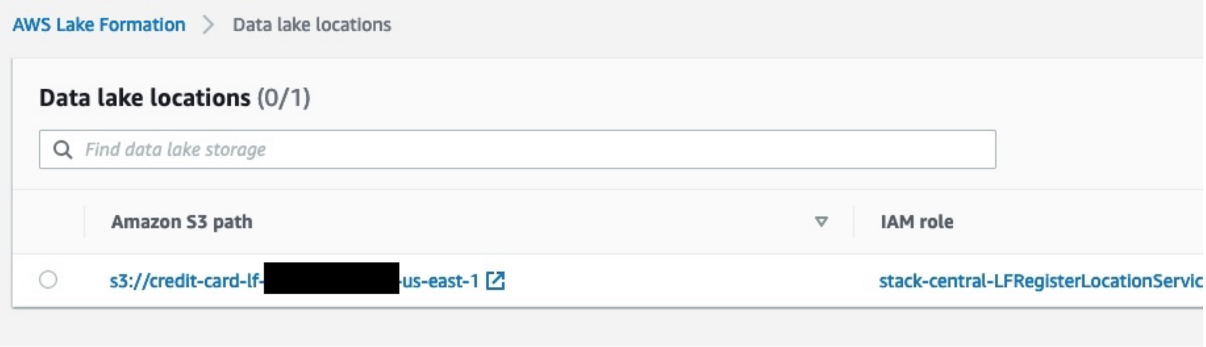

AWS Lake Formation > Data lake locations

Data lake locations (0/1)

Find data lake storage

	Amazon S3 path	IAM role
○	s3://credit-card-lf-[redacted]-us-east-1	stack-central-LFRegisterLocationServic

Configure Lake Formation Data Catalog settings in the central account

After we complete all the prerequisites, we start the data mesh configuration. We log in as `DataMeshOwner` in the central account.

Define LF-tags

`DataMeshOwner` creates the tag ontology by defining LF-tags. Complete the following steps:

1. On the Lake Formation console, under **Permissions** in the navigation pane, under **Administrative roles and tasks**, choose **LF-Tags.**
2. Choose **Add LF-tags**.
3. For **Key**, enter database and for **Values**, choose `credit-card`.
4. Choose **Add** and then **Add LF-tag**.

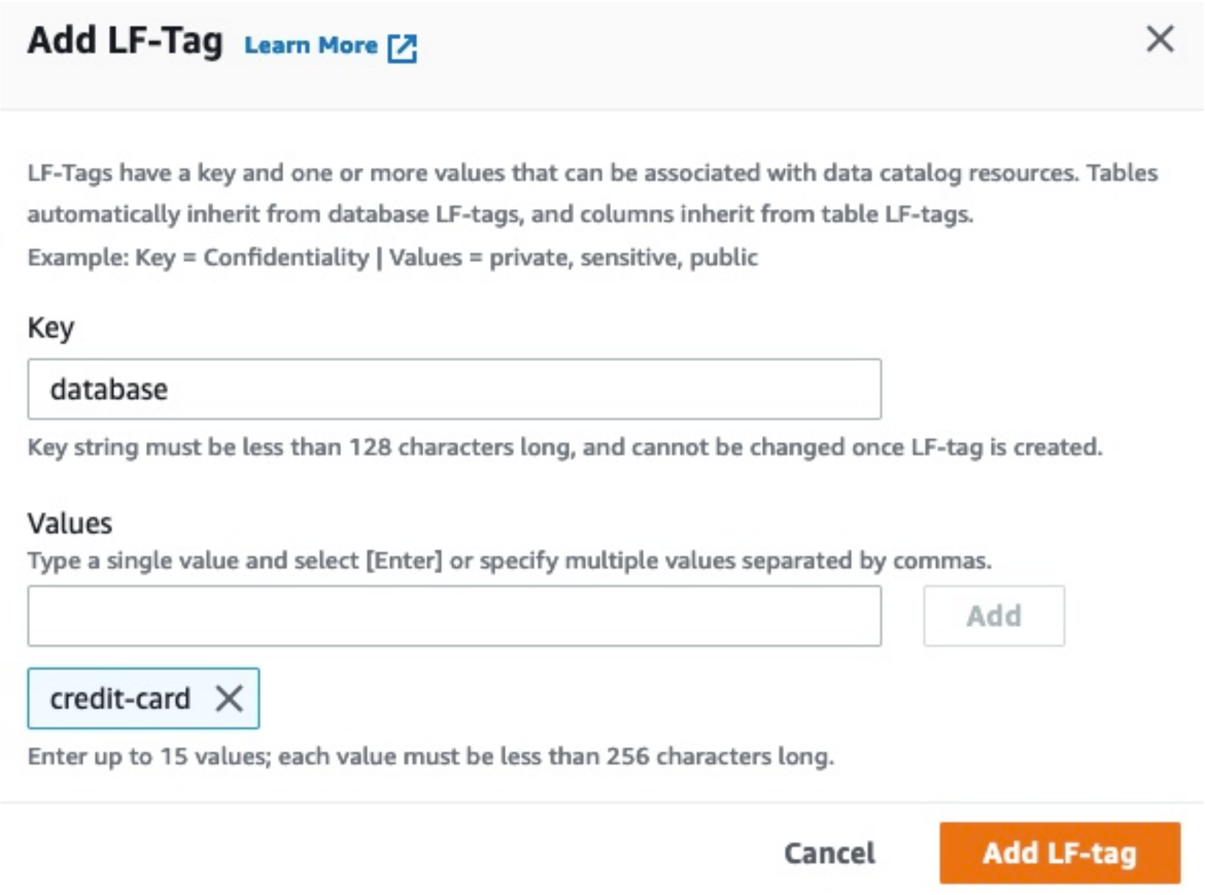

Add LF-Tag Learn More

LF-Tags have a key and one or more values that can be associated with data catalog resources. Tables automatically inherit from database LF-tags, and columns inherit from table LF-tags.
Example: Key = Confidentiality | Values = private, sensitive, public

Key

database

Key string must be less than 128 characters long, and cannot be changed once LF-tag is created.

Values
Type a single value and select [Enter] or specify multiple values separated by commas.

Add

credit-card ✕

Enter up to 15 values; each value must be less than 256 characters long.

Cancel Add LF-tag

5. The result looks like the image below.

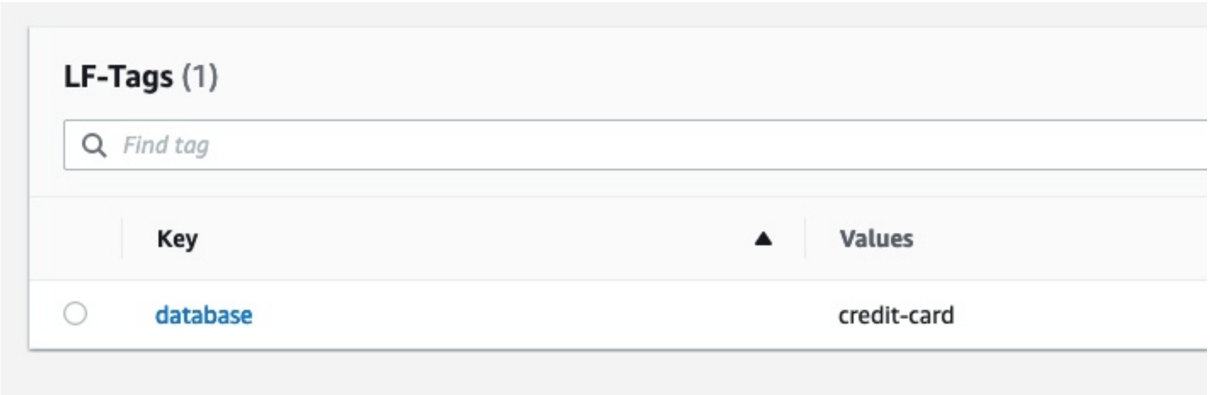

LF-Tags (1)

Find tag

	Key ▲	Values
○	database	credit-card

Grant permissions

We grant `ProducerSteward` in the central accounts describe and associate permissions on the preceding tag ontology. This enables `ProducerSteward` to view the LF-tags and assign them to Data Catalog resources (databases, tables, and columns). `ProducerSteward` in the central account can further grant the permission to `ProducerSteward` in the producer account. For more information, see Granting, Revoking, and Listing LF-Tag Permissions. When you have multiple producers, grant the relevant tags to each steward.

In our situation, we will only have one LF tag assigned which points to the database. This could be further improved by adding extra tags for more granularity on data access. But it is out of the scope for this guide.

1. Under **Permissions** in the navigation pane, under **Administrative roles and tasks**, choose **LF-tag permissions**.
2. Choose **Grant**.
3. For **IAM users and roles**, choose the `ProducerSteward` user.
4. In the **LF-Tags** section, add all three key-values:
   a. Key `database` with values `credit-card`.

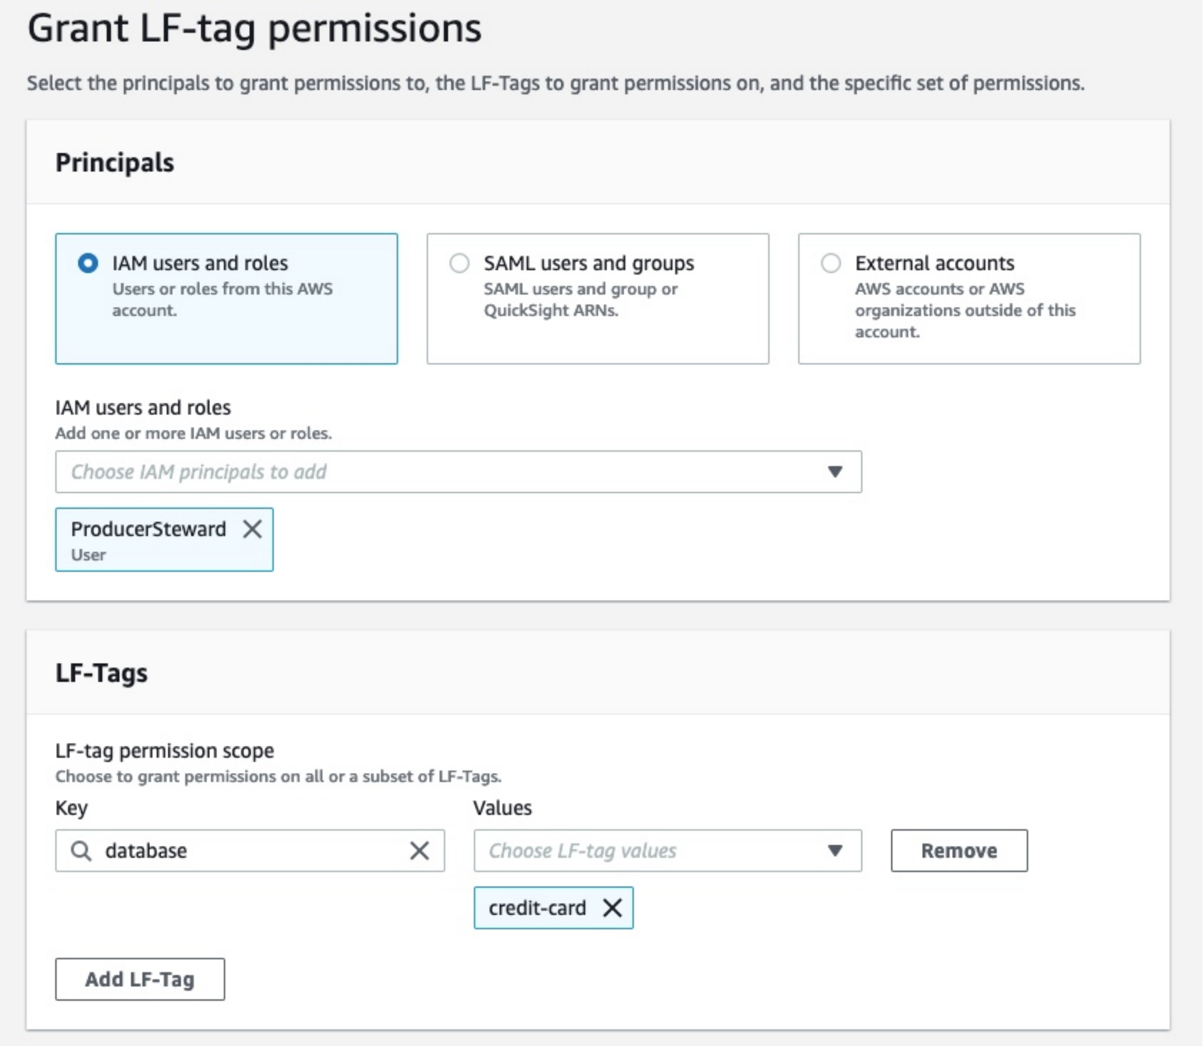

Grant LF-tag permissions

Select the principals to grant permissions to, the LF-Tags to grant permissions on, and the specific set of permissions.

Principals

● IAM users and roles
Users or roles from this AWS account.

○ SAML users and groups
SAML users and group or QuickSight ARNs.

○ External accounts
AWS accounts or AWS organizations outside of this account.

IAM users and roles
Add one or more IAM users or roles.

Choose IAM principals to add

ProducerSteward ✕
User

LF-Tags

LF-tag permission scope
Choose to grant permissions on all or a subset of LF-Tags.

Key: database ✕

Values: Choose LF-tag values

credit-card ✕

5. For **Permissions**, select **Describe** and **Associate** for both **LF-tag permissions** and **Grantable permissions**.

6. Choose **Grant.**

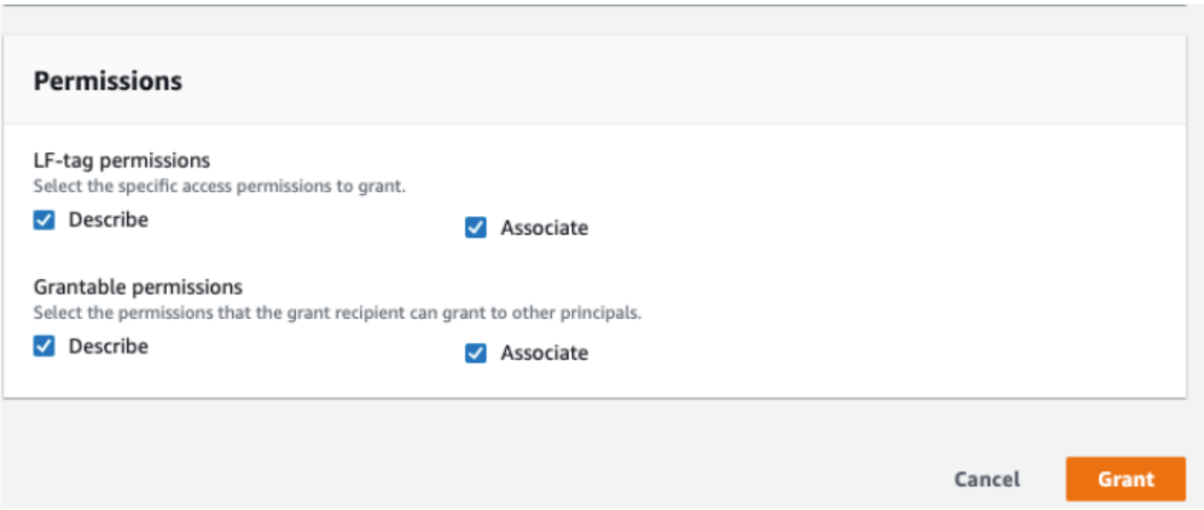

Permissions

LF-tag permissions
Select the specific access permissions to grant.
☑ Describe ☑ Associate

Grantable permissions
Select the permissions that the grant recipient can grant to other principals.
☑ Describe ☑ Associate

Next, we grant `ProducerSteward` tag-based data lake permissions. This enables `ProducerSteward` to create, alter, and drop tables in the databases with corresponding tags. `ProducerSteward` in the producer account can further grant the permission across accounts.

7. In the navigation pane, under **Permissions**, **Data lake permissions**, choose **Grant**.
8. For **Principals**, choose **IAM users and roles**, and choose `ProducerSteward`.
9. For **LF-tags or catalog resources**, select **Resources matched by LF-Tags (recommended)**.
10. Choose **Add LF-Tag**.
11. For **Key**, choose `database` and for **Values**, choose `credit-card`.
12. For **Database permissions**, select the Super permission because `ProducerSteward` owns the producer databases.

This permission allows a principal to perform every supported Lake Formation operation on the database. Use this admin permission when a principal is trusted with all operations.

13. Select **Super** under **Grantable permissions** so the `ProducerSteward` user can grant database-level permissions to the producer and consumer accounts.
14. For **Table permissions**, select **Super**.
15. Select **Super** permission under **Grantable permissions**.
16. Choose **Grant**.

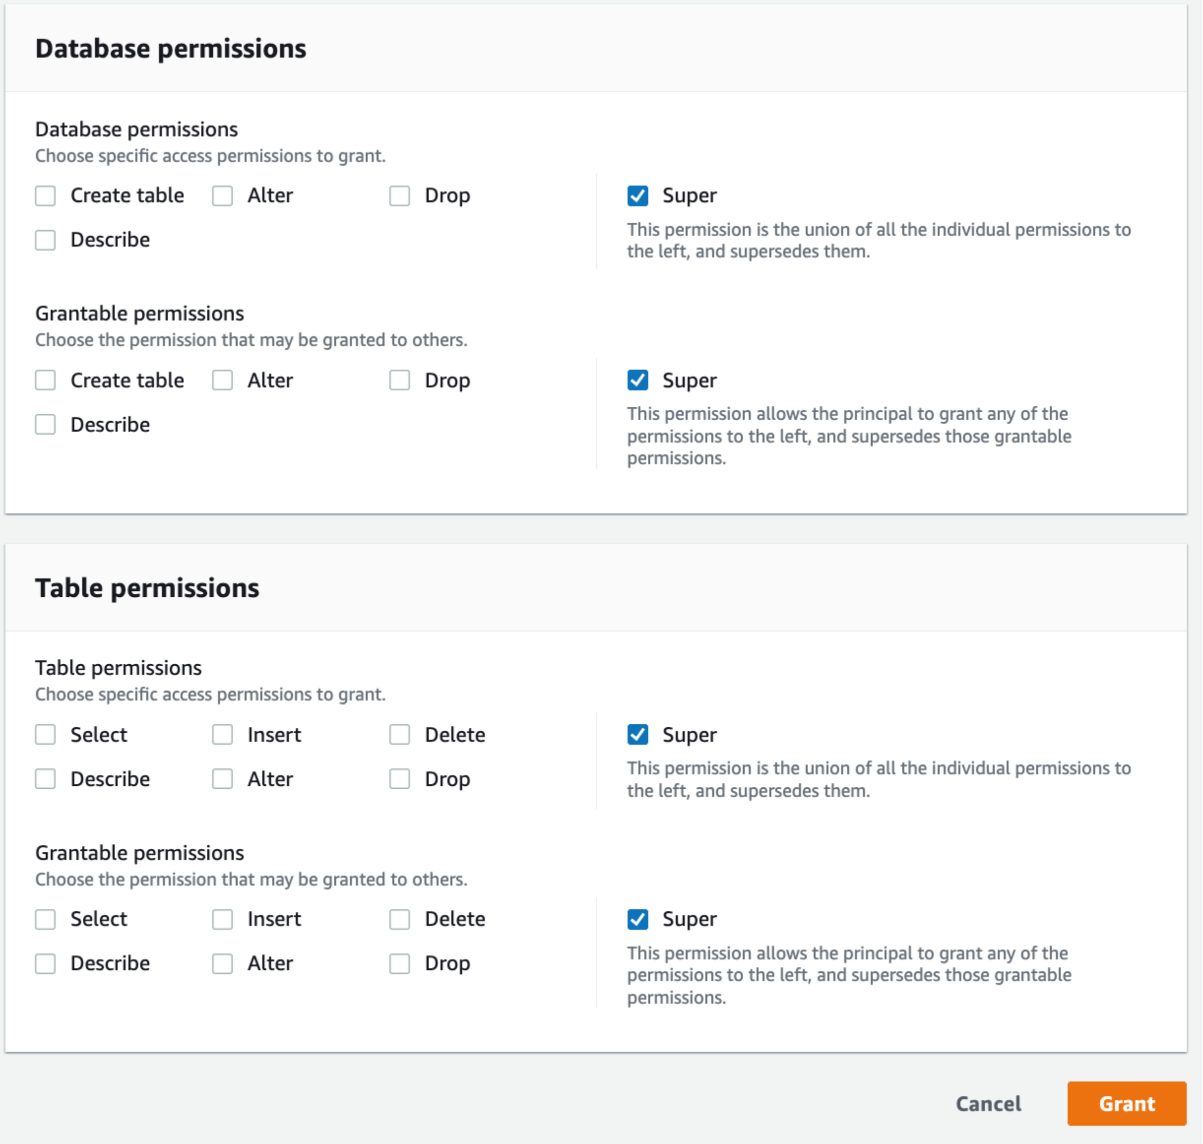

Database permissions

Database permissions
Choose specific access permissions to grant.

□ Create table □ Alter □ Drop □ Describe

☑ Super
This permission is the union of all the individual permissions to the left, and supersedes them.

Grantable permissions
Choose the permission that may be granted to others.

□ Create table □ Alter □ Drop □ Describe

☑ Super
This permission allows the principal to grant any of the permissions to the left, and supersedes those grantable permissions.

Table permissions

Table permissions
Choose specific access permissions to grant.

□ Select □ Insert □ Delete □ Describe □ Alter □ Drop

☑ Super
This permission is the union of all the individual permissions to the left, and supersedes them.

Grantable permissions
Choose the permission that may be granted to others.

□ Select □ Insert □ Delete □ Describe □ Alter □ Drop

☑ Super
This permission allows the principal to grant any of the permissions to the left, and supersedes those grantable permissions.

Producer data steward actions in the central account

Next, we log in as the `ProducerSteward` user in the central account and create skeleton databases.

1. Sign in to the Lake Formation console as `ProducerSteward`.
2. In the navigation pane, under **Data catalog,** select **Databases**.
3. Choose the `credit-card` database.
4. On the **Actions** menu, choose **Edit LF-tags**

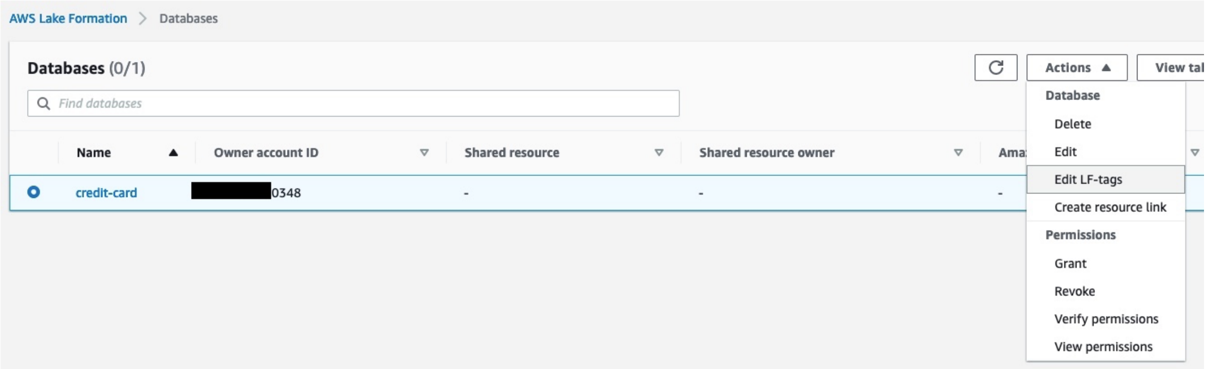

Databases (0/1)

Find databases

	Name	Owner account ID	Shared resource	Shared resource owner	Ama
●	credit-card	[redacted]0348	-	-	-

5. Choose **Assign new LF-tag**.
6. For **Assigned Keys**, enter `database` and for **Values**, choose `credit-card`.
7. Choose **Save.**

This assigns the `database=credit-card` tag to the `credit-card` database.

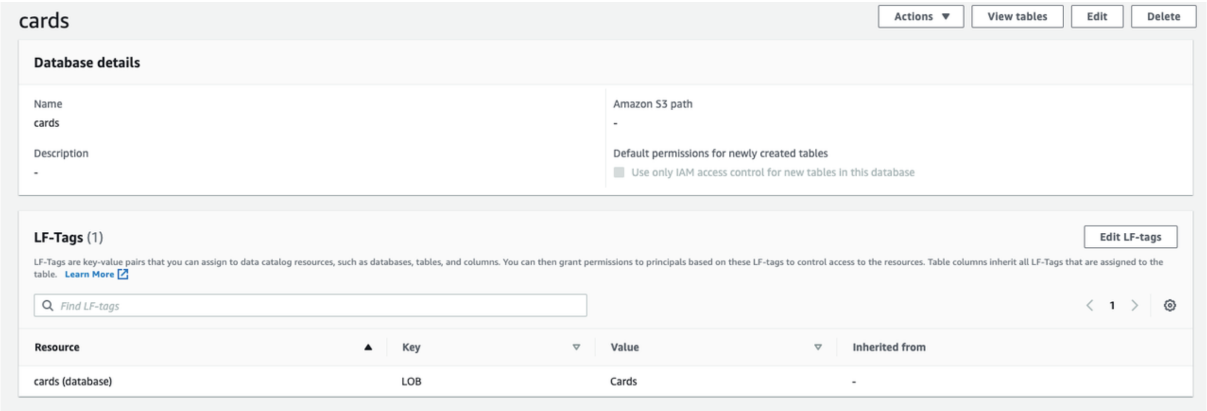

cards

Database details

Name
cards

Description
-

Amazon S3 path
-

Default permissions for newly created tables
□ Use only IAM access control for new tables in this database

LF-Tags (1)

LF-Tags are key-value pairs that you can assign to data catalog resources, such as databases, tables, and columns. You can then grant permissions to principals based on these LF-tags to control access to the resources. Table columns inherit all LF-Tags that are assigned to the table. Learn More

Resource	Key	Value	Inherited from
cards (database)	LOB	Cards	-

Next, we share the LF-tags and data lake permissions with the producer account so that `ProducerSteward` in the producer account can run AWS Glue crawlers and generate tables in the preceding skeleton databases.

12. Under **Permissions** in the navigation pane, under **Administrative roles and tasks**, choose **LF-tag permissions**.
13. Choose **Grant**.
14. For **Principals**, select **External accounts**.
15. For **AWS account or AWS organization**, enter the account ID for the producer account.
16. In the **LF-Tags** section, we only need to add database-level tags.
17. For **Key**, enter `database` and for **Values**, choose `credit-card`.
18. For **Permissions**, choose **Describe** and **Associate** for both **LF-tag permissions** and **Grantable permissions**.

19. Choose **Grant**.

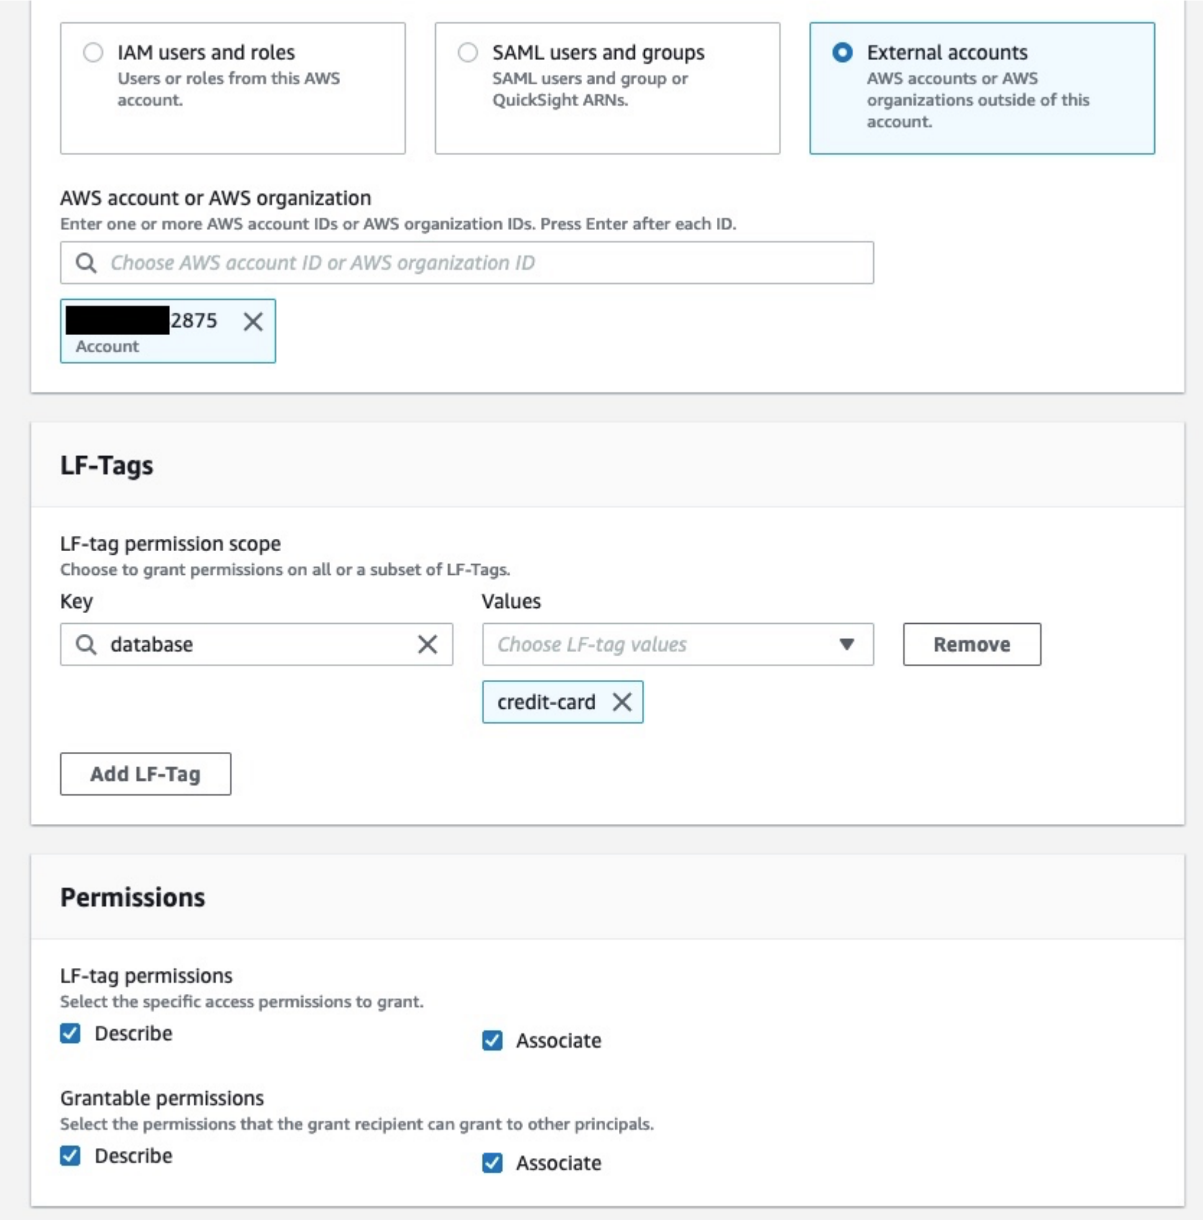

○ IAM users and roles
Users or roles from this AWS account.

○ SAML users and groups
SAML users and group or QuickSight ARNs.

● External accounts
AWS accounts or AWS organizations outside of this account.

AWS account or AWS organization
Enter one or more AWS account IDs or AWS organization IDs. Press Enter after each ID.

Choose AWS account ID or AWS organization ID

[redacted]2875 ✕
Account

LF-Tags

LF-tag permission scope
Choose to grant permissions on all or a subset of LF-Tags.

Key	Values
database ✕	Choose LF-tag values credit-card ✕

Remove

Add LF-Tag

Permissions

LF-tag permissions
Select the specific access permissions to grant.

☑ Describe ☑ Associate

Grantable permissions
Select the permissions that the grant recipient can grant to other principals.

☑ Describe ☑ Associate

20. In the navigation pane, under **Permissions**, **Data lake permissions**, choose **Grant**.
21. For **Principals**, select **External accounts**.
22. For **AWS account or AWS organization**, enter the account ID for the producer account.
23. For **LF-tags or catalog resources**, select **Resources matched by LF-Tags (recommended)**.
24. Choose **Add LF-Tag**.

25. Choose the key `database` and value `credit-card`.

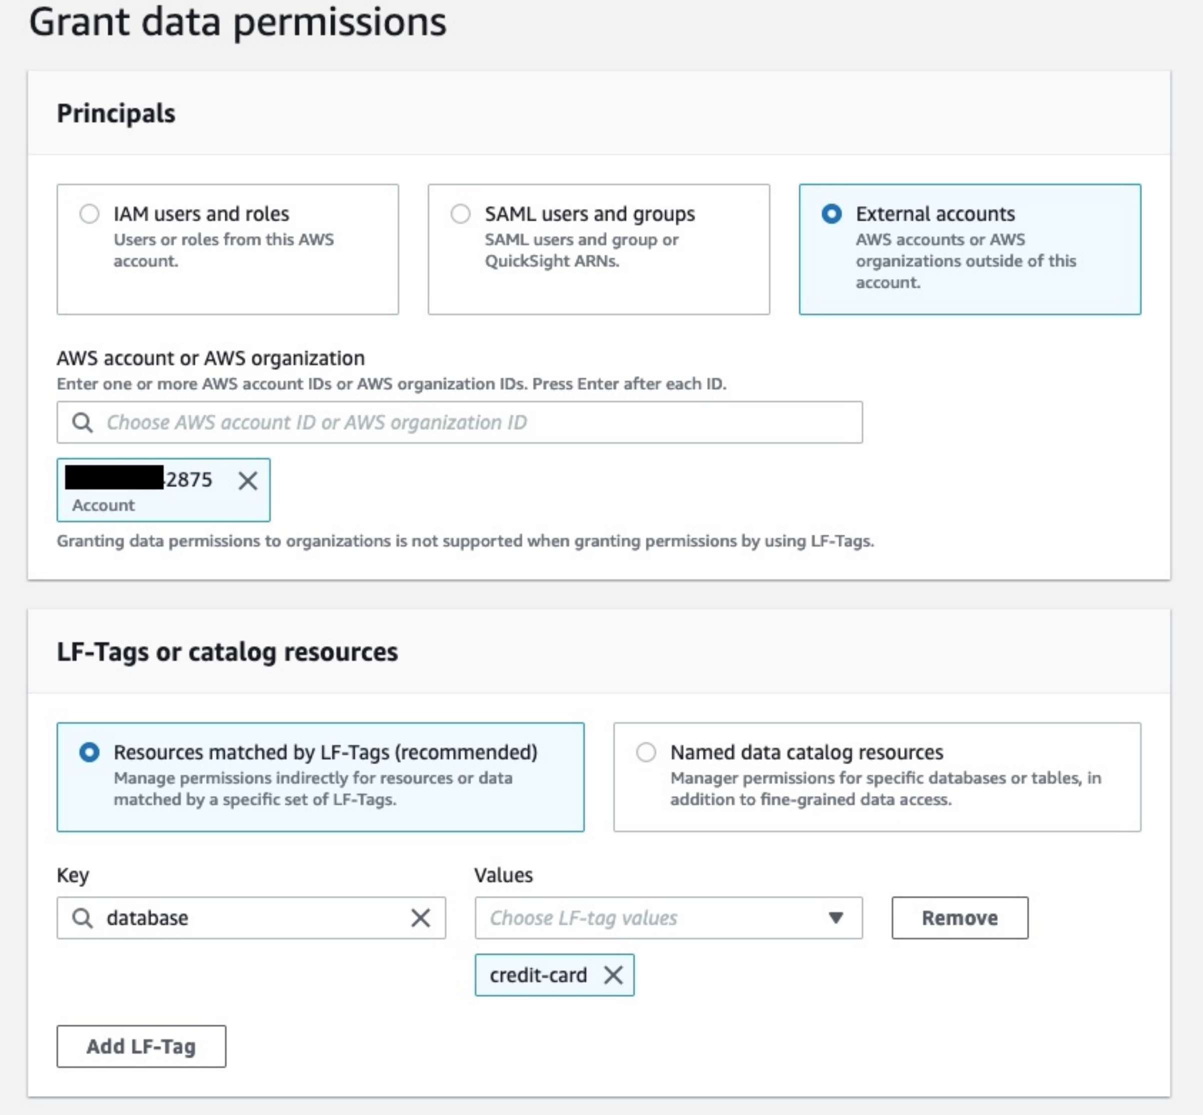

Grant data permissions

Principals

○ IAM users and roles
Users or roles from this AWS account.

○ SAML users and groups
SAML users and group or QuickSight ARNs.

● External accounts
AWS accounts or AWS organizations outside of this account.

AWS account or AWS organization
Enter one or more AWS account IDs or AWS organization IDs. Press Enter after each ID.

Choose AWS account ID or AWS organization ID

[redacted]2875 ✕
Account

Granting data permissions to organizations is not supported when granting permissions by using LF-Tags.

LF-Tags or catalog resources

● Resources matched by LF-Tags (recommended)
Manage permissions indirectly for resources or data matched by a specific set of LF-Tags.

○ Named data catalog resources
Manager permissions for specific databases or tables, in addition to fine-grained data access.

Key: database ✕

Values: Choose LF-tag values ▼
credit-card ✕

26. For **Database permissions**, select **Create table** and **Describe** because the `ProducerSteward` user in the producer account will add tables in the database.
27. Select **Create table** and **Describe** under **Grantable permissions** so the `ProducerSteward` user can further grant the permission to the AWS Glue crawler.
28. For **Table permissions**, select all the permissions.
29. Select all the permissions under **Grantable permissions.**

30. Choose **Grant**.

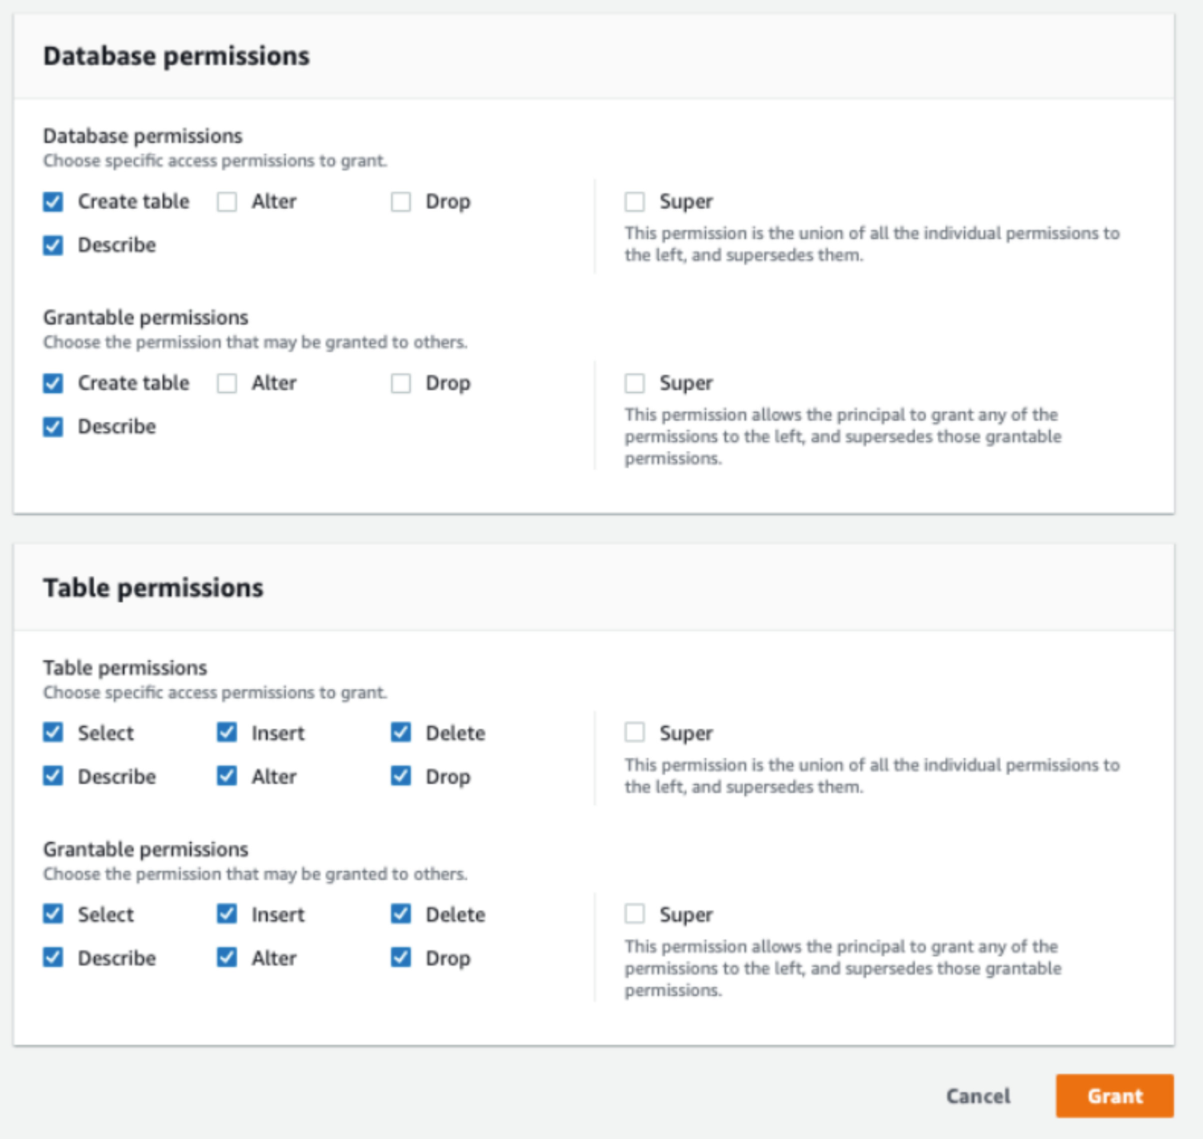

Database permissions

Database permissions
Choose specific access permissions to grant.

☑ Create table ☐ Alter ☐ Drop
☑ Describe

☐ Super
This permission is the union of all the individual permissions to the left, and supersedes them.

Grantable permissions
Choose the permission that may be granted to others.

☑ Create table ☐ Alter ☐ Drop
☑ Describe

☐ Super
This permission allows the principal to grant any of the permissions to the left, and supersedes those grantable permissions.

Table permissions

Table permissions
Choose specific access permissions to grant.

☑ Select ☑ Insert ☑ Delete
☑ Describe ☑ Alter ☑ Drop

☐ Super
This permission is the union of all the individual permissions to the left, and supersedes them.

Grantable permissions
Choose the permission that may be granted to others.

☑ Select ☑ Insert ☑ Delete
☑ Describe ☑ Alter ☑ Drop

☐ Super
This permission allows the principal to grant any of the permissions to the left, and supersedes those grantable permissions.

Now the Lake Formation administrators on the producer account side has the right permissions to add tables.

Crawl source tables in the producer account

Next, we log in as the `ProducerSteward` user in the producer account to crawl the source tables for the `Cards` and `Retail` databases.

1. Sign in to the Lake Formation console as `ProducerSteward`.
2. In the navigation pane, under **Administrative Roles and Tasks**, verify that `ProducerSteward` is configured as the data lake administrator.

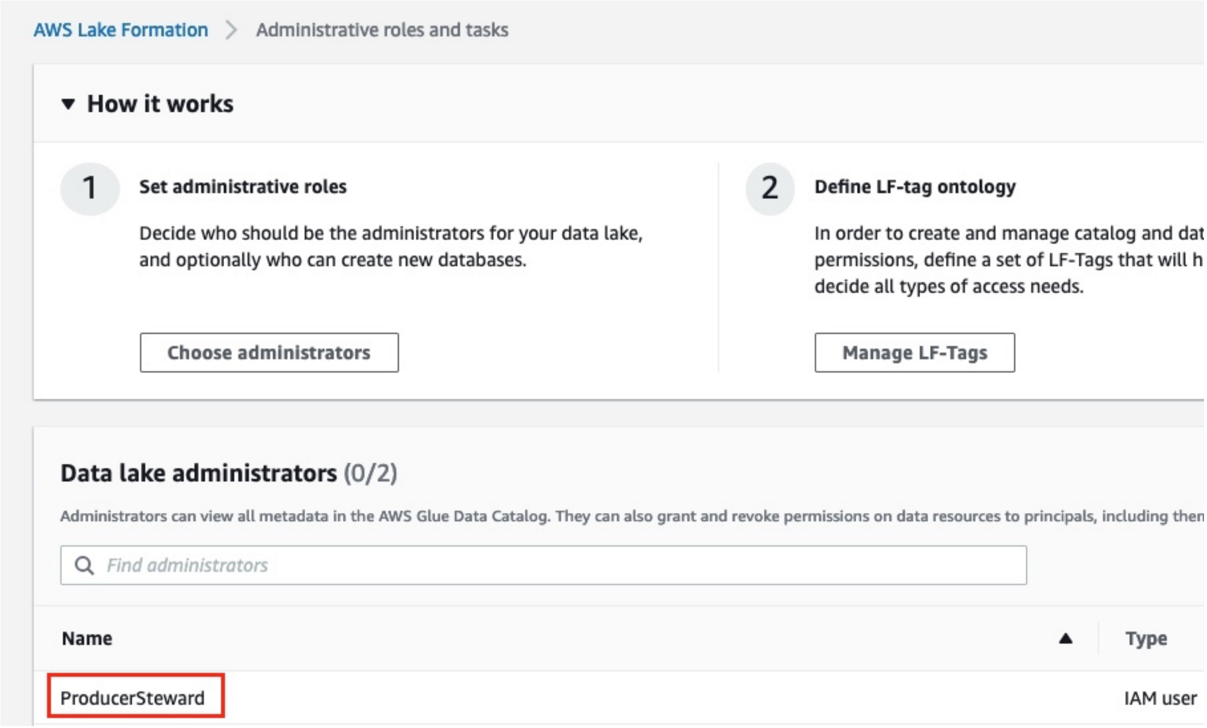

Set administrative roles

Decide who should be the administrators for your data lake, and optionally who can create new databases.

Define LF-tag ontology

In order to create and manage catalog and dat permissions, define a set of LF-Tags that will h decide all types of access needs.

Data lake administrators (0/2)

Administrators can view all metadata in the AWS Glue Data Catalog. They can also grant and revoke permissions on data resources to principals, including then

Name ▲	Type
ProducerSteward	IAM user

3. In the navigation pane, under **Permissions**, then choose **Administrative roles and tasks**, choose **LF-Tags**.

You can verify the `database` tag that was shared with the producer account.

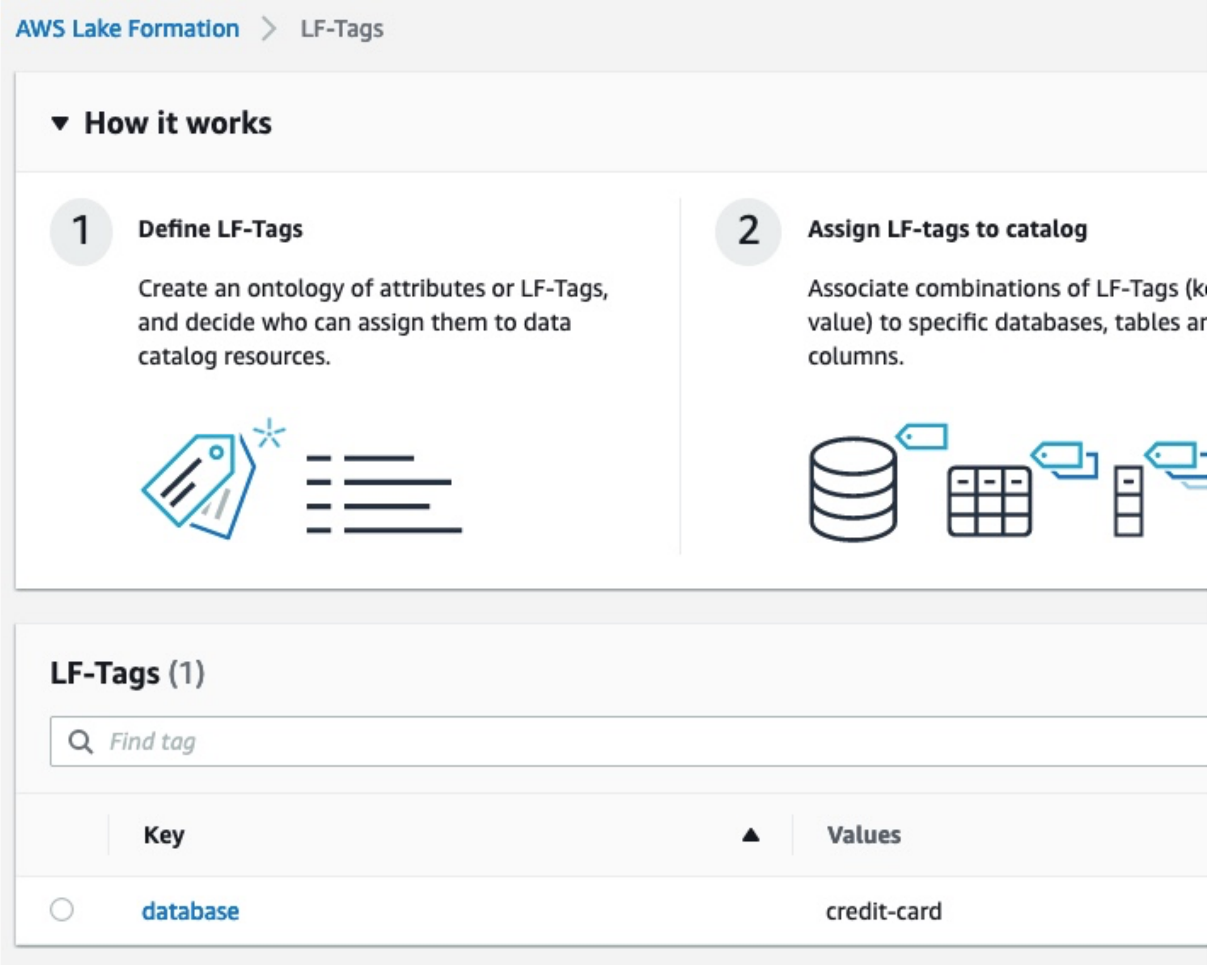

AWS Lake Formation > LF-Tags

▼ How it works

1 **Define LF-Tags**

Create an ontology of attributes or LF-Tags, and decide who can assign them to data catalog resources.

2 **Assign LF-tags to catalog**

Associate combinations of LF-Tags (k value) to specific databases, tables ar columns.

LF-Tags (1)

Find tag

	Key ▲	Values
○	database	credit-card

4. In the navigation pane, under **Data catalog**, select **Databases.**

You can verify the two databases `cards` and `retail` that were shared with the producer account from the previous step.

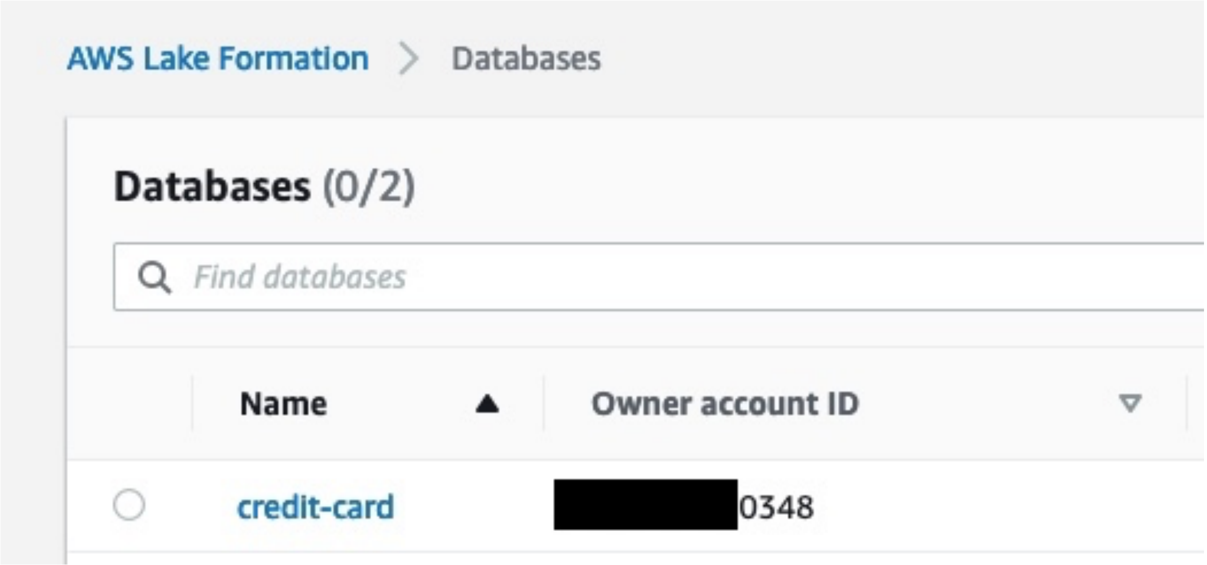

AWS Lake Formation > Databases

Databases (0/2)

Find databases

	Name ▲	Owner account ID ▽
○	credit-card	[redacted]0348

Now, we create a resource link in the producer account for this database. This link point at the shared database and is used by AWS Glue crawler to create the tables. First, we create a resource link for the `credit-card` database.

5. Select the `cards` database and on the **Actions** menu, choose **Create resource link**.

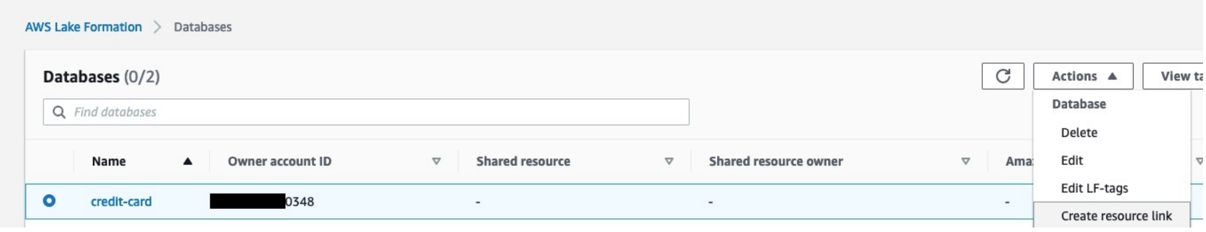

Databases (0/2)

Find databases

	Name	Owner account ID	Shared resource	Shared resource owner	Ama
●	credit-card	[redacted]0348	-	-	-

6. For **Resource link name**, enter `rl_credit-card`.
7. Choose **Create**.

After the resource link creation, you should see both the resource link databases as shown in the following screenshot.

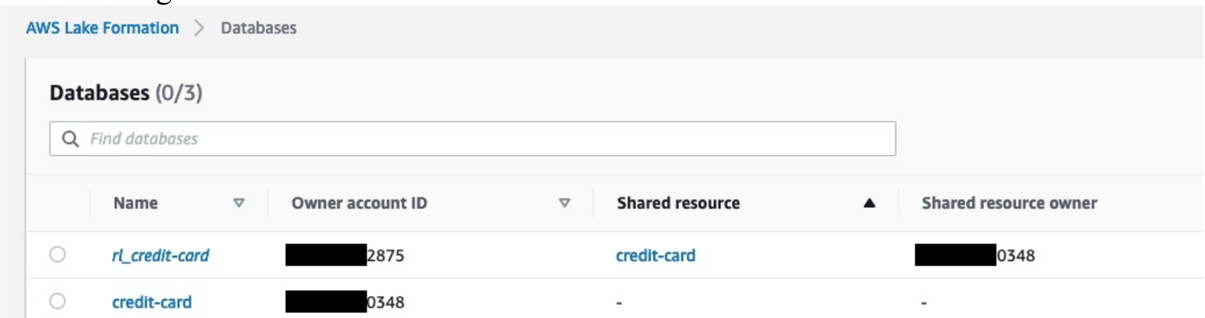

Databases (0/3)

Find databases

	Name	Owner account ID	Shared resource	Shared resource owner
○	*rl_credit-card*	[redacted]2875	credit-card	[redacted]0348
○	credit-card	[redacted]0348	-	-

Next, we need to grant permissions to the AWS Glue crawler role so that the crawler can crawl the source bucket and create the tables.

9. Select the `rl_credit-card` database and on the **Actions** menu, choose **Grant**.
10. In the **Grant data permissions** section, select **IAM users and roles**, and choose the AWS Glue crawler role that was created by the CloudFormation template (for example, `stack-producer-AWSGlueServiceRoleDefault-xxxxxx`).
11. For **Databases**, choose `rl_credit-card`.
12. For **Resource link permissions**, select **Describe**.

13. Choose **Grant**.

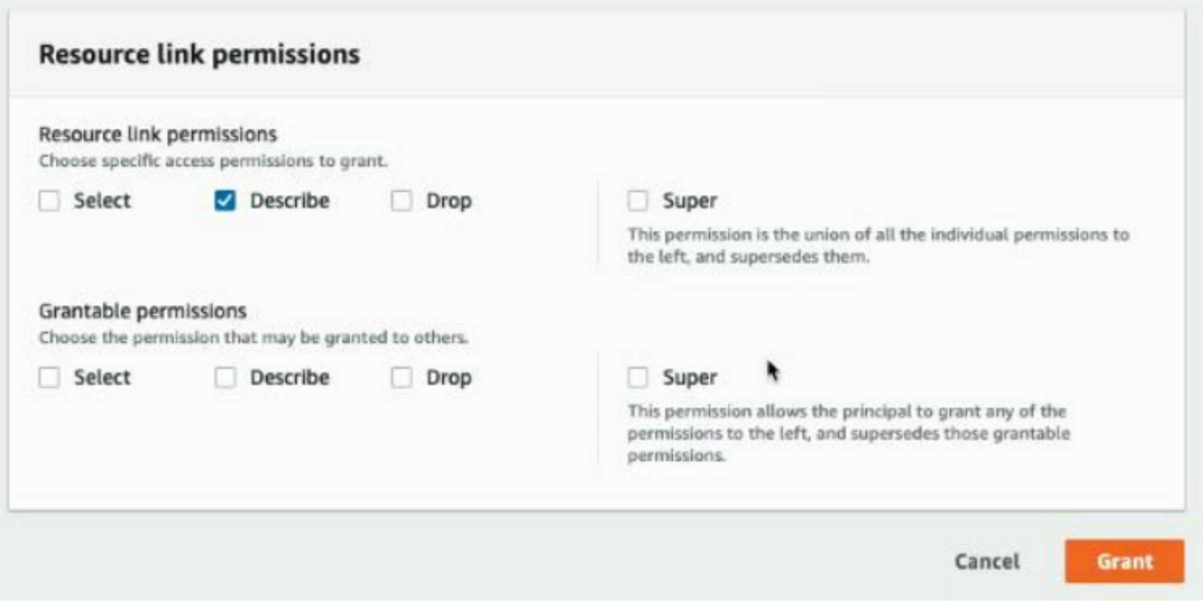

Resource link permissions

Resource link permissions
Choose specific access permissions to grant.

☐ Select ☑ Describe ☐ Drop

☐ Super
This permission is the union of all the individual permissions to the left, and supersedes them.

Grantable permissions
Choose the permission that may be granted to others.

☐ Select ☐ Describe ☐ Drop

☐ Super
This permission allows the principal to grant any of the permissions to the left, and supersedes those grantable permissions.

14. Next, in the navigation pane, choose **Data lake Permissions** and choose **Grant**.
15. For **IAM users and roles**, choose the role `stack-producer-AWSGlueServiceRoleDefault-XXXX`.
16. For **LF-Tags or catalog resources**, select **Resources matched by LF-Tags**.
17. Enter the key `database` and values `credit-card`.
18. For **Database permissions**, select **Create table** and **Describe**.
19. For **Table permissions**, choose **Select**, **Describe**, and **Alter**.
20. Choose **Grant**.

Next, we will verify grant permissions on the S3 bucket locations corresponding to credit-card producer to the AWS Glue crawler role. This is completed by the CloudFormation template.

In the navigation pane, under **Permissions**, on the **Data Locations**, you should see the locations.

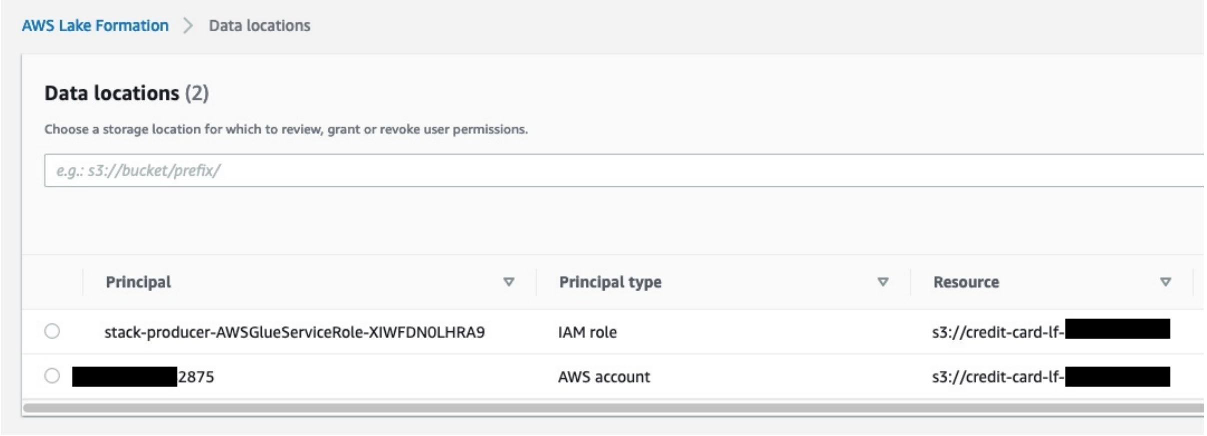

AWS Lake Formation > Data locations

Data locations (2)

Choose a storage location for which to review, grant or revoke user permissions.

e.g.: s3://bucket/prefix/

	Principal	Principal type	Resource
○	stack-producer-AWSGlueServiceRole-XIWFDN0LHRA9	IAM role	s3://credit-card-lf-[redacted]
○	[redacted]2875	AWS account	s3://credit-card-lf-[redacted]

Now we're ready to run the crawler. We configure the crawler that the CloudFormation template created, to point it to the resource link database.

22. On the AWS Glue console, under **Data catalog** in the navigation pane, choose **Crawlers**.

The crawler you created should be listed.

23. Select the crawler for the cards database `CardsCrawler-xxxxxxxxxxxx` and on the **Action** menu, choose **Edit crawler**.

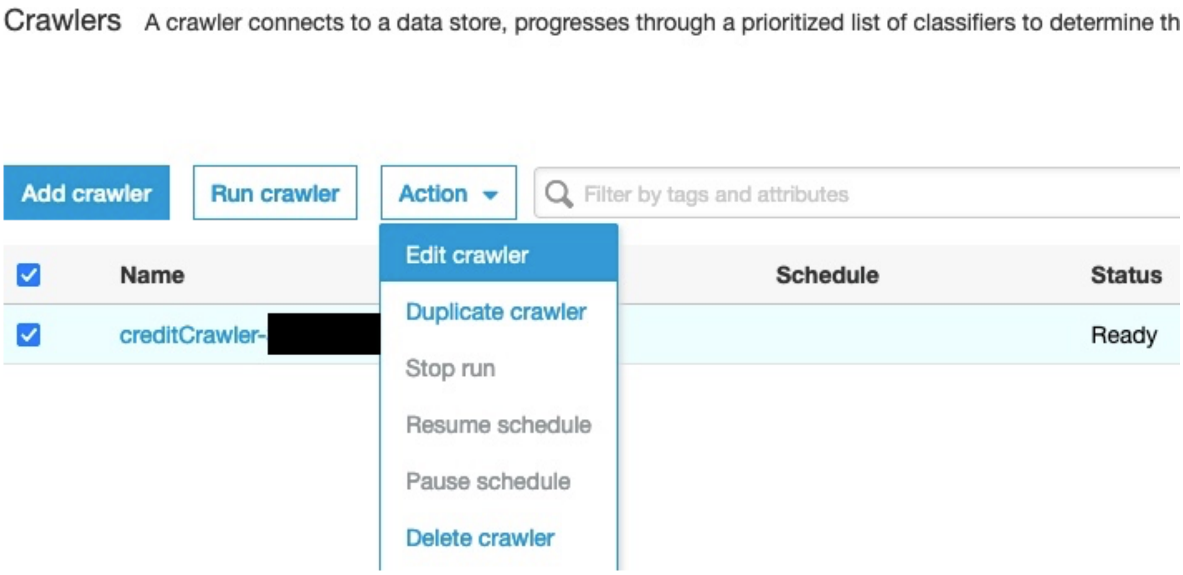

Crawlers A crawler connects to a data store, progresses through a prioritized list of classifiers to determine th

☑	Name	Schedule	Status
☑	creditCrawler-[redacted]		Ready

24. For the input data store, choose the S3 bucket for the `credit-card` producer.
25. For **IAM role**, choose the AWS Glue service role created by the CloudFormation template.
26. For **Schedule**, choose **Run on demand**.
27. For the output database, choose the resource link database `rl_credit-card` corresponding to the `cards` database.

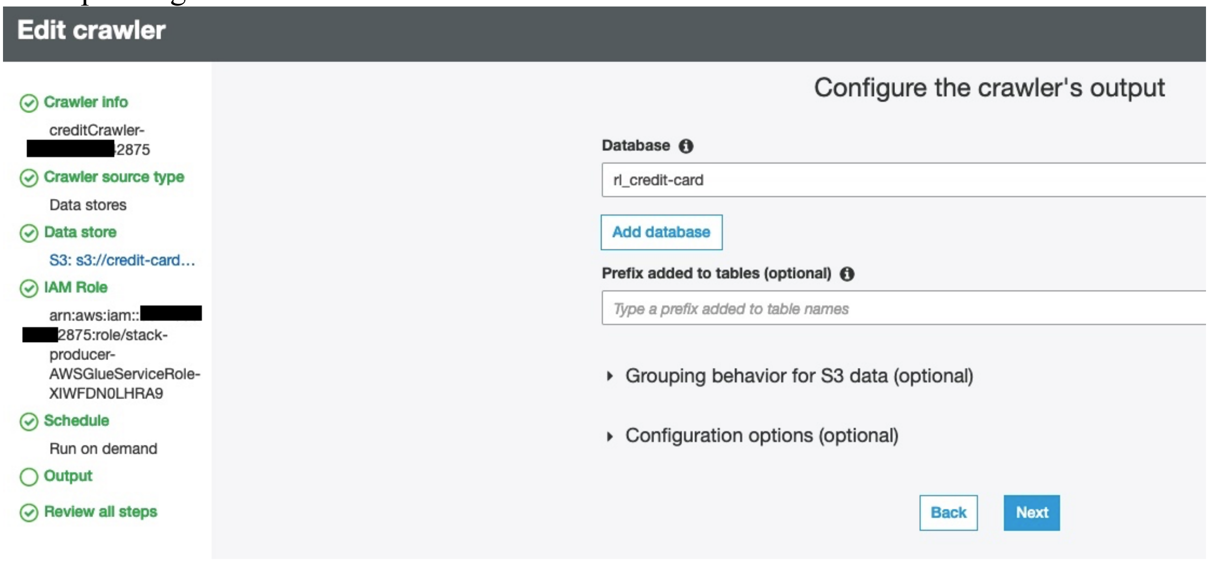

Edit crawler

Crawler info
creditCrawler-[redacted]2875
Crawler source type
Data stores
Data store
S3: s3://credit-card...
IAM Role
arn:aws:iam::[redacted]2875:role/stack-producer-AWSGlueServiceRole-XIWFDN0LHRA9
Schedule
Run on demand
Output
Review all steps

Configure the crawler's output

Database: rl_credit-card

Prefix added to tables (optional): Type a prefix added to table names

Add database
Grouping behavior for S3 data (optional)
Configuration options (optional)
Back
Next

28. Verify all the information and choose **Save**.
29. Select the crawler and choose **Run crawler**.

When the crawler finish, it creates table(s) corresponding to the producer in its respective resource link database. The table schemas are present in the shared database in the central account.

Configure Lake Formation tags in the central account

1. Log in to central account as IAM user `ProducerSteward`.
2. On the Lake Formation console, in the navigation pane, choose **Data catalog** and then choose **Tables.**

You should see the `credit_card` table corresponding to `credit-card` database.

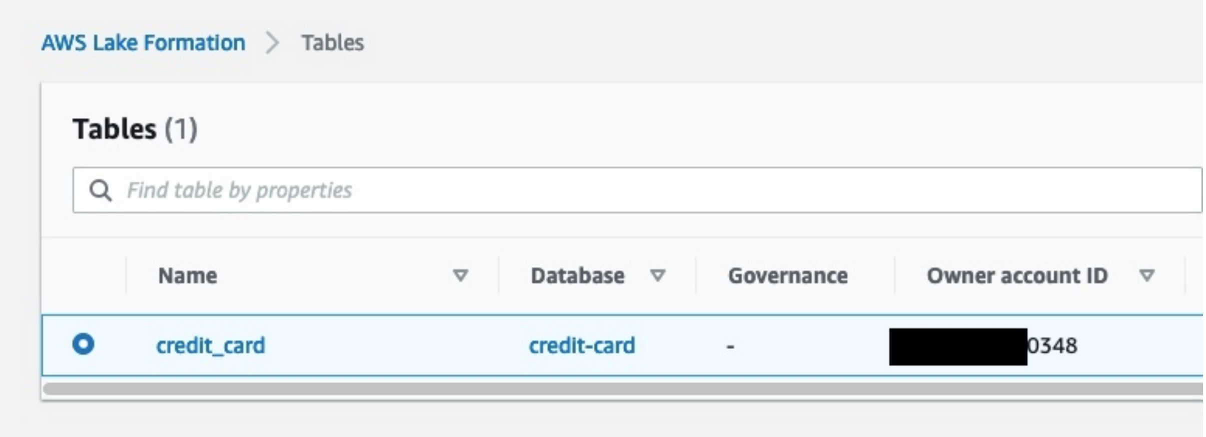

AWS Lake Formation > Tables

Tables (1)

Find table by properties

	Name	Database	Governance	Owner account ID
●	credit_card	credit-card	-	[redacted]0348

Grant tag permissions

Next, grant LF-tag permissions to the external consumer account.

1. On the Lake Formation console, in the navigation pane, choose **Permissions**, then choose **Administrative roles and tasks** and choose **LF-tag permissions**.
2. Choose **Grant**.
3. For **Principals**, select **External accounts**.
4. For **AWS account or AWS organization**, enter the AWS account number corresponding to the consumer account.
5. For **LF-Tags**, choose **Add LF-Tag**.

6. For **Key**, choose `database` and for **Values**, choose `credit-card`.

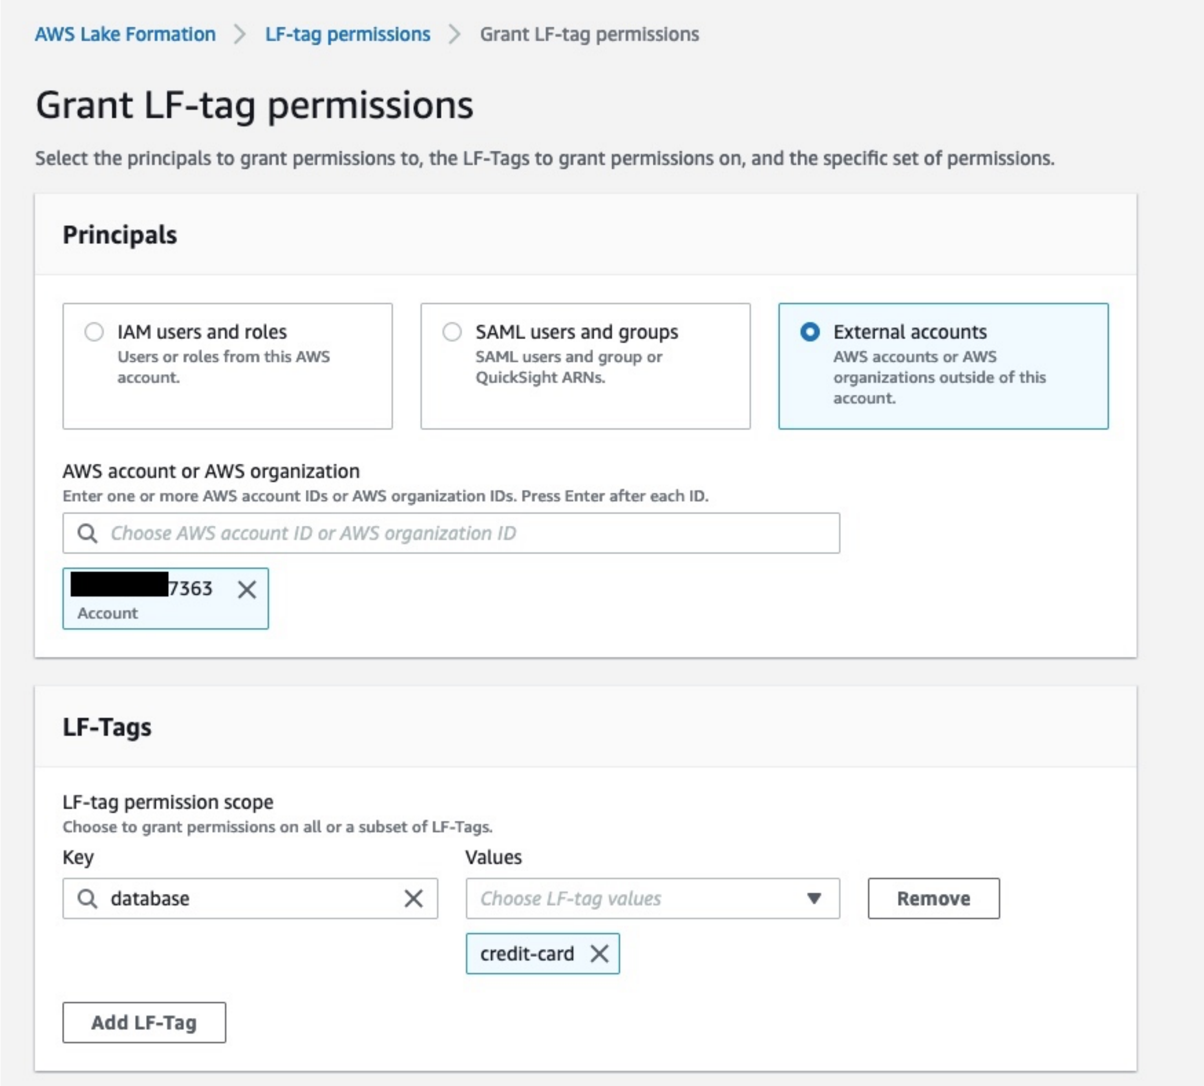

AWS Lake Formation > LF-tag permissions > Grant LF-tag permissions

Grant LF-tag permissions

Select the principals to grant permissions to, the LF-Tags to grant permissions on, and the specific set of permissions.

Principals

○ IAM users and roles
Users or roles from this AWS account.

○ SAML users and groups
SAML users and group or QuickSight ARNs.

● External accounts
AWS accounts or AWS organizations outside of this account.

AWS account or AWS organization
Enter one or more AWS account IDs or AWS organization IDs. Press Enter after each ID.

Choose AWS account ID or AWS organization ID

[redacted]7363 ✕
Account

LF-Tags

LF-tag permission scope
Choose to grant permissions on all or a subset of LF-Tags.

Key: database ✕

Values: Choose LF-tag values ▼

credit-card ✕

7. For **Permissions**, choose **Describe**.
8. For **Grantable permissions**, choose **Describe**.
9. Choose **Grant**.

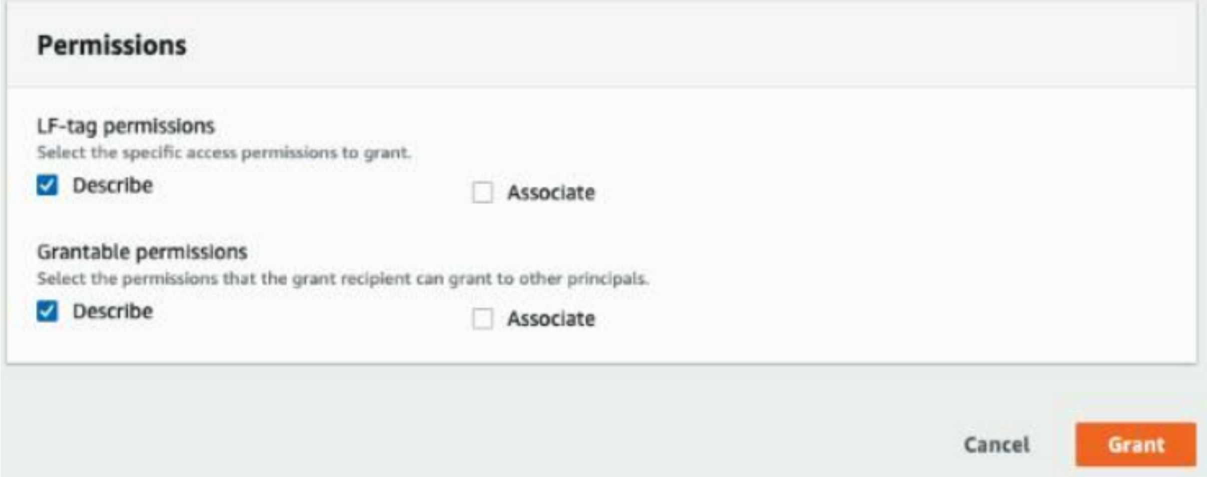

Permissions

LF-tag permissions
Select the specific access permissions to grant.

☑ Describe □ Associate

Grantable permissions
Select the permissions that the grant recipient can grant to other principals.

☑ Describe □ Associate

Next, we grant Lake Formation policy tag expression permissions to the external consumer account.

11. In the navigation pane, choose **Data lake permissions** and choose **Grant**.
12. In the **Principals** section, select **External accounts**.
13. For **AWS account or AWS organization**, enter the AWS account number corresponding to the consumer account.
14. For **LF-Tags or catalog resources**, select **Resources matched by LF-Tags**.
15. Choose **Add LF-Tag**.
16. For **Key**, choose database and for **Values**, choose `credit-card`.
17. For **Database permissions**, select **Describe**.
18. For **Grantable permissions**, select **Describe**.
19. Choose **Grant**.

Next, we grant table permissions.

21. In the navigation pane, choose **Data lake permissions** and choose **Grant**.
22. For **Principals**, select **External accounts**.
23. For **AWS account or AWS organization**, enter the AWS account number corresponding to the consumer account.
24. For **LF-Tags or catalog resources**, select **Resources matched by LF-Tags**.
25. Add key `database` with value `credit-card`
26. For **Table Permissions**, select **Select** and **Describe**.
27. For **Grantable permissions**, select **Select** and **Describe**.
28. Choose **Grant**.

Share and consume tables in the consumer account

When you sign in to the Lake Formation console in the consumer account as `ConsumerAdmin`, you can see the tags and the corresponding value that was shared by the producer.

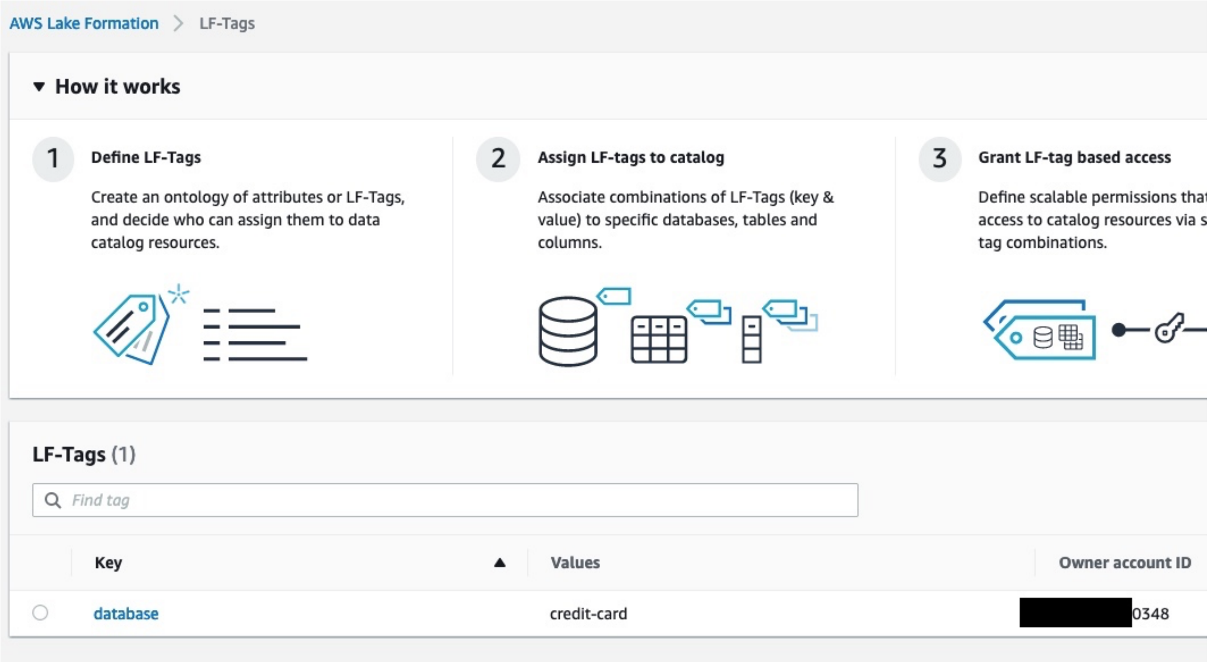

AWS Lake Formation > LF-Tags

▼ How it works

1 **Define LF-Tags**

Create an ontology of attributes or LF-Tags, and decide who can assign them to data catalog resources.

2 **Assign LF-tags to catalog**

Associate combinations of LF-Tags (key & value) to specific databases, tables and columns.

3 **Grant LF-tag based access**

Define scalable permissions that access to catalog resources via s tag combinations.

LF-Tags (1)

Find tag

	Key ▲	Values	Owner account ID
○	database	credit-card	[redacted]0348

In these next steps, we share and consume the table in the consumer account.

Create a resource link to the shared database

On the **Databases** page on the Lake Formation console, you can see all the databases that were shared to the consumer account. To create a resource link, complete the following steps:

1. On the **Databases** page, select the `credit-card` database and on the **Actions** menu, choose **Create resource link**.

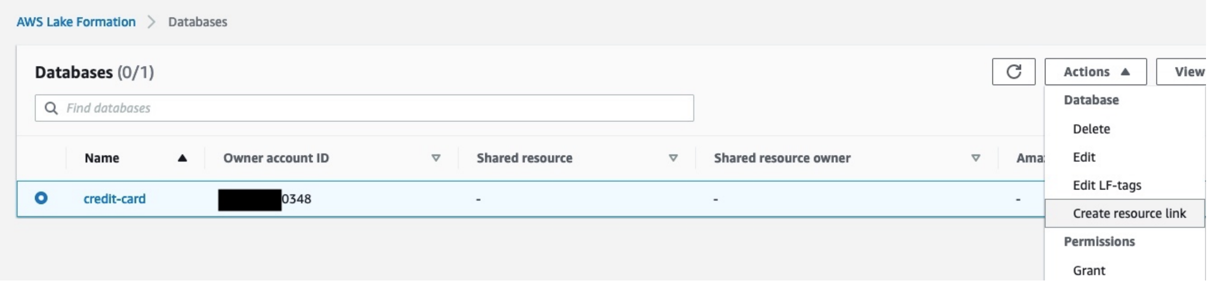

	Name ▲	Owner account ID ▽	Shared resource ▽	Shared resource owner ▽	Ama
●	credit-card	[redacted]0348	-	-	-

2. Enter the resource link name as `rl_credit-card`.
3. Leave the shared database and shared database's owner ID as default.
4. Choose **Create**.

Grant Describe permission to SageMaker (SM) role used by the SageMaker Studio user

To grant Describe permissions on resource link databases to SM Studio user, complete the following steps:

1. On the **Databases** page, select the resource database `rl_credit-card` and on the **Actions** menu, choose **Grant**.
2. In the **Grant data permissions** section, select **IAM users and roles**.
3. Choose the role corresponding to the SageMaker Studio user.
4. In the **Resource link permissions** section, select **Describe**.
5. Choose **Grant**.

   Grant Tag permissions to ConsumerAnalyst1

To grant Tag permissions on the `database:credit-card` tag to SM Studio user to access the credit_card table, complete the following steps:

1. On the Lake Formation console, on the **Data permission** page, choose **Grant**.
2. In the **Grant data permissions** section, select **IAM users and roles**.
3. Choose the role corresponding to the SageMaker Studio user
4. For **LF-Tags or catalog resources**, select **Resources matched by LF-Tags**.
5. Add the key database with value `credit-card`
6. For **Table permissions**, select **Select** and **Describe**.
7. Choose **Grant**.

The last step for the SageMaker Studio user to be able to access the data is to update the Studio user role by adding the following policy to the SageMaker role used by the Studio user

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "lakeformation:GetDataAccess",
        "glue:GetDatabase"
      ],
      "Resource": "*"
    }
  ]
}
```