
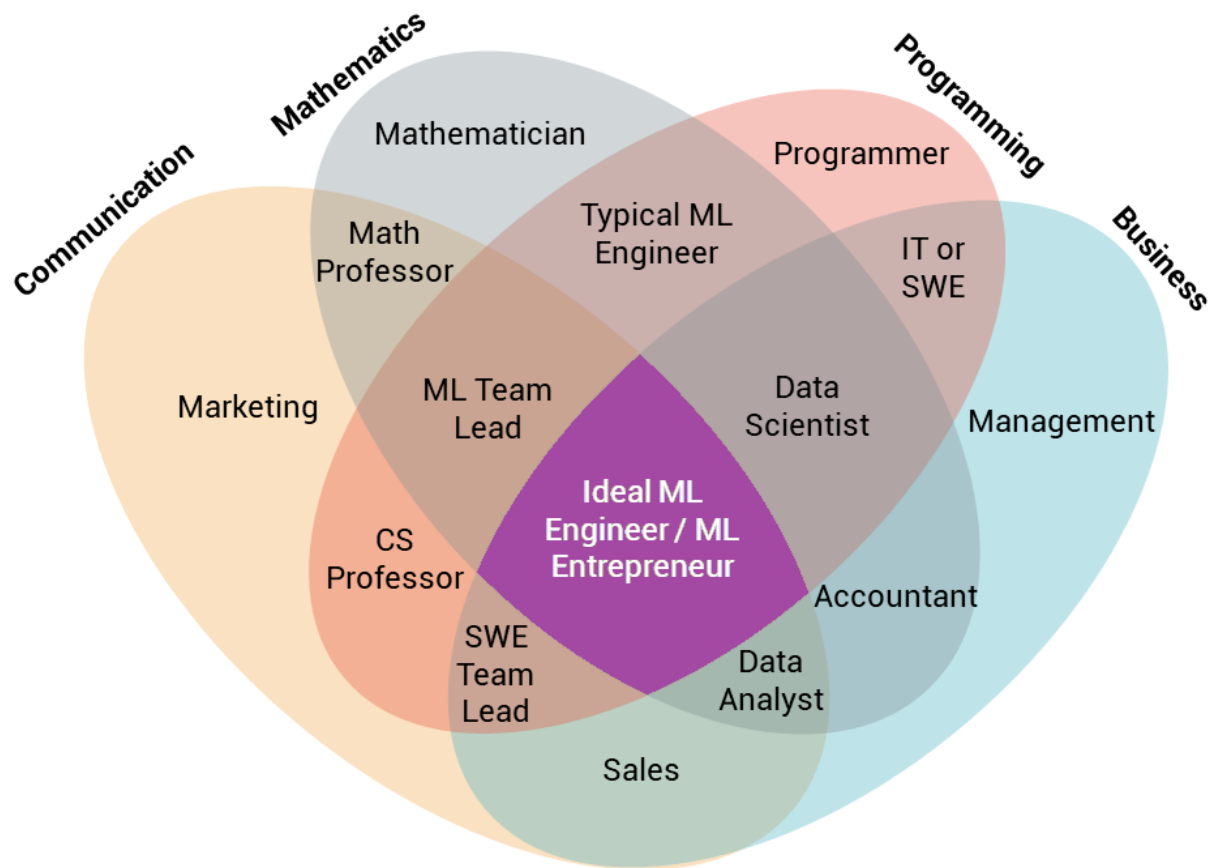
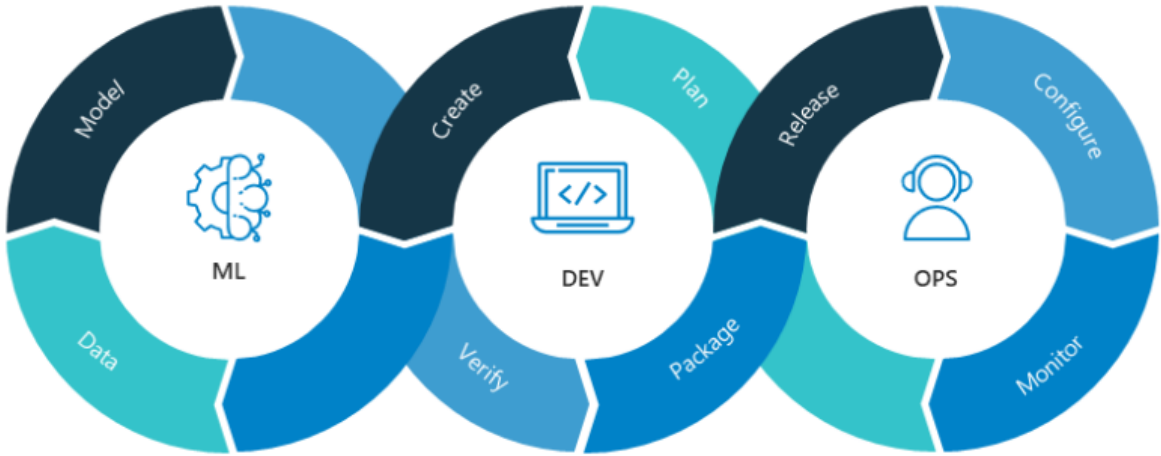
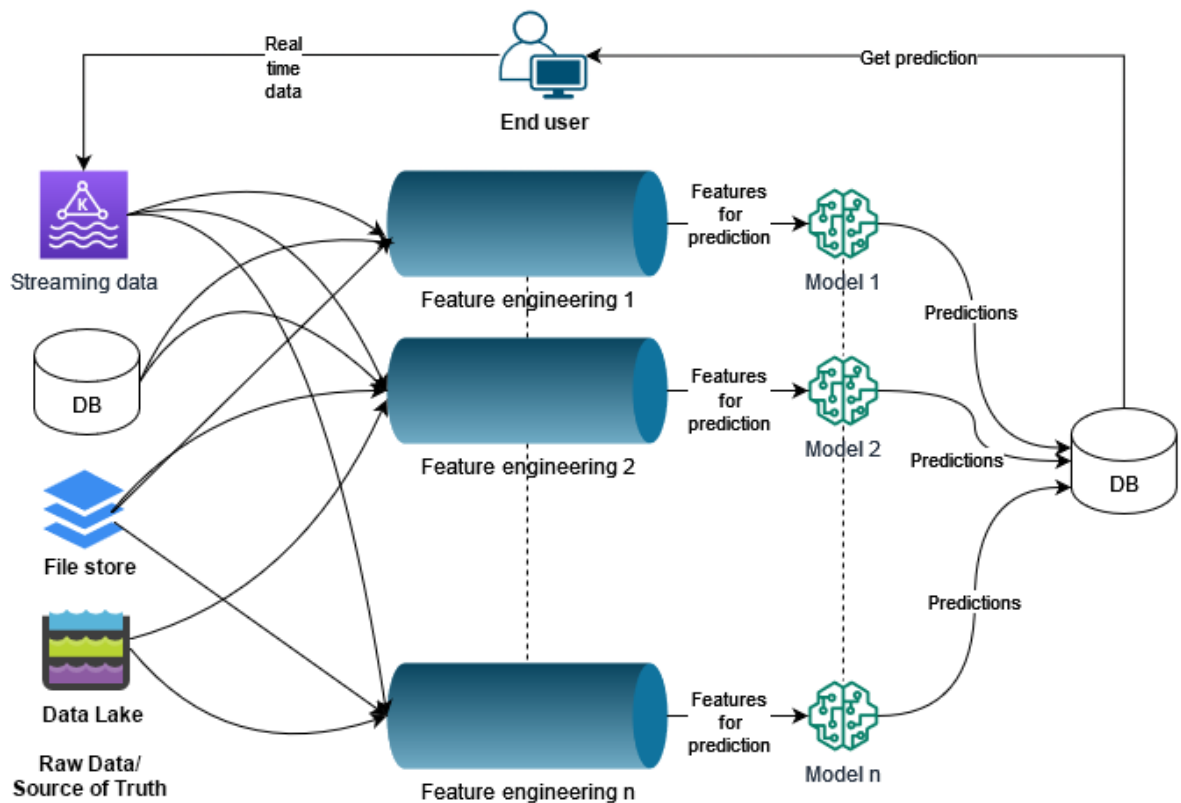
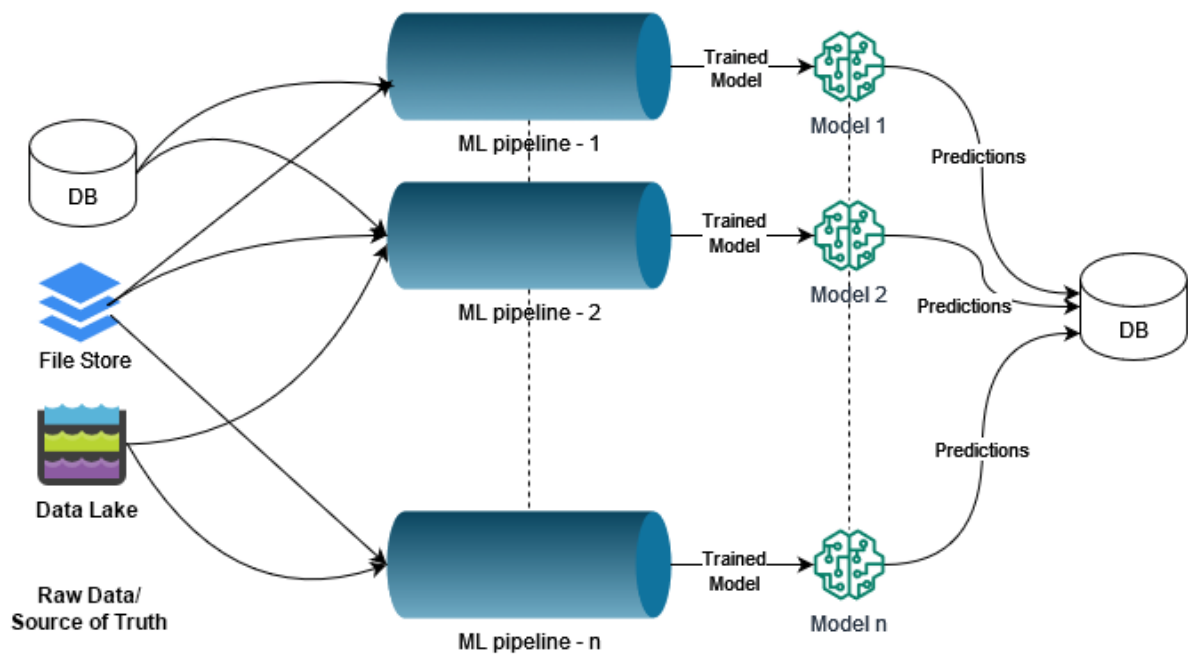




Chapter 1: Overview of the Machine Learning Life Cycle





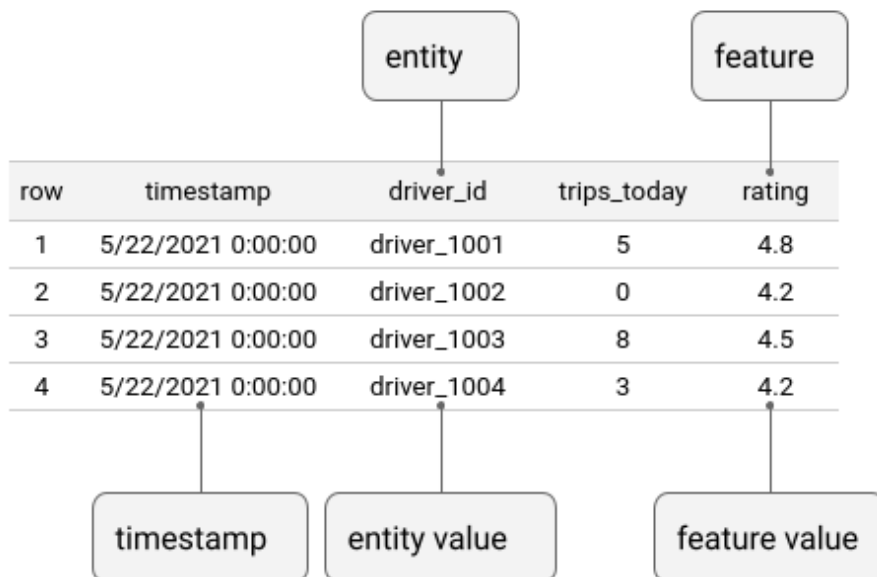
Chapter 2: What Do Feature Stores Solve?





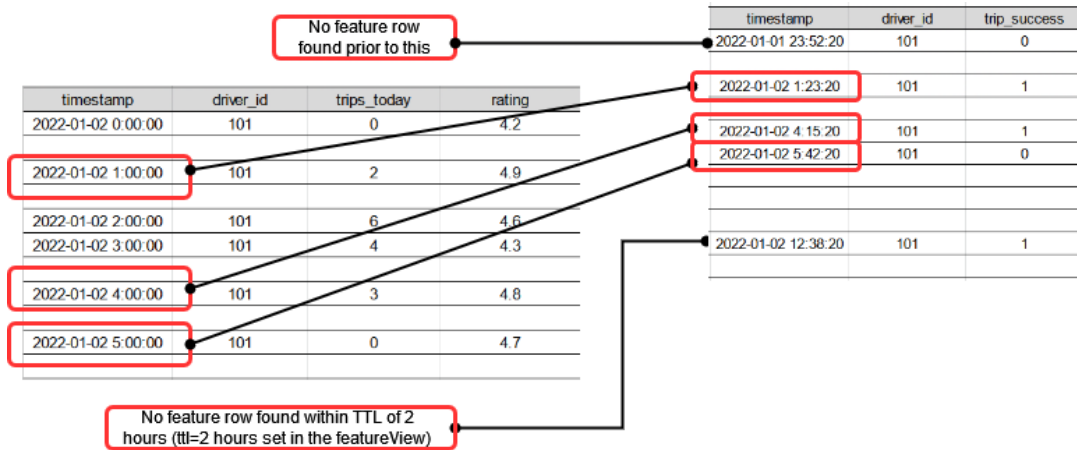
	iPhone 13 Pro	Google Pixel 6 Pro
Size	146.7 x 71.5 x 7.7 mm (5.78 x 2.81 x 0.30 inches)	163.9 x 75.9 x 8.9 mm (6.45 x 2.99 x 0.35 inches)
Weight	204 grams (7.2 ounces)	210 grams (7.41 ounces)
Screen size	6.1-inch Super Retina XDR OLED	6.71-inch LTPO AMOLED
Screen resolution	2532 x 1170 pixels (460 ppi)	3120 x 1440 pixels (512 ppi)
Operating system	iOS 15	Android 12
Storage	128 GB, 256 GB, 512 GB, 1TB	128 GB, 256 GB, 512 GB
MicroSD card slot	No	No
Processor	Apple A15 Bionic (5nm)	Google Tensor (5nm)
RAM	6 GB	12 GB
Camera	Triple lens 12-megapixel wide, 12 MP ultrawide, and 12 MP telephoto rear, 12 MP front	Triple lens 50MP wide, 12 MP ultrawide, and 48 MP telephoto rear, 11.1 MP front
Video	4K at up to 60 fps, 1080p at 120 fps	4K at up to 60 fps, 1080p at 240 fps
Bluetooth version	Bluetooth 5.1	Bluetooth 5.2
Ports	Lightning	USB-C
Fingerprint sensor	No, FaceID instead	Yes (in-display ultrasonic)
Water resistance	IP68	IP68
Battery	3,125 mAh 20 W wired charging (no charger included in the box) 15 W MagSafe charging 7.5 W wireless charging	5,003 mAh 30 W wired charging (no charger included in the box) 23 W wireless charging
App marketplace	Apple App Store	Google Play Store
Network support	All major carriers	All major carriers
Colors	Graphite, Gold, Silver, Sierra Blue	Cloudy White, Sorta Sunny, Stormy Black
Price	Starting at \$999	Starting at \$899
Buy from	Apple	Google
Review score	4.5 out of 5 stars	4 stars out of 5

Chapter 3: Feature Store Fundamentals, Terminology, and Usage

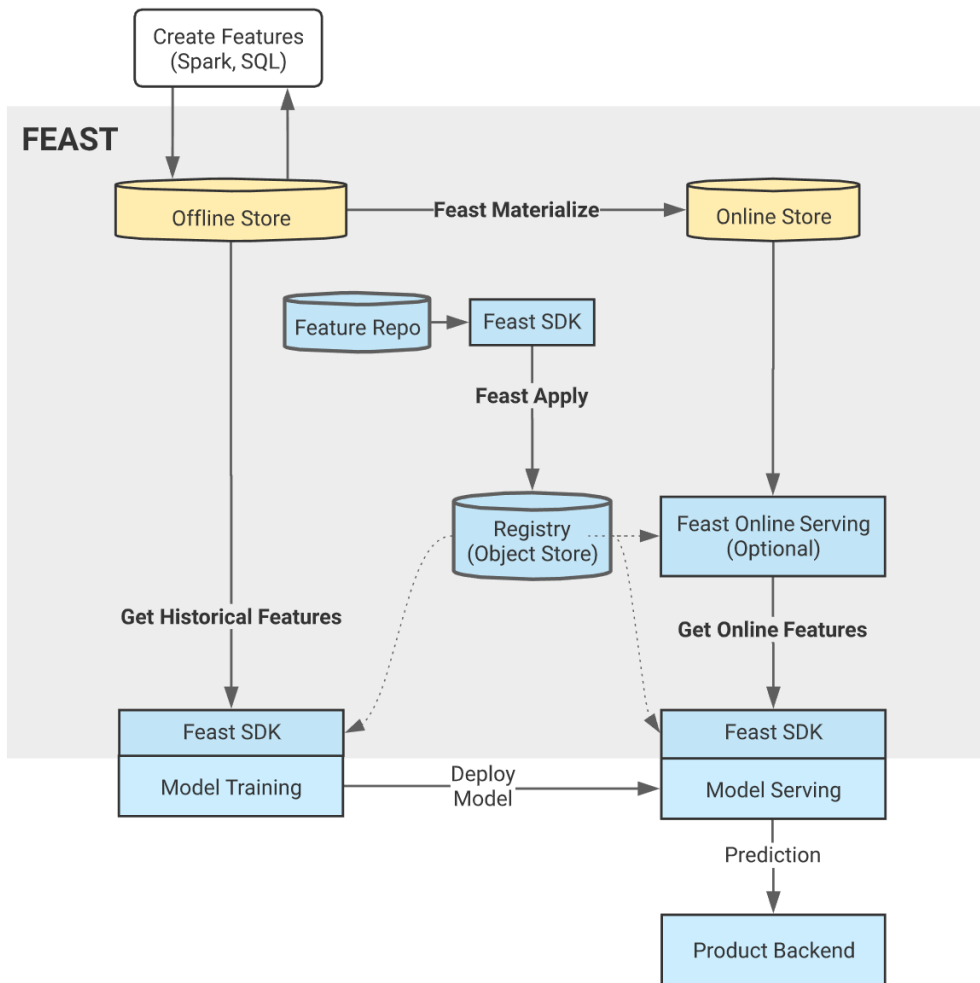


timestamp	driver_id	trips_today	rating
2022-01-02 0:00:00	101	0	4.2
2022-01-02 1:00:00	101	2	4.9
2022-01-02 2:00:00	101	6	4.6
2022-01-02 3:00:00	101	4	4.3
2022-01-02 4:00:00	101	3	4.8
2022-01-02 5:00:00	101	0	4.7

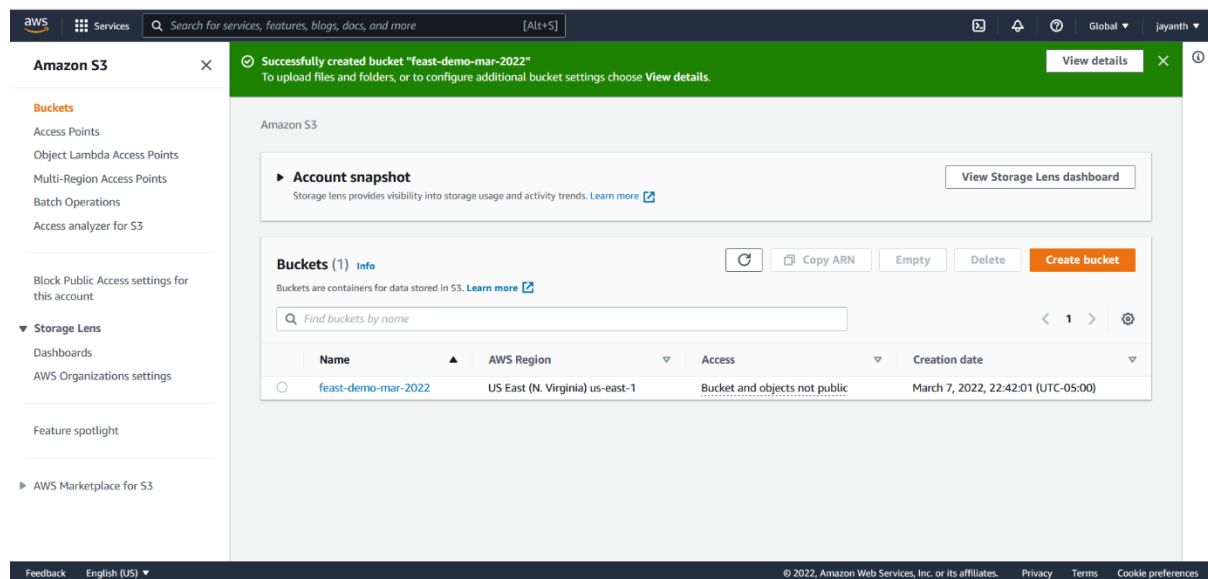
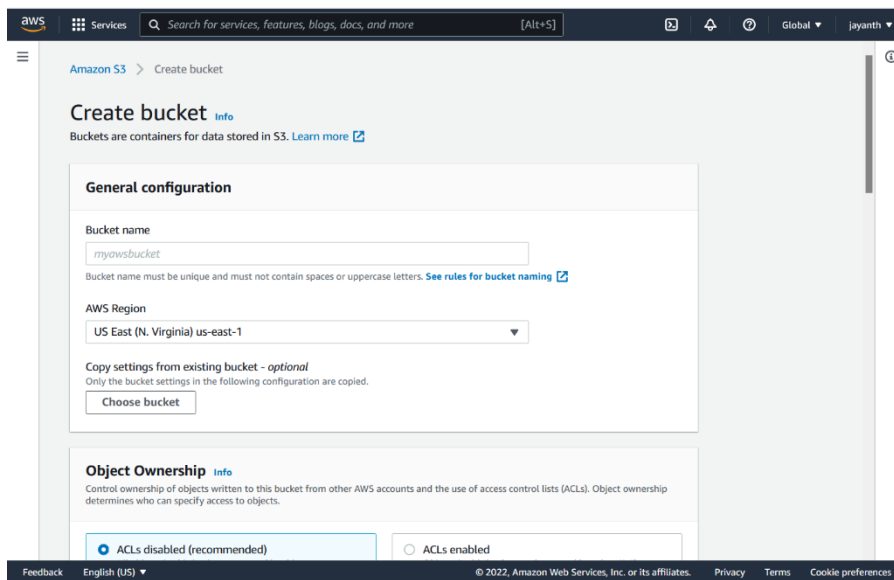
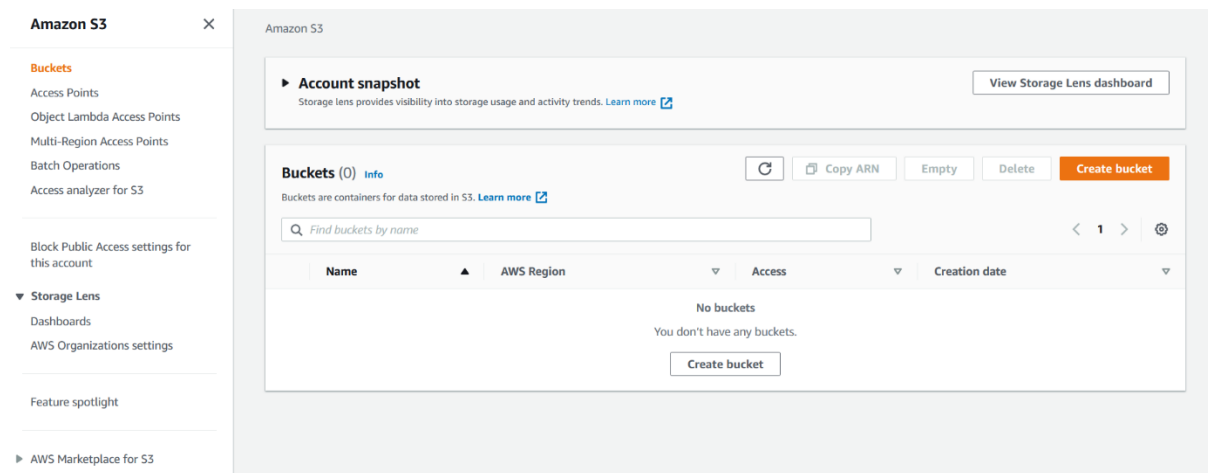
timestamp	driver_id	trip_success
2022-01-01 23:52:20	101	0
2022-01-02 1:23:20	101	1
2022-01-02 4:15:20	101	1
2022-01-02 5:42:20	101	0
2022-01-02 12:38:20	101	1



timestamp	driver_id	trip_success	trips_today	rating
2022-01-01 23:52:20	101	0	NULL	NULL
2022-01-02 1:23:20	101	1	2	4.9
2022-01-03 4:15:20	101	1	3	4.8
2022-01-04 5:42:20	101	0	0	4.7
2022-01-05 12:38:20	101	1	NULL	NULL



Chapter 4: Adding Feature Store to ML Models



aws Services Search for services, features, blogs, docs, and more [Alt+S] Global jayanth

Identity and Access Management (IAM)

Introducing the new IAM roles experience
We've redesigned the IAM roles experience to make it easier to use. [Let us know what you think.](#)

Search IAM

Dashboard

Access management

- User groups
- Users
- Roles**
- Policies
- Identity providers
- Account settings

Access reports

- Access analyzer
- Archive rules
- Analyzers
- Settings
- Credential report
- Organization activity
- Service control policies (SCPs)

IAM > Roles

Roles (2) Info

An IAM role is an identity you can create that has specific permissions with credentials that are valid for short durations. Roles can be assumed by entities that you trust.

Search

<input type="checkbox"/>	Role name	Trusted entities	Last activity
<input type="checkbox"/>	AWSServiceRoleForSupport	AWS Service: support (Service-Linked Role)	-
<input type="checkbox"/>	AWSServiceRoleForTrustedAdvisor	AWS Service: trustedadvisor (Service-Linked Role)	-

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] Global jayanth

IAM > Roles > Create role

Step 1
Select trusted entity

Step 2
Add permissions

Step 3
Name, review, and create

Select trusted entity

Trusted entity type

- AWS service**
Allow AWS services like EC2, Lambda, or others to perform actions in this account.
- AWS account**
Allow entities in other AWS accounts belonging to you or a 3rd party to perform actions in this account.
- Web identity**
Allows users federated by the specified external web identity provider to assume this role to perform actions in this account.
- SAML 2.0 federation**
Allow users federated with SAML 2.0 from a corporate directory to perform actions in this account.
- Custom trust policy**
Create a custom trust policy to enable others to perform actions in this account.

Custom trust policy
Create a custom trust policy to enable others to perform actions in this account.

```

1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Effect": "Allow",
6       "Principal": {},
7       "Action": "sts:AssumeRole"
8     }
9   ]

```

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] Global jayanth

IAM > Roles > Create role

Step 1
Select trusted entity

Step 2
Add permissions

Step 3
Name, review, and create

Add permissions

Permissions policies (Selected 2/733)

Choose one or more policies to attach to your new role.

Filter policies by property or policy name and press enter

<input type="checkbox"/>	Policy name	Type	Description
<input type="checkbox"/>	AWSDirectConnectReadOnlyAccess	AWS m...	Provides read only access to AWS Direct Connect via the AWS Management C...
<input type="checkbox"/>	AmazonGlacierReadOnlyAccess	AWS m...	Provides read only access to Amazon Glacier via the AWS Management Cons...
<input type="checkbox"/>	AWSMarketplaceFullAccess	AWS m...	Provides the ability to subscribe and unsubscribe to AWS Marketplace software...
<input type="checkbox"/>	AWSSSODirectoryAdministrator	AWS m...	Administrator access for SSO Directory
<input type="checkbox"/>	AWSIoT1ClickReadOnlyAccess	AWS m...	Provides read only access to AWS IoT 1-Click.
<input type="checkbox"/>	AutoScalingConsoleReadOnlyAccess	AWS m...	Provides read-only access to Auto Scaling via the AWS Management Console.
<input type="checkbox"/>	AmazonDMSRedshiftS3Role	AWS m...	Provides access to manage S3 settings for Redshift endpoints for DMS.
<input type="checkbox"/>	AWSQuickSightListIAM	AWS m...	Allow QuickSight to list IAM entities
<input type="checkbox"/>	AWSHealthFullAccess	AWS m...	Allows full access to the AWS Health Apis and Notifications and the Personal H...

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] Global jayanth

IAM > Roles > Create role

Step 1
Select trusted entity

Step 2
Add permissions

Step 3
Name, review, and create

Name, review, and create

Role details

Role name
Enter a meaningful name to identify this role.

feast-demo-mar-2022-spectrum-role
Maximum 128 characters. Use alphanumeric and "+=, @, _" characters.

Description
Add a short explanation for this policy.

Maximum 1000 characters. Use alphanumeric and "+=, @, _" characters.

Step 1: Select trusted entities Edit

```

1- {
2-   "Version": "2012-10-17",
3-   "Statement": [
4-     {
5-       "Effect": "Allow",
6-       "Principal": {

```

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Amazon Redshift provisioned clusters

Redshift serverless **New**

Provisioned clusters dashboard

Clusters

- Reserved nodes
- Snapshots

Query editor

Query editor v2

Queries and loads

Datashares

Configurations

Amazon Redshift > Clusters

In my account From other accounts

Connect to Redshift clusters

Clusters (0) Info Refresh Query data Actions Create cluster

Filter clusters by property or value

Cluster	Cluster namespace	Status	Storage capacity us...	CPU utilization	Snapshots	Not
No clusters Create an Amazon Redshift cluster Create cluster						

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Amazon Redshift provisioned clusters

Redshift serverless **New**

Provisioned clusters dashboard

Clusters

- Reserved nodes
- Snapshots

Query editor

Query editor v2

Queries and loads

Datashares

Configurations

- Workload management
- Subnet groups

Amazon Redshift > Clusters > Create cluster

Create cluster Info

Cluster configuration

Cluster identifier
This is the unique key that identifies a cluster.

feast-demo-mar-2022
The identifier must be from 1-63 characters. Valid characters are a-z (lowercase only) and - (hyphen).

What are you planning to use this cluster for?

Production
Configure for fast and consistent performance at the best price.

Free trial
Configure for learning about Amazon Redshift. This configuration is free for a limited time if your organization has never created an Amazon Redshift cluster.

Info When the free trial ends, delete your cluster to avoid incurring charges at on-demand rate for compute and storage. If you want to take a final snapshot of your cluster and store the snapshot on an S3, our on-demand rate applies.

Calculated configuration summary
dc2.large | 1 node

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Sample data is loaded with your Redshift cluster.

Tickit (28 MB)
 Tickit is the sample data set that uses a sample database called TICKIT. Tickit contains individual sample data files: two fact tables and five dimensions.

Database configurations

Admin user name
 Enter a login ID for the admin user of your DB instance.

 The name must be 1-128 alphanumeric characters, and it can't be a [reserved word](#).

Auto generate password
 Amazon Redshift can generate a password for you, or you can specify your own password.

Admin user password

 Show password
 Must be 8-64 characters long. Must contain at least one uppercase letter, one lowercase letter and one number. Can be any printable ASCII character except "/", "*", or "@".

Cancel **Create cluster**

Amazon Redshift query editor v2 is now available
 Query editor v2 provides new features such as multistatement query execution, query parameterization, query versioning, visualizations, and query sharing. [Learn more](#)

Go to query editor v2

Amazon Redshift > Clusters > feast-demo-mar-2022

feast-demo-mar-2022

Actions Edit Add partner integration **Query data**

General information

Cluster identifier feast-demo-mar-2022	Status Available	Node type dc2.large	Endpoint feast-demo-mar-2022
Cluster namespace [REDACTED]	Date created March 10, 2022, 20:26 (UTC-05:00)	Number of nodes 1	JDBC URL jdbc:redshift://feast-demo-mar-2022
	Storage used 0.21% (0.34 of 160 GB used)	AQUA Not available	ODBC URL Driver={Amazon Redshift (x64)}; Server=feast-...

Cluster performance Query monitoring Schedules Maintenance **Properties**

arn Highlight All Match Case Match Duplicates Whole Words 1 of 3 matches Reached end of page, continued from top

Create, associate, or remove an IAM role. You can associate up to 50 IAM roles. You can also choose an IAM role and set it as the default for this cluster.

Associated IAM roles (0) info

Search for associated IAM role by name, status, or role type

Set default Manage IAM roles

IAM roles	Status	Role type
No resources No associated IAM roles		

Associate IAM role

Granted accounts (0)
 VPCs in other accounts that are allowed to access this cluster. [Learn more](#)

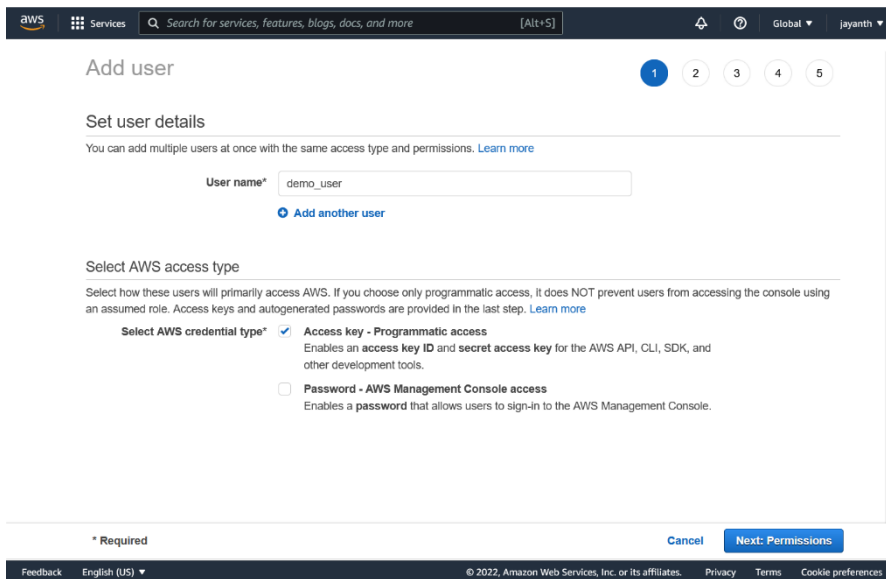
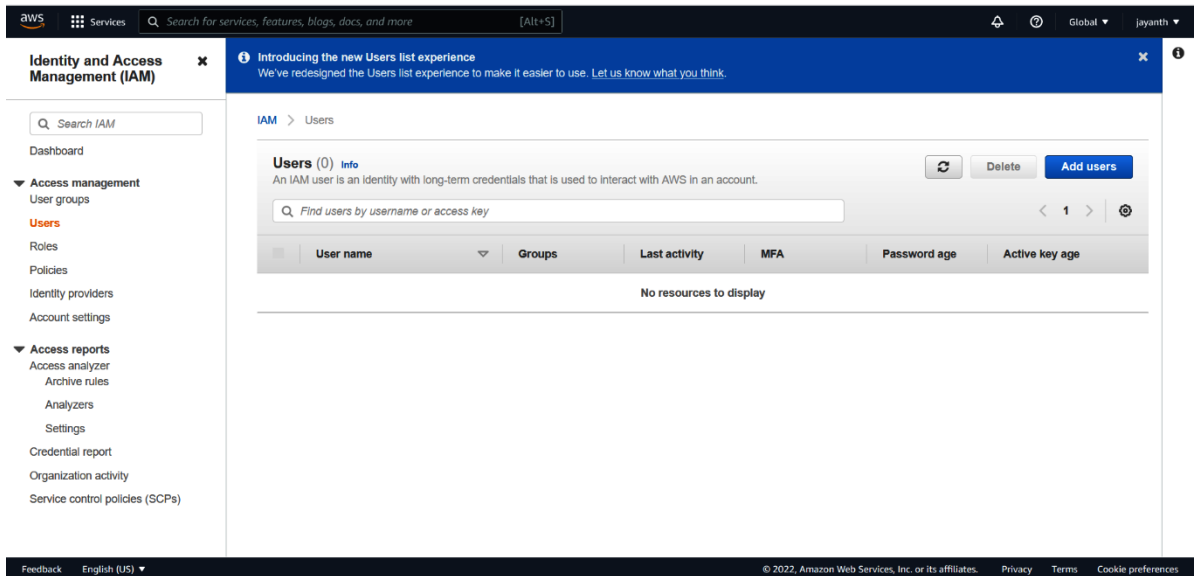
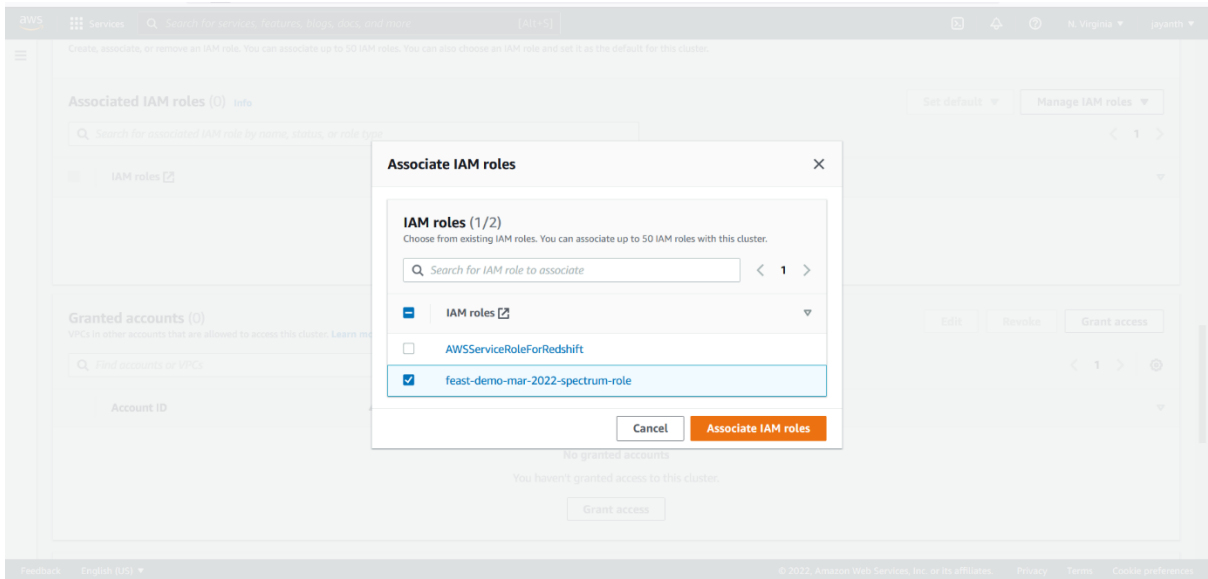
Edit Revoke Grant access

Find accounts or VPCs

Account ID	VPC	Endpoints created
No granted accounts You haven't granted access to this cluster.		

Grant access

arn Highlight All Match Case Match Duplicates Whole Words 1 of 3 matches Reached end of page, continued from top



aws Services Search for services, features, blogs, docs, and more [Alt+S] Global jayanth

Add user

1 2 3 4 5

Set permissions

Add user to group Copy permissions from existing user Attach existing policies directly

Get started with groups
 You haven't created any groups yet. Using groups is a best-practice way to manage users' permissions by job functions, AWS service access, or your custom permissions. Get started by creating a group. [Learn more](#)

Create group

Set permissions boundary

Cancel Previous Next: Tags

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] Global jayanth

Add user

1 2 3 4 5

Review

Review your choices. After you create the user, you can view and download the autogenerated password and access key.

User details

User name demo_user
 AWS access type Programmatic access - with an access key
 Permissions boundary Permissions boundary is not set

Permissions summary

The following policies will be attached to the user shown above.

Type	Name
Managed policy	AmazonRedshiftFullAccess
Managed policy	AmazonS3FullAccess
Managed policy	AmazonDynamoDBFullAccess

Tags

No tags were added.

Cancel Previous Create user

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] Global jayanth

Add user

1 2 3 4 5

Success
 You successfully created the users shown below. You can view and download user security credentials. You can also email users instructions for signing in to the AWS Management Console. This is the last time these credentials will be available to download. However, you can create new credentials at any time.

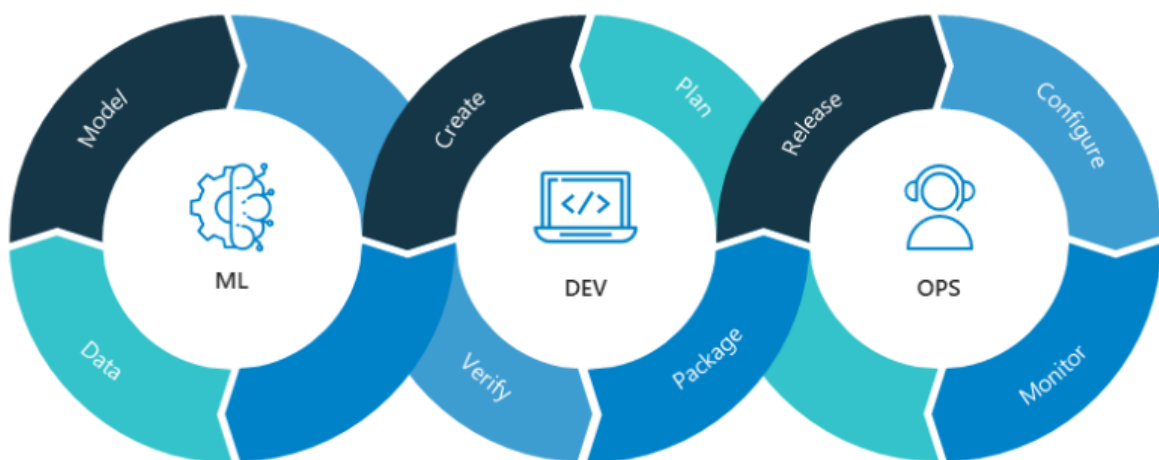
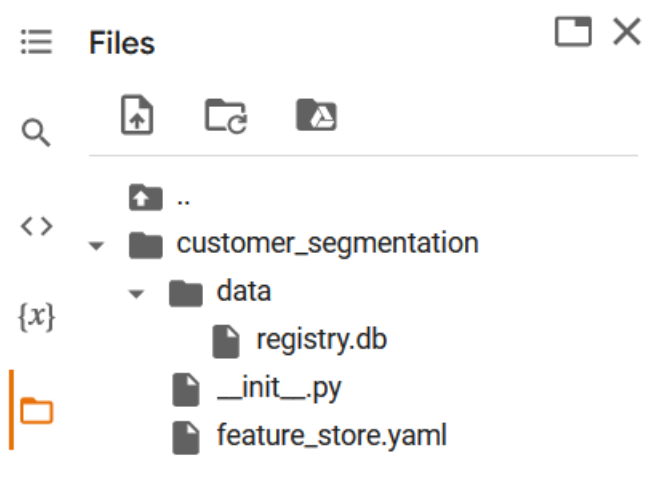
Users with AWS Management Console access can sign-in at: <https://456464162939.signin.aws.amazon.com/console>

Download .csv

User	Access key ID	Secret access key
demo_user	AKIAWUR3CMB5Q5H7ZEJL	***** Show

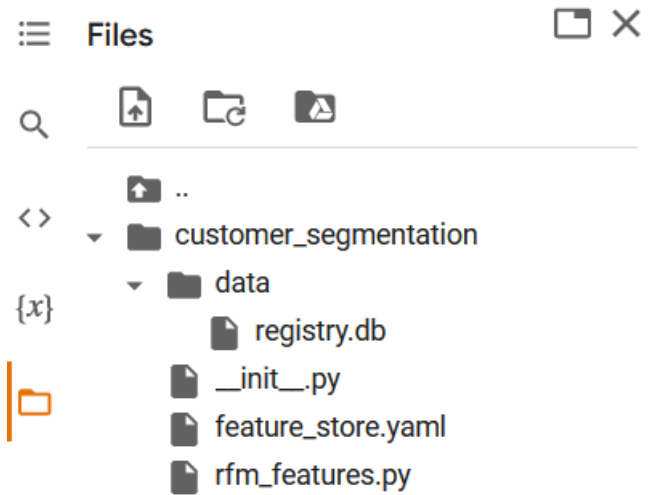
Close

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences



customerid	recency	frequency	monetaryvalue	r	f	m	rfmscore	revenue6m	ltvcluster	segmenthighvalue	segmentlowvalue	segmentmidvalue
0	12747.0	7	35	1082.09	3	2	3	8	1666.11	1	1	0
1	12748.0	1	582	4336.73	3	3	3	9	18679.01	2	1	0
2	12749.0	8	54	782.10	3	3	3	9	2323.04	1	1	0
4	12823.0	63	1	459.00	0	0	2	2	765.00	1	0	1
7	12836.0	28	62	814.71	1	3	3	7	951.46	1	1	0

CustomerID	Recency	Frequency	MonetaryValue	R	F	M	RFMScore	Revenue6m	LTVCcluster	SegmentHighValue	SegmentLowValue	SegmentMidValue
0	12747.0	7	35	1082.09	3	2	3	8	1666.11	0	1	0
1	12748.0	1	582	4336.73	3	3	3	9	18679.01	1	1	0
2	12749.0	8	54	782.10	3	3	3	9	2323.04	0	1	0
4	12823.0	63	1	459.00	0	0	2	2	765.00	0	0	1
7	12836.0	28	62	814.71	1	3	3	7	951.46	0	1	0



aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

AWS Lake Formation

Dashboard

Data catalog

Databases

Tables

Data filters

Settings

Register and ingest

Data lake locations

Blueprints

Crawlers

Jobs

Permissions

Administrative roles and tasks

LF-Tags

LF-tag permissions

Data lake permissions

Data locations

External data filtering

AWS Lake Formation > Databases

Databases

Find databases

Name	Owner account ID	Shared resource	Shared resource owner	Amazon S3 path	Description
No databases					

Create database

© 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Database details

Create a database in the AWS Glue Data Catalog.

Database
Create a database in my account.

Resource link
Create a resource link to a shared database.

Name

dev

Location - optional
Choose an Amazon S3 path for this database, which eliminates the need to grant data location permissions on catalog table paths that are this location's children

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

Description - optional
Enter a description

Descriptions can be up to 2048 characters long.

Default permissions for newly created tables
This setting maintains existing AWS Glue Data Catalog behavior. You can still set individual permissions, 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 tables in this database

Cancel Create database

© 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

AWS Lake Formation × AWS Lake Formation > Tables

Tables ↻ Actions ▼ Create table using a crawler ↗ Create table

Find table by properties < 1 > ⊙

Name	Database	Governance	Owner account ID	Shared resou...	Shared resource ow...	Location
No tables						

Create table

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

AWS Lake Formation > Tables > Create table

Create table

Table details
Create a table in the AWS Glue Data Catalog.

Table
Create a table in my account.
 Resource link
Create a resource link to a shared table.

Name

Name may contain letters (A-Z), numbers (0-9), hyphens (-), or underscores (_), and must be less than 256 characters long.

Database
Table is contained within this database.
 ↻
[Create database](#)

Description - optional

Descriptions can be up to 2048 characters long.

Data management and security

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Data management and security

ⓘ **Governance choice cannot be changed after table creation**
 Governed tables support ACID (atomic, consistent, isolated, durable) transactions to guarantee data integrity with concurrent changes across multiple tables. Governed tables also come with automatic data layout optimizations by compacting small and delta files in order to maximize query performance.

Enable governed data access and management
Check this option to enable transactional workloads and automatic data compaction.

Automatic compaction
Enable this option to turn on automatic compaction on a governed table.

Data store

Data is located in
 Specified path in my account
 Specified path in another account

Include path
Path must be in the form s3://bucket/prefix. It must end with a slash (/) and not include any files.
 Browse

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia Jayanth

Browse

Data format

Classification

- Avro
- CSV
- JSON
- XML
- PARQUET
- ORC

Choose the format of the data in your table.

Schema Upload Schema Delete Edit Add column

< 1 > ⊗

No available schema

Cancel Submit

#	Column Name	Data type	Partition key	Comment	LF-Tags
1	customerid	string	-	-	-
2	recency	bigint	-	-	-
3	frequency	bigint	-	-	-
4	monetaryvalue	double	-	-	-
5	r	bigint	-	-	-
6	f	bigint	-	-	-
7	m	bigint	-	-	-
8	rfmscore	bigint	-	-	-
9	revenue6m	double	-	-	-
10	ltvcluster	int	-	-	-
11	segmenthighvalue	int	-	-	-
12	segmentlowvalue	int	-	-	-
13	segmentmidvalue	int	-	-	-
14	event_timestamp	timestamp	-	-	-
15	created_timestamp	timestamp	-	-	-

Amazon Redshift query editor v2 is now available
 Query editor v2 provides new features such as multistatement query execution, query parameterization, query versioning, visualizations, and query sharing. [Learn more](#)

Amazon Redshift > Clusters > feast-demo-mar-2022

feast-demo-mar-2022

Actions Edit Add partner integration Query data

Query in query editor
 Query in query editor v2

General information

Cluster identifier feast-demo-mar-2022	Status Available	Node type dc2.large	Endpoint [REDACTED]
Cluster namespace [REDACTED]	Date created March 13, 2022, 14:24 (UTC-04:00)	Number of nodes 1	JDBC URL [REDACTED]
	Storage used 0.21% (0.34 of 160 GB used)	AQUA Not available	ODBC URL Driver=(Amazon Redshift (x64)); Serv...

Cluster performance Query monitoring Schedules Maintenance Properties

Recommendations (0)

Redshift query editor v2

Cluster: feast-demo-mar-2022 (awsuser) Database: dev

Untitled 1 x

Run Limit 100 Explain Save Shortcuts

1

Identity and Access Management (IAM)

Search IAM

Dashboard

- Access management
 - User groups
 - Users
 - Roles**
 - Policies
 - Identity providers
 - Account settings
- Access reports
 - Access analyzer
 - Archive rules
 - Analyzers
 - Settings
- Credential report
- Organization activity
- Service control policies (SCPs)

IAM > Roles > feast-demo-mar-2022-spectrum-role

feast-demo-mar-2022-spectrum-role

Delete Edit

Summary

Creation date
March 06, 2022, 20:12 (UTC-05:00)

ARN
arn:aws:iam::[REDACTED]:role/feast-demo-mar-2022-spectrum-role

Last activity
None

Maximum session duration
1 hour

Permissions Trust relationships Tags Access Advisor Revoke sessions

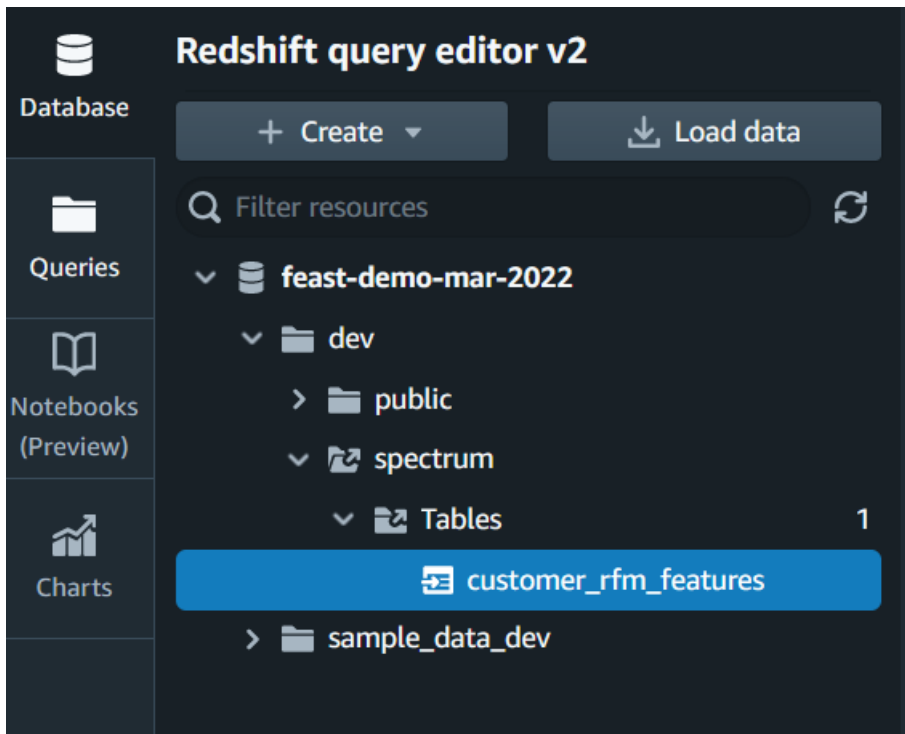
Permissions policies (2)

You can attach up to 10 managed policies.

Simulate Remove Add permissions

Filter policies by property or policy name and press enter

Policy name	Type	Description
AmazonS3FullAccess	AWS managed	Provides full access to all buckets v
AWSGlueConsoleFullAccess	AWS managed	Provides full access to AWS Glue v

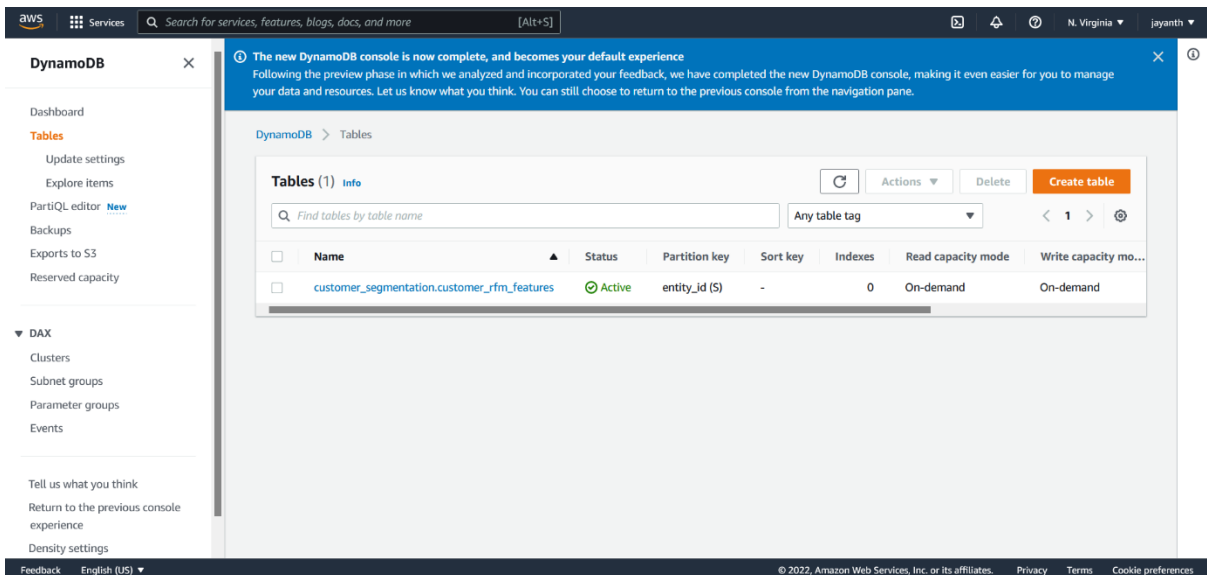


```

- /content/customer_segmentation
/usr/local/lib/python3.7/dist-packages/scipy/fft/__init__.py:97: DeprecationWarning: The module num
from numpy.dual import register_func
/usr/local/lib/python3.7/dist-packages/scipy/sparse/sputils.py:17: DeprecationWarning: `np.typeDict
supported dtypes = [np.typeDict[x] for x in supported dtypes]
/usr/local/lib/python3.7/dist-packages/feast/infra/offline_stores/redshift_source.py:55: Deprecati
DeprecationWarning,
03/14/2022 03:01:38 AM INFO:Found credentials in environment variables.
Created data source
Created entity customer
Created feature view customer_rfm_features

```

Deploying infrastructure for **customer_rfm_features**



Redshift query editor v2

Cluster: feast-demo-mar-2022 (awsuser) Database: dev

Query: 1 SELECT * from spectrum.customer_rfm_features limit 5

Result 1 (5)

customerid	recency	frequency	monetaryvalue	r	f
12747.0	7	35	1082.09	3	2
12748.0	1	582	4336.73	3	3
12749.0	8	54	782.1	3	3
12823.0	63	1	459	0	0
12836.0	28	62	814.71	1	3

Elapsed time: 744 ms Total rows: 5

↳

	CustomerID	event_timestamp
0	12747.0	2022-03-15 00:58:45.594038
1	12748.0	2022-03-15 00:58:45.594038
2	12749.0	2022-03-15 00:58:45.594038

	customerid	event_timestamp	recency	frequency	monetaryvalue	r	f	m
0	12747.0	2022-03-15 00:58:45.594038	7	35	1082.09	3	2	3
1	12749.0	2022-03-15 00:58:45.594038	8	54	782.10	3	3	3
2	12748.0	2022-03-15 00:58:45.594038	1	582	4336.73	3	3	3

Chapter 5: Model Training and Inference

	CustomerID	event_timestamp
0	14620.0	2022-03-26 16:24:21.490847
1	14740.0	2022-03-26 16:24:21.490847
2	13880.0	2022-03-26 16:24:21.490847
3	16462.0	2022-03-26 16:24:21.490847
4	17068.0	2022-03-26 16:24:21.490847

	customerid	event_timestamp	recency	monetaryvalue	r	m	rfmscore	segmenthighvalue	segmentlowvalue	segmentmidvalue	ltvcluster
0	12748.0	2022-03-26 18:05:06.863532	1.0	4336.73	3.0	3.0	9.0	1.0	0.0	0.0	1.0
2	12839.0	2022-03-26 18:05:06.863532	7.0	303.58	3.0	1.0	6.0	0.0	0.0	1.0	2.0
4	12878.0	2022-03-26 18:05:06.863532	45.0	854.99	1.0	3.0	6.0	0.0	0.0	1.0	0.0
5	12891.0	2022-03-26 18:05:06.863532	29.0	204.00	1.0	0.0	1.0	0.0	1.0	0.0	0.0
6	12916.0	2022-03-26 18:05:06.863532	10.0	1441.13	3.0	3.0	8.0	1.0	0.0	0.0	0.0

	precision	recall	f1-score	support
0.0	0.93	0.95	0.94	119
1.0	0.67	0.67	0.67	3
2.0	0.50	0.44	0.47	16
accuracy			0.88	138
macro avg	0.70	0.68	0.69	138
weighted avg	0.88	0.88	0.88	138

	customerid	event_timestamp	r	m	f	segmenthighvalue	segmentlowvalue	segmentmidvalue	ltvcluster
0	12748.0	2022-03-26 23:53:55.691200	3.0	3.0	3.0	1.0	0.0	0.0	1.0
2	12839.0	2022-03-26 23:53:55.691200	3.0	1.0	2.0	0.0	0.0	1.0	2.0
4	12878.0	2022-03-26 23:53:55.691200	1.0	3.0	2.0	0.0	0.0	1.0	0.0
5	12891.0	2022-03-26 23:53:55.691200	1.0	0.0	0.0	0.0	1.0	0.0	0.0
6	12916.0	2022-03-26 23:53:55.691200	3.0	3.0	2.0	1.0	0.0	0.0	0.0

	precision	recall	f1-score	support
0.0	0.86	1.00	0.93	119
1.0	0.00	0.00	0.00	1
2.0	0.00	0.00	0.00	18
accuracy			0.86	138
macro avg	0.29	0.33	0.31	138
weighted avg	0.74	0.86	0.80	138

FastAPI 0.1.0 OAS3

/openapi.json

default

- GET /ping Ping
- POST /invocations Inference

Schemas

HTTPValidationError >

ValidationError >

POST /invocations Inference

Parameters Cancel Reset

No parameters

Request body required application/json

```
{ "customer_list": ["12747.0", "12841.0"] }
```

Execute Clear

Responses

Curl

```
curl -X 'POST' \
  'http://localhost:8000/invocations' \
  -H 'accept: application/json' \
  -H 'Content-Type: application/json' \
  -d '{"customer_list":["12747.0", "12841.0"]}'
```

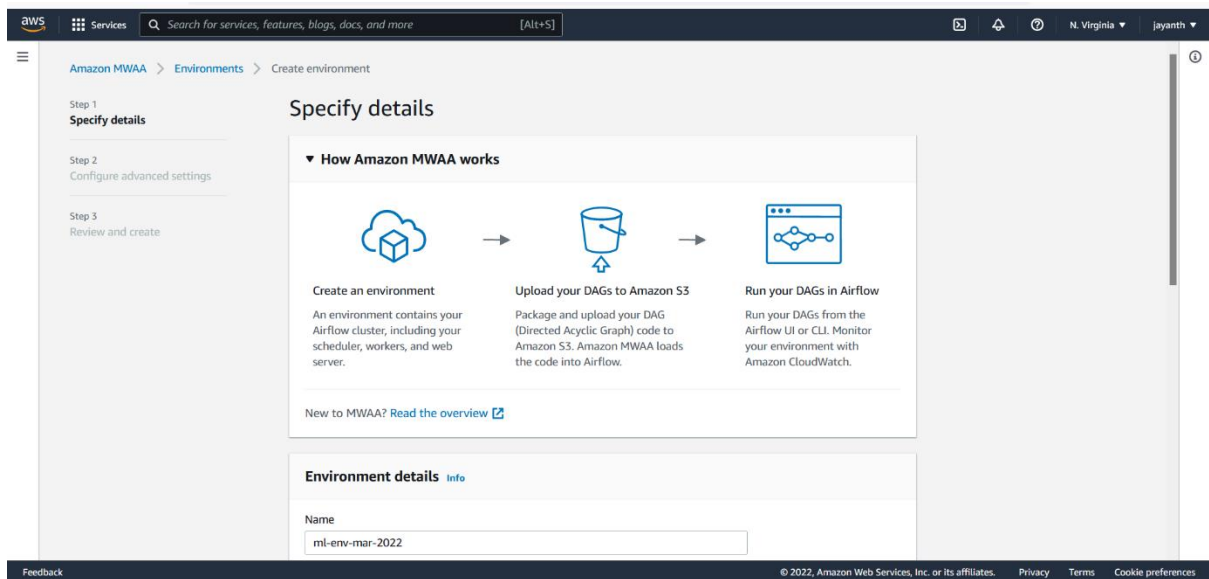
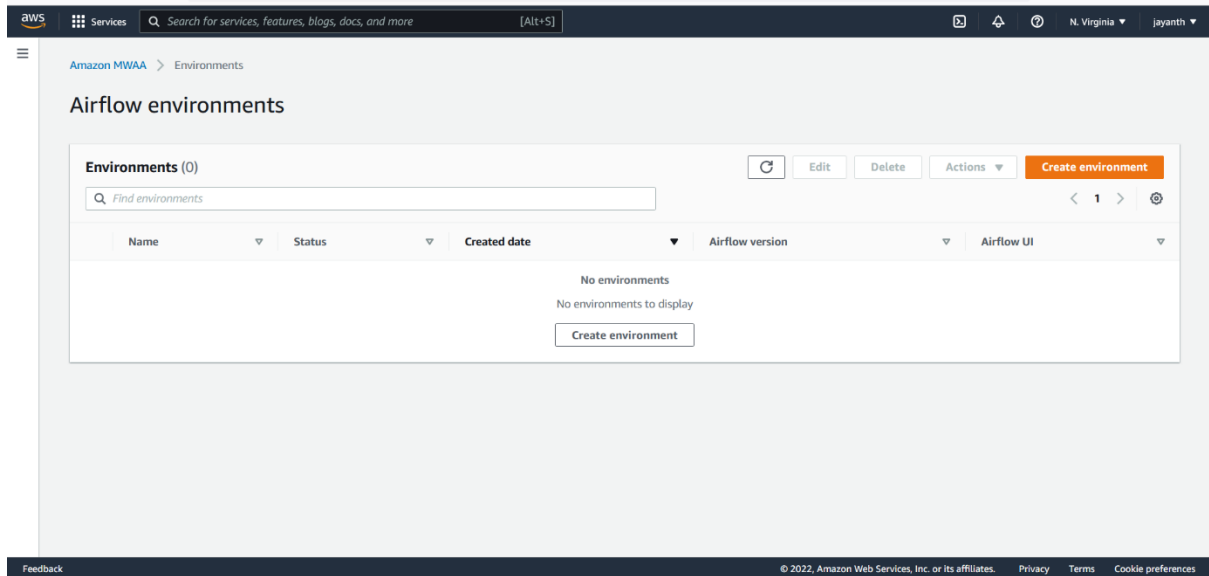
Request URL

```
http://localhost:8000/invocations
```

Server response

Code	Details
200	<p>Response body</p> <pre>{ "predictions": [0, 0] }</pre> <p>Response headers</p> <pre>content-length: 25 content-type: application/json date: Sun, 27 Mar 2022 23:28:41 GMT server: uvicorn</pre>

Chapter 6: Model to Production and Beyond



aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Create or specify an S3 bucket to store your DAG code. The bucket name must have versioning enabled. You can create a new bucket in the [Amazon S3 console](#)

S3 Bucket
The S3 bucket where your source code is stored. Enter an S3 URI or browse and select a bucket.

 Format: s3://mybucketname

DAGs folder
The S3 bucket folder that contains your DAG code. Enter an S3 URI or browse and select a folder.

 Format: s3://mybucketname/mydagfolder

Plugins file - optional
The S3 bucket ZIP file that contains your DAG plugins. Enter an S3 URI or browse and select a file object and version.

 Format: s3://mybucketname/myplugins.zip

Requirements file - optional
The S3 bucket file that contains your DAG requirements.txt. Enter an S3 URI or browse and select a file object and version.

 Format: s3://mybucketname/myrequirements.txt

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services vpc

Amazon MWAAs > Environments > Create environment

Step 1 Specify details
 Step 2 **Configure advanced settings**
 Step 3 Review and create

Configure advanced settings

Networking info

Virtual private cloud (VPC)
 Defines the networking infrastructure setup of your Airflow environment. An environment needs 2 private subnets in different availability zones. To create a new VPC with private subnets, choose Create MWAAs VPC. [Learn more](#)

Subnet 1
 Private subnet for the first availability zone. Each environment occupies 2 availability zones.

Subnet 2
 Private subnet for the second availability zone. Each environment occupies 2 availability zones.

Web server access

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Airflow DAGs Data Profiling Browse Admin Docs About 2022-04-12 22:26:12 UTC

DAGs

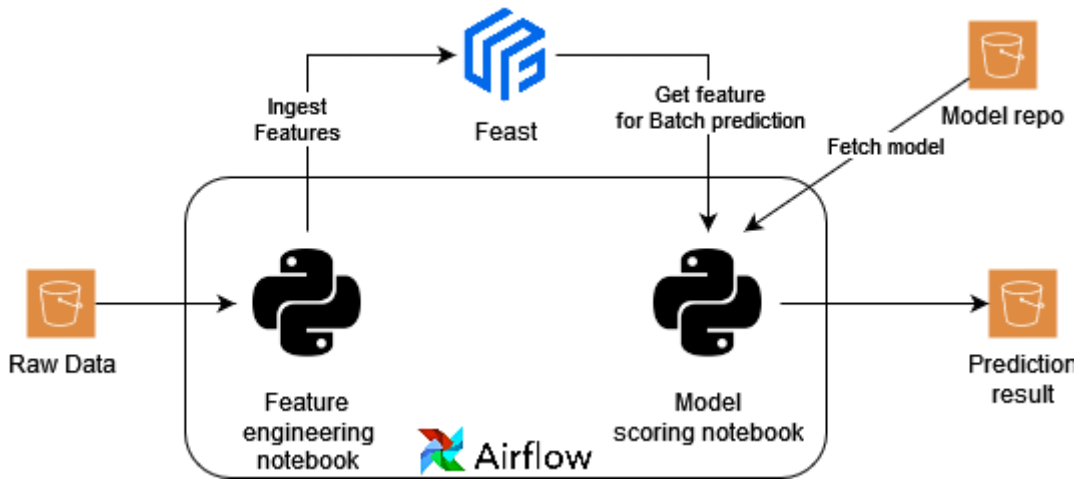
Search:

<input type="checkbox"/>	DAG	Schedule	Owner	Recent Tasks	Last Run	DAG Runs	Links
No data available in table							

Showing 0 to 0 of 0 entries

« ‹ › »

[Hide Paused DAGs](#)



Airflow DAGs Security Browse Admin Docs 00:23 UTC

DAGs

All 2 Active 1 Paused 1

Filter DAGs by tag Search DAGs

DAG	Owner	Runs	Schedule	Last Run	Next Run	Recent Tasks	Actions	Links
customer_segmentation_batch_model	airflow	1	@daily	2022-04-17, 00:19:56	2022-04-17, 00:00:00	...	▶ 🛑	...
notebook_test	airflow	0	@daily		2022-04-16, 00:00:00	...	▶ 🛑	...

Showing 1-2 of 2 DAGs

Airflow DAGs Security Browse Admin Docs 00:21 UTC

DAG: customer_segmentation_batch_model Batch model pipeline

success Schedule: @daily Next Run: 2022-04-17, 00:00:00

Tree Graph Calendar Task Duration Task Tries Landing Times Gantt Details Code

2022-04-17T00:19:57Z Runs 25 Run manual__2022-04-17T00:19:56.834354+00:00 Layout Left > Right Update Find Task...

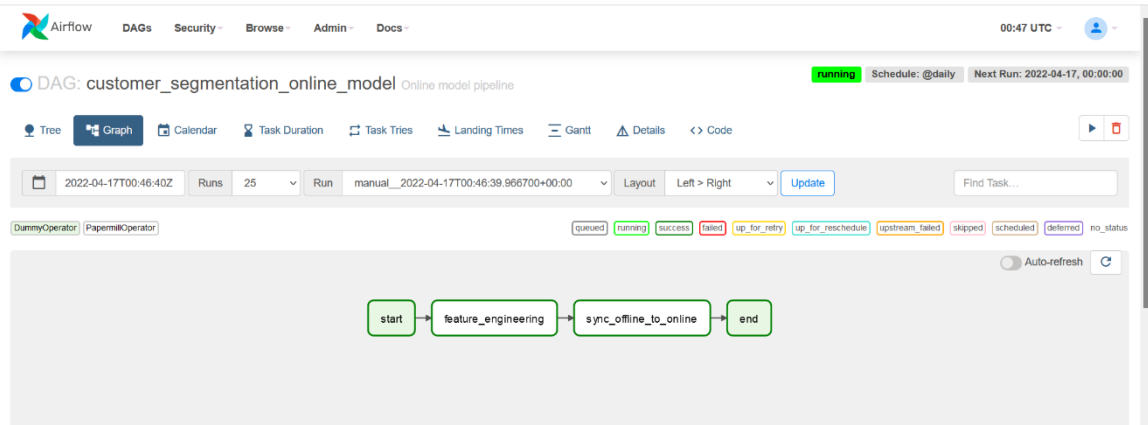
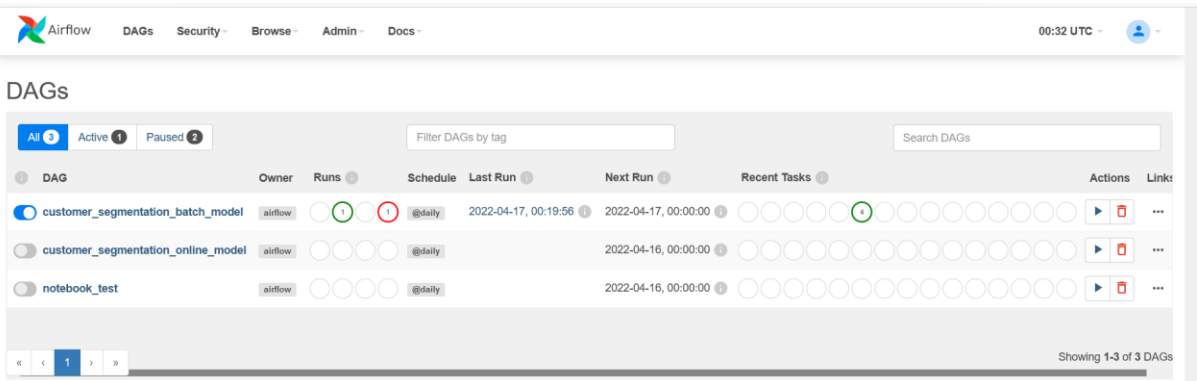
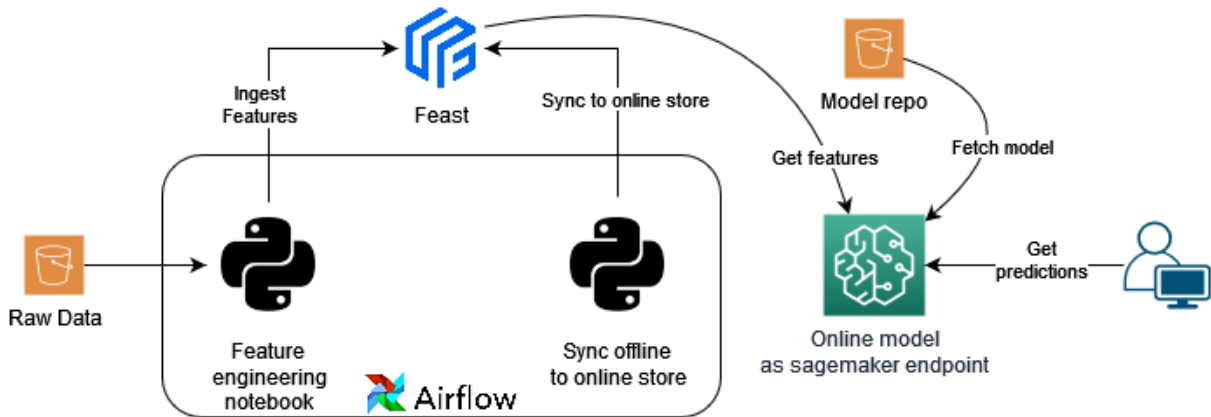
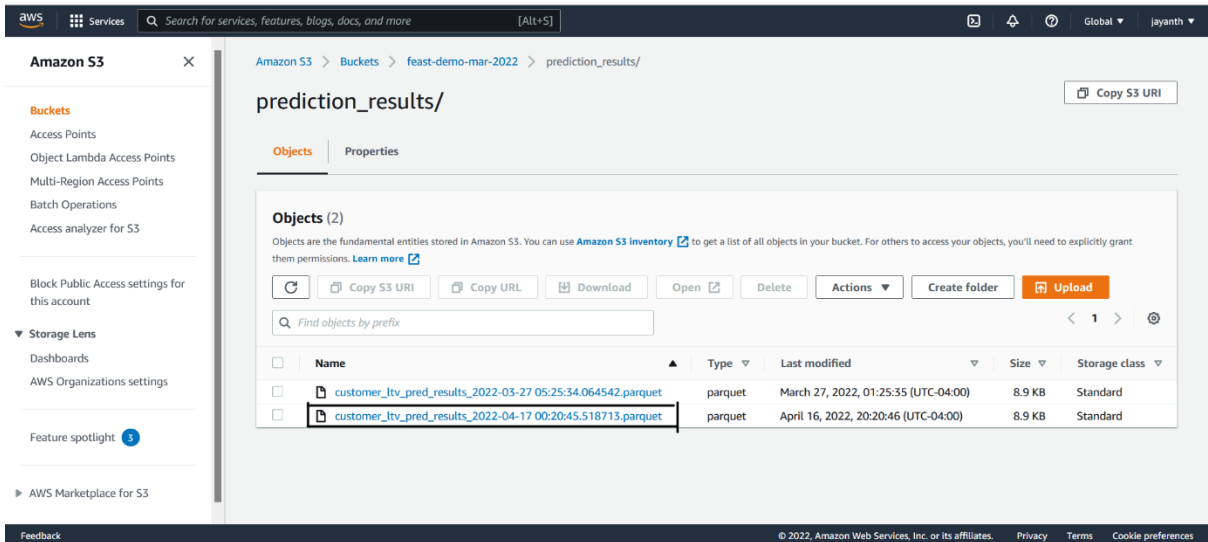
DummyOperator PapermillOperator

queued running success failed up_for_retry up_for_reschedule upstream_failed skipped scheduled deferred no_status

Auto-refresh

```

graph LR
    start --> feature_engineering
    feature_engineering --> model_prediction
    model_prediction --> end
  
```



aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Amazon Elastic Container Registry

- Private registry
- Public registry
- Repositories
- Getting started
- Documentation
- Public gallery

Amazon ECR > Repositories

Private Public

Private repositories

Find repositories

Repository name	URI	Created at	Tag immutability	Scan frequency	Encryption type	Pull through cache
No repositories						
No repositories were found						

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Amazon ECR > Repositories > Create repository

Create repository

General settings

Visibility settings info
Choose the visibility setting for the repository.

Private
Access is managed by IAM and repository policy permissions.

Public
Publicly visible and accessible for image pulls.

Repository name
Provide a concise name. A developer should be able to identify the repository contents by the name.

456464162939.dkr.ecr.us-east-1.amazonaws.com/

21 out of 256 characters maximum (2 minimum). The name must start with a letter and can only contain lowercase letters, numbers, hyphens, underscores, and forward slashes.

Tag immutability info
Enable tag immutability to prevent image tags from being overwritten by subsequent image pushes using the same tag. Disable tag immutability to allow image tags to be overwritten.

Disabled

Once a repository is created, the visibility setting of the repository can't be changed.

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Amazon Elastic Container Registry

- Private registry
- Public registry
- Repositories
- Summary
- Images
- Permissions
- Lifecycle Policy
- Tags
- Getting started
- Documentation
- Public gallery

Amazon ECR > Repositories > customer-segmentation

customer-segmentation

View push commands Edit

Images (0)

Find images

Image tag	Pushed at	Size (MB)	Image URI	Digest	Scan status	Vulnerabilities
No images						
No images to display						

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Push commands for customer-segmentation

macOS / Linux | Windows

Make sure that you have the latest version of the AWS CLI and Docker installed. For more information, see [Getting Started with Amazon ECR](#).

Use the following steps to authenticate and push an image to your repository. For additional registry authentication methods, including the Amazon ECR credential helper, see [Registry Authentication](#).

- Retrieve an authentication token and authenticate your Docker client to your registry. Use the AWS CLI:


```
aws ecr get-login-password --region us-east-1 | docker login --username AWS --password-stdin [REDACTED].dkr.ecr.us-east-1.amazonaws.com
```

Note: If you receive an error using the AWS CLI, make sure that you have the latest version of the AWS CLI and Docker installed.
- Build your Docker image using the following command. For information on building a Docker file from scratch see the instructions [here](#). You can skip this step if your image is already built:


```
docker build -t customer-segmentation .
```
- After the build completes, tag your image so you can push the image to this repository:


```
docker tag customer-segmentation:latest [REDACTED].dkr.ecr.us-east-1.amazonaws.com/customer-segmentation:latest
```
- Run the following command to push this image to your newly created AWS repository:


```
docker push [REDACTED].dkr.ecr.us-east-1.amazonaws.com/customer-segmentation:latest
```

Amazon Elastic Container Registry

customer-segmentation

View push commands | Edit

Images (1)

Find images

Image tag	Pushed at	Size (MB)	Image URI	Digest	Scan status	Vulnerabilities
latest	April 17, 2022, 15:55:02 (UTC-04)	647.80	Copy URI	sha256-[REDACTED]	-	-

Amazon SageMaker

Models

Create endpoint | Create endpoint configuration | Actions | Create model

Search models

Name	ARN	Creation time
There are currently no resources.		

© 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Create model

To deploy a model to Amazon SageMaker, first create the model by providing the location of the model artifacts and inference code. See [Deploying a Model on Amazon SageMaker Hosting Services](#) or [Learn more about the API](#).

Model settings

Model name

 Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.

IAM role
 Amazon SageMaker requires permissions to call other services on your behalf. Choose a role or let us create a role that has the [AmazonSageMakerFullAccess](#) IAM policy attached.

Create a new role
 Enter a custom IAM role ARN

Use existing role

Container input options

Provide model artifacts and inference image location
 Use this for models trained using built-in algorithms, BYO algorithms, or models trained outside Amazon SageMaker.

Use a model package resource
 Use this for model packages that contain inference images and artifacts from AWS Marketplace subscribed algorithms.

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Amazon SageMaker Models Create model

Create model

To deploy a model to Amazon SageMaker, first create the model by providing the location of the model artifacts and inference code. See [Deploying a Model on Amazon SageMaker Hosting Services](#) or [Learn more about the API](#).

Model settings

Model name

 Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.

IAM role
 Amazon SageMaker requires permissions to call other services on your behalf. Choose a role or let us create a role that has the [AmazonSageMakerFullAccess](#) IAM policy attached.

Create a new role
 Please provide a valid ARN.

Container definition 1

Container input options

Provide model artifacts and inference image location
 Use this for models trained using built-in algorithms, BYO algorithms, or models trained outside Amazon SageMaker.

Use a model package resource
 Use this for model packages that contain inference images and artifacts from AWS Marketplace subscribed algorithms.

Create an IAM role

Passing an IAM role gives Amazon SageMaker permission to perform actions in other AWS services on your behalf. Creating a role here will grant permissions described by the [AmazonSageMakerFullAccess](#) IAM policy to the role you create.

The IAM role you create will provide access to:

- S3 buckets you specify - optional**
 - Any S3 bucket**
 Allow users that have access to your notebook instance access to any bucket and its contents in your account.
 - Specific S3 buckets**

 Comma delimited. ARNs, "*" and "/" are not supported.
 - None
- Any S3 bucket with "sagemaker" in the name
- Any S3 object with "sagemaker" in the name
- Any S3 object with the tag "sagemaker" and value "true" [See Object tagging](#)
- S3 bucket with a Bucket Policy allowing access to SageMaker [See S3 bucket policies](#)

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Create model

To deploy a model to Amazon SageMaker, first create the model by providing the location of the model artifacts and inference code. See [Deploying a Model on Amazon SageMaker Hosting Services](#) or [Learn more about the API](#).

Model settings

Model name

 Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.

IAM role
 Amazon SageMaker requires permissions to call other services on your behalf. Choose a role or let us create a role that has the [AmazonSageMakerFullAccess](#) IAM policy attached.

Success! You created an IAM role.
[AmazonSageMaker-ExecutionRole-20220417T182620](#)

Container definition 1

Container input options

Provide model artifacts and inference image location
 Use this for models trained using built-in algorithms, BYO algorithms, or models trained outside Amazon SageMaker.

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Container definition 1

Container input options

- Provide model artifacts and inference image location
Use this for models trained using built-in algorithms, BYO algorithms, or models trained outside Amazon SageMaker.
- Use a model package resource
Use this for model packages that contain inference images and artifacts from AWS Marketplace subscribed algorithms.
- Use a model package subscription from AWS Marketplace
Use this for model packages published by vendors from AWS Marketplace.

Provide model artifacts and inference image options

- Use a single model
Use this to host a single model in this container.
- Use multiple models
Use this to host multiple models in this container.

Location of inference code image
Type the registry path where the inference code image is stored in Amazon ECR.
[redacted].us-east-1.amazonaws.com/customer-segmentation:latest

Location of model artifacts - optional
Type the URL where model artifacts are stored in S3.
s3://bucket/path-to-your-data/
The path must point to a single gzip compressed tar archive (.tar.gz suffix).

Container host name - optional

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Network

Enable network isolation
Containers that run with network isolation can't make any outbound network calls. This is required for AWS ML Marketplace products.

VPC - optional
For better security, we recommend that you use a private VPC.
Default vpc: [redacted]

Subnet(s)
Choose a subnet in an availability zone supported by Amazon SageMaker.
subnet-[redacted] 2.31.48.0/20 | us-east-1e X

Security group(s)
Configure your VPC to ensure that your models have access to protected resources. [Allow Models to Access Resources in Your Private VPC.](#)
sg-[redacted] (default) X

Tags - optional

Key	Value	Remove
<input type="text"/>	<input type="text"/>	<input type="button" value="Remove"/>

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

aws Services Search for services, features, blogs, docs, and more [Alt+S] N. Virginia jayanth

Amazon SageMaker Endpoint configuration

Apply to endpoint Clone Actions **Create endpoint configuration**

Search endpoint configuration < 1 >

Name	ARN	Creation time
There are currently no resources.		

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Amazon SageMaker > Endpoint configuration > Create endpoint configuration

Create endpoint configuration

To deploy models to Amazon SageMaker, first create an endpoint configuration. In the configuration, specify which models to deploy, and the relative traffic weighting and hardware requirements for each. See [Deploying a Model on Amazon SageMaker](#) [Hosting Services](#) [Learn more about the API](#)

Endpoint configuration

Endpoint configuration name

Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.

Type of endpoint
 Provisioned
 Serverless (In Preview)

Encryption key - optional
Encrypt your data. Choose an existing KMS key or enter a key's ARN.

▼ Async Invocation Config - optional
Max concurrent invocations per instance - optional

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

▼ Data capture - optional
 Enable data capture
By enabling this feature, Amazon SageMaker will capture and store information from your training jobs. Learn more [Learn more](#)

Production variants

Model name Training job

Add model

▼ Tags - optional

Key Value Remove

Add tag

Add model

Search resources < 1 >

Name	Creation time
customer-segmentation	Apr 17, 2022 23:06 UTC

Cancel Save

Cancel Create endpoint configuration

Amazon SageMaker > Endpoints

Endpoints

Update endpoint Actions Create endpoint

Search endpoints < 1 >

Name	ARN	Creation time	Status	Last updated
There are currently no resources.				

Feedback © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Endpoint

Endpoint name
Your application uses this name to access this endpoint.

Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.

Attach endpoint configuration

Use an existing endpoint configuration
Use an existing endpoint configuration or clone an endpoint configuration.

Create a new endpoint configuration
Add models and configure the instance and initial weight for each model.

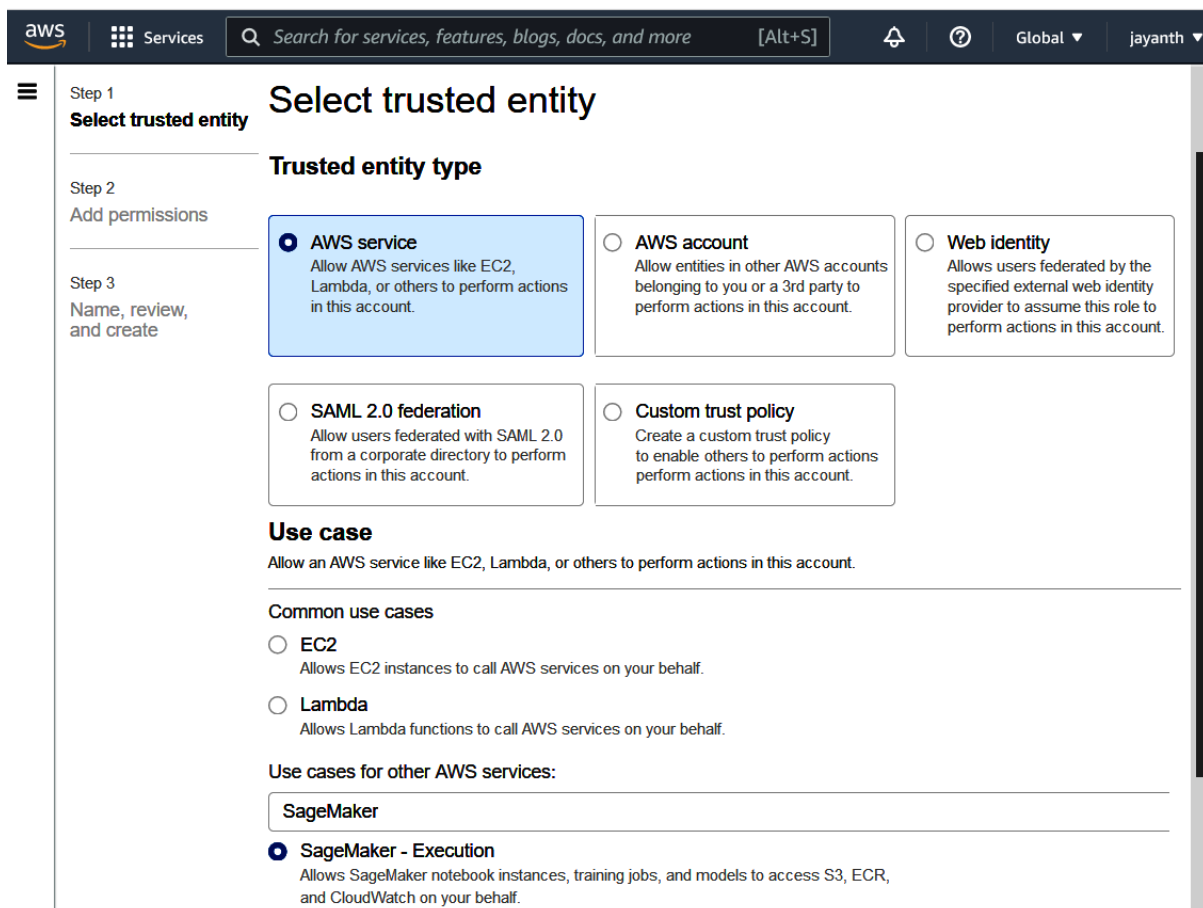
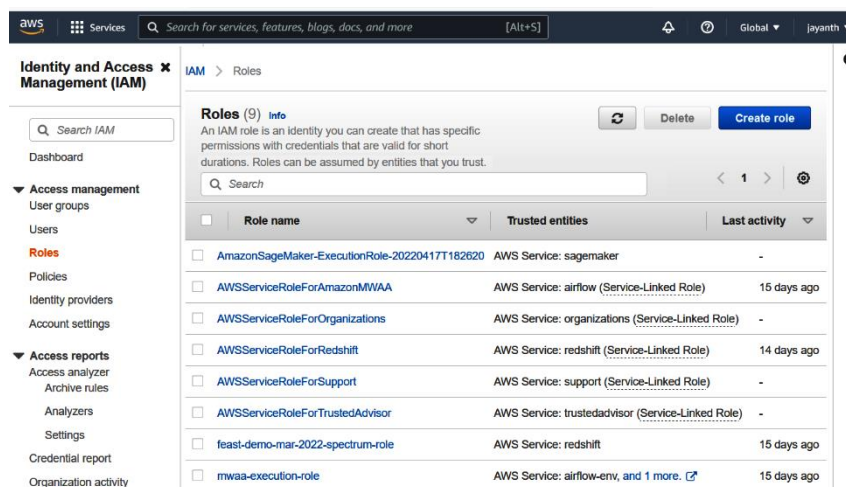
Endpoint configuration

< 1 >

Name	ARN	Creation time
<input checked="" type="radio"/> customer-segmenation-config	arn:aws:sagemaker:us-east-1:456464162939:endpoint-config/customer-segmenation-config	Apr 18, 2022 01:54 UTC

Select endpoint configuration

Chapter 7: Feast Alternatives and ML Best Practices





Step 1
Select trusted entity

Name, review, and create

Step 2
Add permissions

Role details

Step 3
Name, review, and create

Role name
Enter a meaningful name to identify this role.

Maximum 128 characters. Use alphanumeric and '+=, @-_' characters.

Description
Add a short explanation for this policy.

Allows SageMaker notebook instances, training jobs, and models to access S3, ECR, and CloudWatch on your behalf.

Maximum 1000 characters. Use alphanumeric and '+=, @-_' characters.

Step 1: Select trusted entities Edit

```

1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Sid": "",
6       "Effect": "Allow"
    }
  ]
}
    
```

	customerid	recency	frequency	monetaryvalue	r	f	m	rfmscore	revenue6m	ltvcluster	segmenthighvalue	segmentlowValue	segmentmidvalue	event_timestamp
0	12747.0	7	35	1082.09	3	2	3	8	1666.11	1	1	0	0	1.651455e+09
1	12748.0	1	582	4336.73	3	3	3	9	18679.01	2	1	0	0	1.651455e+09
2	12749.0	8	54	782.10	3	3	3	9	2323.04	1	1	0	0	1.651455e+09
4	12823.0	63	1	459.00	0	0	2	2	765.00	1	0	1	0	1.651455e+09
7	12836.0	28	62	814.71	1	3	3	7	951.46	1	1	0	0	1.651455e+09

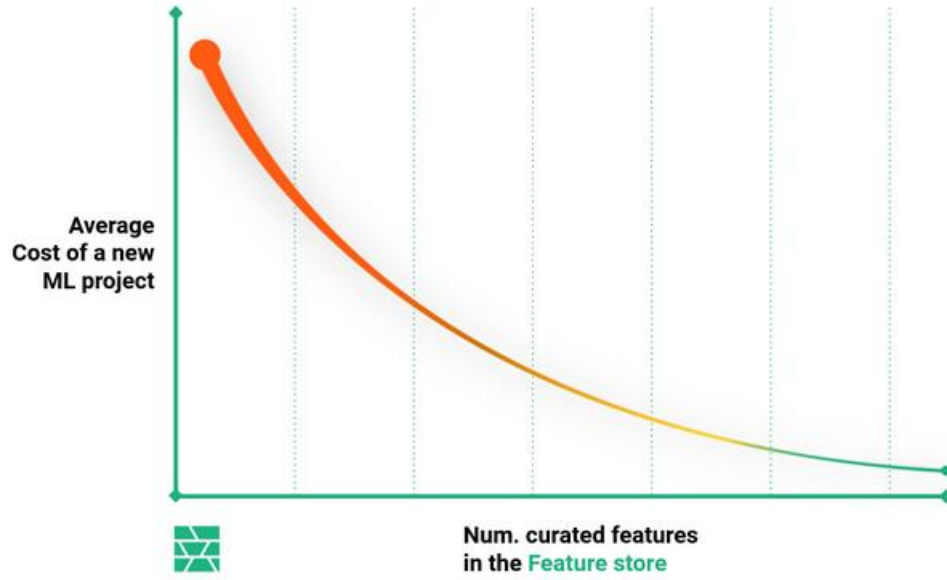
```

[FeatureDefinition(feature_name='customerid', feature_type=<FeatureTypeEnum.FRACTIONAL: 'Fractional'>),
FeatureDefinition(feature_name='recency', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='frequency', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='monetaryvalue', feature_type=<FeatureTypeEnum.FRACTIONAL: 'Fractional'>),
FeatureDefinition(feature_name='r', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='f', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='m', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='rfmscore', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='revenue6m', feature_type=<FeatureTypeEnum.FRACTIONAL: 'Fractional'>),
FeatureDefinition(feature_name='ltvcluster', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='segmenthighvalue', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='segmentlowValue', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='segmentmidvalue', feature_type=<FeatureTypeEnum.INTEGRAL: 'Integral'>),
FeatureDefinition(feature_name='event_timestamp', feature_type=<FeatureTypeEnum.FRACTIONAL: 'Fractional'>)]
    
```

```

{'FeatureGroupArn': 'arn:aws:sagemaker:us-east-1:XXXXXXXXXX:feature-group/customer-rfm-features',
'ResponseMetadata': {'HTTPHeaders': {'content-length': '98',
'content-type': 'application/x-amz-json-1.1',
'date': 'Mon, 02 May 2022 01:30:15 GMT',
'x-amzn-requestid': '4fb00252-deXXXXXXXXXX'},
'HTTPStatusCode': 200,
'RequestId': '4fb00252-deXXXXXXXXXX',
'RetryAttempts': 0}}
    
```

customerid	recency	frequency	monetaryvalue	r	f	m	rfmscore	revenue6m	ltvcluster	segmenthighvalue	segmentlowvalue	segmentmidvalue	event_timestamp	write_time
12747.0	7	35	1082.09	3	2	3	8	1666.11	1	1	0	0	1.651455e+09	2022-05-02 01:36:31.398
12823.0	63	1	459.00	0	0	2	2	765.00	1	0	1	0	1.651455e+09	2022-05-02 01:36:31.406
12836.0	28	62	814.71	1	3	3	7	951.46	1	1	0	0	1.651455e+09	2022-05-02 01:36:31.391
12748.0	1	582	4336.73	3	3	3	9	18679.01	2	1	0	0	1.651455e+09	2022-05-02 01:36:31.400
12749.0	8	54	782.10	3	3	3	9	2323.04	1	1	0	0	1.651455e+09	2022-05-02 01:36:31.397



Chapter 8: Use Case – Customer Churn Prediction

	customer_id	gender	SeniorCitizen	Partner	Dependents	tenure	PhoneService	MultipleLines	InternetService	OnlineSecurity	OnlineBackup
0	7590-VHVEG	Female	0	Yes	No	1	No	No phone service	DSL	No	Yes
1	5575-GNVDE	Male	0	No	No	34	Yes	No	DSL	Yes	No
2	3668-QPYBK	Male	0	No	No	2	Yes	No	DSL	Yes	Yes
3	7795-CFOCW	Male	0	No	No	45	No	No phone service	DSL	Yes	No
4	9237-HQITU	Female	0	No	No	2	Yes	No	Fiber optic	No	No

	DeviceProtection	TechSupport	StreamingTV	StreamingMovies	Contract	PaperlessBilling	PaymentMethod	MonthlyCharges	TotalCharges	Churn
0	No	No	No	No	Month-to-month	Yes	Electronic check	29.85	29.85	No
1	Yes	No	No	No	One year	No	Mailed check	56.95	1889.5	No
2	No	No	No	No	Month-to-month	Yes	Mailed check	53.85	108.15	Yes
3	Yes	Yes	No	No	One year	No	Bank transfer (automatic)	42.30	1840.75	No
4	No	No	No	No	Month-to-month	Yes	Electronic check	70.70	151.65	Yes

	customerID	gender	SeniorCitizen	Partner	Dependents	tenure	PhoneService	MultipleLines	InternetService	OnlineSecurity	OnlineBackup	DeviceProtection
4850	3950-VPYJB	Male	0	Yes	Yes	57	Yes	No	DSL	Yes	Yes	No
2622	4213-HKBJO	Female	0	No	No	33	Yes	Yes	No	No	No	No
4773	9814-AOUDH	Male	0	No	No	53	Yes	No	No	No	No	No
490	6778-YSNIH	Female	0	No	No	2	Yes	No	DSL	No	Yes	No
1085	4546-FOKWR	Female	0	No	No	16	Yes	Yes	Fiber optic	No	No	No

	TechSupport	StreamingTV	StreamingMovies	Contract	PaperlessBilling	PaymentMethod	MonthlyCharges	TotalCharges	Churn	tenure_group
4850	Yes	No	No	One year	No	Mailed check	59.60	3509.40	No	48-end
2622	No	No	No	Two year	No	Credit card (automatic)	25.70	826.10	No	24-48
4773	No	No	No	Two year	No	Credit card (automatic)	19.50	1050.50	No	48-end
490	No	Yes	No	Month-to-month	Yes	Electronic check	59.00	114.15	No	0-24
1085	No	No	No	Month-to-month	Yes	Credit card (automatic)	74.75	1129.35	No	0-24

```

customerID      7032
gender          2
SeniorCitizen   2
Partner         2
Dependents      2
tenure          72
PhoneService    2
MultipleLines   3
InternetService 3
OnlineSecurity  2
OnlineBackup    2
DeviceProtection 2
TechSupport     2
StreamingTV     2
StreamingMovies 2
Contract        3
PaperlessBilling 2
PaymentMethod   4
MonthlyCharges  1584
TotalCharges    6530
Churn           2
tenure_group    3
dtype: int64
    
```

	customerid	gender	seniorcitizen	partner	dependents	tenure	phoneservice	onlinesecurity	onlinebackup	deviceprotection	techsupport	streamingtv
0	7590-VHVEG	0	0	1	0	-1.280248	0	0	1	0	0	0
1	5575-GNVDE	1	0	0	0	0.064303	1	1	0	1	0	0
2	3668-QPYBK	1	0	0	0	-1.239504	1	1	1	0	0	0
3	7795-CFOCW	1	0	0	0	0.512486	0	1	0	1	1	0
4	9237-HQITU	0	0	0	0	-1.239504	1	0	0	0	0	0

	streamingmovies	paperlessbilling	monthlycharges	totalcharges	churn	multiplelines_no	multiplelines_no_phone_service	multiplelines_yes	internetservice_dsl	
0	0	1	-1.161694	-0.994194	0	0		1	0	1
1	0	0	-0.260878	-0.173740	0	1		0	0	1
2	0	1	-0.363923	-0.959649	1	1		0	0	1
3	0	0	-0.747850	-0.195248	0	0		1	0	1
4	0	1	0.196178	-0.940457	1	1		0	0	0

	internetservice_fiber_optic	internetservice_no	contract_month_to_month	contract_one_year	contract_two_year	paymentmethod_bank_transfer_automatic
0		0		1	0	0
1		0		0	1	0
2		0		1	0	0
3		0		0	1	1
4		1		1	0	0

	paymentmethod_credit_card_automatic	paymentmethod_electronic_check	paymentmethod_mailed_check	tenure_group_0_24	tenure_group_24_48	tenure_group_48_end
0		0	1	0	1	0
1		0	0	1	0	1
2		0	0	1	1	0
3		0	0	0	0	1
4		0	1	0	1	0

	customerid	predicted_churn	churn
0	7590-VHVEG	0	0
1	5575-GNVDE	0	0
2	3668-QPYBK	0	1
3	7795-CFOCW	0	0
4	9237-HQITU	1	1

Classification report :

	precision	recall	f1-score	support
0	0.82	0.90	0.86	1268
1	0.65	0.48	0.56	490
accuracy			0.78	1758
macro avg	0.74	0.69	0.71	1758
weighted avg	0.77	0.78	0.77	1758

Accuracy Score : 0.7849829351535836

Area under curve : 0.6925465138736883