

Unable to locate subtitle

AWS Data Exchange User Guide



Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AWS Data Exchange User Guide: ***Unable to locate subtitle***

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

What is AWS Data Exchange?	1
What is a data grant in AWS Data Exchange	1
What is an AWS Marketplace data product?	2
Supported data sets	3
Accessing AWS Data Exchange	3
Data receivers	3
Data senders and providers	3
Supported Regions	4
Related services	5
Setting up	6
Sign up for an AWS account	6
Sign up for an AWS account	6
Create a user with administrative access	6
Create a user	8
Using Open Data on AWS data sets	10
Getting started using the AWS Data Exchange console	11
Step 1: Find an Open Data on AWS data set	11
Step 2: Use an Open Data on AWS data set	. 12
Getting started without an AWS account	12
Step 1: Find an Open Data on AWS data set	12
Step 2: Use an Open Data on AWS data set	. 13
Data in AWS Data Exchange	14
Assets	14
Asset structure	. 14
Asset types	15
Revisions	17
Revision structure	18
Data sets	19
Owned data sets	20
Entitled data sets	20
Data set types	20
Amazon S3 data access data set	21
AWS Lake Formation data set (Preview)	22
AWS Regions and data sets	22

Data set structure	22
Data set best practices	23
Tags	23
Creating data grants	25
Programmatic access	26
Containing file-based data	26
Step 1: Create assets	26
Step 2: Create a data set	27
Step 3: Create a revision	27
Step 4: Import assets to a revision	28
Step 5: Create a new data grant	29
Containing APIs	30
Prerequisites	31
Step 1: Update the API resource policy	33
Step 2: Create an API data set	34
Step 3: Create a revision	35
Step 4: Add API assets to a revision	37
Step 5: Create a new data grant containing APIs	42
Containing Amazon Redshift data sets	43
Step 1: Create an Amazon Redshift datashare asset	44
Step 2: Create an Amazon Redshift data set	44
Step 3: Create a revision	45
Step 4: Add Amazon Redshift datashare assets to a revision	45
Step 5: Create a new data grant	46
Containing Amazon S3 data access	48
Step 1: Create an Amazon S3 data set	57
Step 2: Configure Amazon S3 data access	57
Step 3: Review and finalize the data set	59
Step 4: Create a new data grant	59
Containing AWS Lake Formation data permission data sets (Preview)	60
Step 1: Create an AWS Lake Formation data set (Preview)	61
Step 2: Create an AWS Lake Formation data permission (Preview)	62
Step 3: Review and finalize	62
Step 4: Create a revision	63
Step 5:Create a new data grant containing AWS Lake Formation data sets (Preview)	63

Considerations when creating data grants containing an AWS Lake Formation data	
permission data set (Preview)	64
Accepting data grants and accessing data on AWS Data Exchange	66
Related topics	67
Access a data set after accepting a data grant	67
Containing file-based data	67
Containing APIs	68
Containing Amazon Redshift data sets	70
Containing Amazon S3 data access	71
Containing AWS Lake Formation data sets (Preview)	74
Sharing a data grant license in an organization	75
Prerequisites for license sharing	75
Viewing your licenses	76
Sharing your licenses	77
Subscribing to data products	78
Related topics	79
Product subscriptions	79
Data sets and revisions	82
Data dictionaries and samples	83
Getting started as a subscriber	83
Step 1: Set up AWS Data Exchange	83
Step 2: Browse the catalog	84
Step 3: (Optional) Request a recommendation for a data product	84
Step 4: (Optional) Evaluate products containing data dictionaries and samples	85
Step 5: Subscribe to and access a product	88
Subscribing to a product	88
Containing file-based data	89
Containing APIs	93
Containing Amazon Redshift data sets	97
Containing Amazon S3 data access	100
Containing AWS Lake Formation data sets (Preview)	105
Viewing and downloading a data dictionary	107
Subscription verification for subscribers	107
Completing a subscription request	108
Reviewing your pending subscription requests	108
Email notifications	109

	Sharing license subscriptions in an organization	109
	Prerequisites for license sharing	110
	Step 1: View your licenses	110
	Step 2: Share your licenses	111
	BYOS offers	111
	Private products and offers	113
	Managing subscriptions	114
	Viewing your subscriptions	114
	Turning subscription auto-renewal on or off	115
	Unsubscribing from a product	115
	Products for learning about interacting with AWS Data Exchange	116
	AWS Data Exchange Heartbeat	116
	AWS Data Exchange for APIs	119
	Worldwide Event Attendance	122
	AWS Data Exchange for AWS Lake Formation (Preview)	125
	AWS Data Exchange for Amazon S3	130
	AWS Data Exchange Provider-Generated Notifications	133
Pr	oviding data products on AWS Marketplace	137
	Extended Provider Program (EPP)	138
	Programmatic access	139
	Related topics	139
	Getting started as a provider	140
	Step 1: Confirm your eligibility	140
	Step 2: Register to be a provider	142
	Step 3: Confirm eligibility of your data	143
	Publishing guidelines	143
	Publishing a new product	146
	Containing file-based data	146
	Containing APIs	153
	Containing Amazon Redshift data sets	170
	Containing Amazon S3 data access	177
	Containing AWS Lake Formation data permission data sets (Preview)	193
	Product best practices	199
	Product visibility	200
	Sensitive categories of information	200
	Product details	202

Revision access rules	205
Data dictionaries	206
Samples	207
Product description templates	207
Generic template	208
Financial services template	210
Healthcare and life sciences template	213
Marketing and advertising template	215
Media and entertainment template	217
Public sector template	220
Retail and location template	222
Creating offers	224
Offer pricing	225
US sales and use tax	225
Data Subscription Agreement	225
Refund policy	226
Subscription verification	226
Offer auto-renewal	226
Private offers	227
BYOS offers	228
Viewing subscriptions	230
Updating products	231
Updating product and offer details	232
Updating a data dictionary	234
Updating a sample	235
Updating custom metadata	235
Publishing a new data set revision	236
Unpublish a product	238
Removing a revision	239
Revoking revisions	239
Subscription verification for providers	244
Email notifications	245
Viewing subscription verification requests	246
Approve or decline requests	246
Provider-generated notifications	249
Provider financials on AWS Marketplace	249

	Payments	250
	US sales and use tax	250
	AWS Marketplace seller reports	250
	Subscriber refund requests	251
Jo	bs in AWS Data Exchange	252
	Job properties	252
	AWS Regions and jobs	253
	Importing assets	254
	From an S3 bucket	254
	From a signed URL	256
	From an Amazon API Gateway API	257
	From a datashare for Amazon Redshift	259
	From an AWS Lake Formation (Preview)	260
	Exporting assets	262
	To an S3 bucket	262
	To a signed URL	266
	Exporting revisions	268
	Key patterns when exporting revisions	269
	Using AWS SDKs	270
	Using the console (Subscriber)	271
	Using the console (Provider)	272
	Automatically exporting revisions (Subscriber)	272
Q	uotas	278
	Service quotas	278
	Service endpoints	278
	Export and import job guidelines	278
	Constraints for resource fields	280
Lc	ogging and monitoring	281
	Monitoring	281
	Amazon EventBridge events for AWS Data Exchange	282
	Events for adding file-based data sets	285
	Events for adding Amazon S3 data access data sets	
	Events for adding AWS Lake Formation data permission data sets	
	Events for adding Amazon Redshift datashare data sets	288
	Events for adding Amazon API Gateway API data sets	289
	Events for adding revisions	290

Events for adding Amazon S3 data access data set revisions	291
Events for adding AWS Lake Formation data permission data set revisions (Preview)	292
Events for adding Amazon Redshift datashare data set revisions	293
Events for adding Amazon API Gateway API data set revisions	294
Events for revoking revisions	297
Events for an action performed on an Amazon Redshift resource	298
Events for losing access to an Amazon Redshift datashare	299
Events for an auto-export job completed	300
Events for an auto-export job failed	301
Events for a provider-generated notification of a data update	302
Events for a provider-generated notification of a schema change	303
Events for a provider-generated notification of a data delay	305
Events for a provider-generated notification of a data deprecation	306
Events for accepting a data grant	307
Events for extending data grants	308
Events for revoking a data grant	308
AWS User Notifications for AWS Data Exchange events	309
Logging AWS Data Exchange API calls with AWS CloudTrail	311
AWS Data Exchange information in CloudTrail	313
Understanding AWS Data Exchange log file entries	314
Upcoming changes in AWS Data Exchange CloudTrail logging	315
Security	319
Data protection	319
Encryption at rest	320
Encryption in transit	321
Restrict access to content	321
Key management for Amazon S3 data access	321
Creating AWS KMS grants	321
Encryption context and grant constraints	322
Monitoring your AWS KMS keys in AWS Data Exchange	322
Identity and access management	326
Authentication	327
Access control	328
API permissions reference	335
AWS managed policies	343
Using service-linked roles	350

Creating a service-linked role for AWS Data Exchange	351
Editing a service-linked role for AWS Data Exchange	352
Deleting a service-linked role for AWS Data Exchange	352
Supported Regions for AWS Data Exchange service-linked roles	352
Service-linked role for license management	353
Service-linked role for AWS Organization discovery	354
Compliance validation	355
PCI DSS compliance	356
Resilience	356
Infrastructure security	357
VPC endpoints (AWS PrivateLink)	357
Considerations for AWS Data Exchange VPC endpoints	358
Creating an interface VPC endpoint for AWS Data Exchange	358
Creating a VPC endpoint policy for AWS Data Exchange	358
AWS Marketplace Catalog API	361
AddDataSets	361
Tutorial: Adding new data sets to a published data product	362
AddDataSets exceptions	365
Document history	367

What is AWS Data Exchange?

AWS Data Exchange is a service that helps AWS customers easily share and manage data entitlements from other organizations at scale.

As a data receiver, you can track and manage all of your data grants and AWS Marketplace data subscriptions in one place. When you have access to an AWS Data Exchange data set, you can use compatible AWS or partner analytics and machine learning to extract insights from it. For information about purchasing data products from AWS Marketplace, see Subscribing to AWS Data Exchange data products on AWS Data Exchange.

For data senders, AWS Data Exchange eliminates the need to build and maintain any data delivery and entitlement infrastructure. Anyone with an AWS account can create and send data grants to data receivers. To sell your data as a product in AWS Marketplace, make sure that you follow the guidelines to determine eligibility. For more information, see Providing AWS Data Exchange data products on AWS Marketplace.

In addition, anyone, with or without an AWS account, can find and use publicly available data sets that are part of the <u>Open Data on AWS</u> program. For more information, see <u>Using Open Data on AWS</u> data sets with AWS Data Exchange.

Topics

- What is a data grant in AWS Data Exchange
- What is an AWS Marketplace data product?
- Supported data sets
- Accessing AWS Data Exchange
- Supported Regions
- Related services

What is a data grant in AWS Data Exchange

A data grant is the unit of exchange in AWS Data Exchange that is created by a data sender in order to grant a data receiver access to a data set. When a data sender creates a data grant, a grant request is sent to the data receiver's AWS account. A data receiver accepts the data grant to gain access to the underlying data.

A grant has the following parts:

- Data set A data set in AWS Data Exchange is a resource curated by the sender. It contains
 the data assets a receiver will gain access to after accepting a data grant. AWS Data Exchange
 supports five types of data sets: Files, API, Amazon Redshift, Amazon S3, and AWS Lake
 Formation (Preview).
- **Data grant details** This information includes a name and description of the data grant that will be visible to data receivers.
- **Recipient access details** This information includes the receiver's AWS account ID and specifies how long the receiver should have access to the data.

What is an AWS Marketplace data product?

A product is the unit of exchange in AWS Marketplace that is published by a provider and made available for use to subscribers. A data product is a product that includes AWS Data Exchange data sets. When a data provider publishes a data product, that product is listed in the AWS Marketplace product catalog after being reviewed by AWS against our guidelines and terms and conditions. Each product published is uniquely identified by its product ID.

A data product has the following parts:

- **Product details** This information includes name, descriptions (both short and long), data samples, a logo image, and support contact information. Providers complete the product details.
 - For more information as a subscriber, see Product subscriptions in AWS Data Exchange.
 - For more information as a provider, see <u>Product best practices in AWS Data Exchange</u>.
- Product offers Offers define the terms that subscribers are agreeing to when they subscribe
 to a product. To make a product available in the public AWS Marketplace Catalog, providers
 must define a public offer. This offer includes prices and durations, data subscription agreement,
 refund policy, and the option to create custom offers.
 - For more information as a subscriber, see <u>Accepting private products and offers in AWS Data</u> Exchange and Accepting Bring Your Own Subscription (BYOS) offers in AWS Data Exchange
 - For more information as a provider, see <u>Creating an offer for AWS Data Exchange products</u>.
- Data sets A product can contain one or more data sets. A data set in AWS Data Exchange is a resource curated by the data provider and contains the data assets a receiver will gain access to after accepting a data grant. AWS Data Exchange supports five types of data sets: Files, API, Amazon Redshift, Amazon S3, and AWS Lake Formation (Preview).

- For more information as a subscriber, see Data sets and revisions.
- For more information as a provider, see Data in AWS Data Exchange.

Supported data sets

AWS Marketplace takes a responsible approach to facilitating data transactions by promoting transparency through use of the service. AWS Marketplace reviews permitted data types, restricting products that are not permitted. Providers are limited to distributing data sets that meet the legal eligibility requirements set forth in the Terms and Conditions for AWS Marketplace Sellers.

For more information about permitted data types, see Publishing guidelines for AWS Data Exchange.

Important

As an AWS customer, you are encouraged to conduct your own additional due-diligence to ensure compliance with any applicable data privacy laws. If you suspect that a product or other resources on AWS Data Exchange are being used for abusive or illegal purposes, report it using the Report Amazon AWS abuse form.

Accessing AWS Data Exchange

Data receivers

As a data receiver, you can view all of your current, pending, and expired data grants from the AWS Data Exchange console.

You can also discover and subscribe to new third-party data sets available through AWS Data Exchange from the AWS Marketplace catalog.

Data senders and providers

As a data sender or provider, you can access AWS Data Exchange through the following options:

- Directly through the AWS Data Exchange console (Publish data)
- Data providers with data products available in AWS Marketplace can access programmatically using the following APIs:

Supported data sets

- AWS Data Exchange API Use the API operations to create, view, update, and delete data sets and revisions. You can also use these API operations to import and export assets to and from those revisions. For more information, see the AWS Data Exchange API Reference.
- AWS Marketplace Catalog API Use the API operations to view and update data products published to AWS Marketplace. For more information, see the AWS Marketplace Catalog API Reference.

Supported Regions

AWS Data Exchange data grants, subscriptions, data sets, revisions, and assets are Region resources that can be managed programmatically or through the AWS Data Exchange console in supported Regions. Data products published to AWS Marketplace are available in a single product catalog. Subscribers can see the same catalog regardless of which supported AWS Region they are using.

The following regions are supported:

- US East (N. Virginia)
- US East (Ohio)
- US West (Oregon)
- US West (N. California)
- EU (Ireland)
- EU (Frankfurt)
- EU (London)
- Asia Pacific (Singapore)
- Asia Pacific (Tokyo)
- Asia Pacific (Sydney)
- Asia Pacific (Seoul)



Note

For more information about supported regions, see the Global Infrastructure Region Table.

Supported Regions

Related services

The following services are related to AWS Data Exchange:

- Amazon S3 AWS Data Exchange allows providers to import and store data files from their Amazon S3 buckets. Data recipients can export these files to Amazon S3 programmatically. AWS Data Exchange also enables recipients to directly access and use providers' Amazon S3 buckets.
 For more information, see What is Amazon S3? in the Amazon Simple Storage Service User Guide.
- Amazon API Gateway Another supported asset type for data sets is APIs. Data recipients can
 call the API programmatically, call the API from the AWS Data Exchange console, or download
 the OpenAPI specification file. For more information, see What is Amazon API Gateway? in the
 Amazon API Gateway Developer Guide.
- Amazon Redshift AWS Data Exchange supports Amazon Redshift data sets. Data recipients can
 get read-only access to query the data in Amazon Redshift without extracting, transforming, and
 loading data. For more information, see <u>Getting started with Amazon Redshift</u> in the *Amazon*Redshift Getting Started Guide and <u>Amazon Redshift system overview</u> in the *Amazon Redshift*Database Developer Guide.
- AWS Marketplace AWS Data Exchange allows data sets to be published as products in AWS
 Marketplace. AWS Data Exchange data providers must be registered as AWS Marketplace sellers,
 and can use the AWS Marketplace Management Portal or the AWS Marketplace Catalog API. For
 information about becoming an AWS Marketplace subscriber, see What Is AWS Marketplace in the AWS Marketplace seller,
 see What Is AWS Marketplace? in the AWS Marketplace Seller Guide.
- AWS Lake Formation AWS Data Exchange supports AWS Lake Formation data permission data sets (Preview). Data recipients get access to data stored in a data provider's AWS Lake Formation data lake and can query, transform, and share access to this data from their own AWS Lake Formation data set. For more information, see AWS Lake Formation.

Related services 5

Setting up AWS Data Exchange

Before you can use any AWS service, including AWS Data Exchange, you must complete the following tasks:

Tasks

- Sign up for an AWS account
- Create a user

Sign up for an AWS account

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

- 1. Open https://portal.aws.amazon.com/billing/signup.
- 2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an AWS account root user is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform tasks that require root user access.

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to https://aws.amazon.com/ and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Sign up for an AWS account 6

Secure your AWS account root user

1. Sign in to the <u>AWS Management Console</u> as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see <u>Signing in as the root user</u> in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see <u>Enable a virtual MFA device for your AWS account root user (console)</u> in the *IAM User Guide*.

Create a user with administrative access

1. Enable IAM Identity Center.

For instructions, see <u>Enabling AWS IAM Identity Center</u> in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to a user.

For a tutorial about using the IAM Identity Center directory as your identity source, see <u>Configure user access with the default IAM Identity Center directory</u> in the AWS IAM Identity <u>Center User Guide</u>.

Sign in as the user with administrative access

 To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see <u>Signing in to the AWS access portal</u> in the *AWS Sign-In User Guide*.

Assign access to additional users

 In IAM Identity Center, create a permission set that follows the best practice of applying leastprivilege permissions.

For instructions, see Create a permission set in the AWS IAM Identity Center User Guide.

2. Assign users to a group, and then assign single sign-on access to the group.

For instructions, see Add groups in the AWS IAM Identity Center User Guide.

Create a user

To create an administrator user, choose one of the following options.

Choose one way to manage your administrator	То	Ву	You can also
In IAM Identity Center (Recommeded)	Use short-term credentials to access AWS. This aligns with the security best practices. For information about best practices, see Security best practices in IAM in the IAM User Guide.	Following the instructions in Getting started in the AWS IAM Identity Center User Guide.	Configure programmatic access by Configuring the AWS CLI to use AWS IAM Identity Center in the AWS Command Line Interface User Guide.
In IAM (Not recommer ed)	Use long-term credentials to access AWS.	Following the instructions in <u>Create an IAM user for emergency access</u> in the <i>IAM User Guide</i> .	Configure programmatic access by Manage access keys for IAM users in the IAM User Guide.

Create a user 8



Note

Only AWS accounts that are registered to provide data products on AWS Marketplace and AWS Data Exchange can create products.

Create a user

Using Open Data on AWS data sets with AWS Data **Exchange**

The Open Data on AWS program is a collection of over 300 free, publicly available data sets. You can use AWS Marketplace to find Open Data on AWS data sets, along with other no-cost and paid products, all in one place.

The Open Data on AWS data sets available in the catalog are part of the following affiliated programs:

- Open Data Sponsorship Program This AWS program covers the cost of storage for publicly available high-value cloud-optimized datasets.
- Amazon Sustainability Data Initiative (ASDI) This AWS program minimizes the cost and time required to acquire and analyze large sustainability datasets.

Anyone can search and find these free Open Data on AWS data sets, with or without an AWS account, no subscription required.

Anyone can analyze and build services on top of an Open Data data set by using compute and data analytics services including Amazon Elastic Compute Cloud (Amazon EC2), Amazon Athena, AWS Lambda, and Amazon EMR.



Note

Unless specifically stated in the applicable data set documentation, data sets that are available through the Registry of Open Data on AWS are not provided and maintained by AWS. Data sets are provided and maintained by a variety of third parties under a variety of licenses. To determine if a data set can be used for your application, check data set licenses and related documentation.

The following topics explain how to get started with Open Data on AWS data sets.

Topics

- Getting started with Open Data on AWS data sets using the AWS Data Exchange console
- Getting started with Open Data on AWS data sets without an AWS account

Getting started with Open Data on AWS data sets using the AWS Data Exchange console

The following topics describe how you can find and use an Open Data on AWS data set on AWS Data Exchange by using the AWS Data Exchange console. You must have an AWS account to complete this process.

The process has the following steps:

Steps

- Step 1: Find an Open Data on AWS data set
- Step 2: Use an Open Data on AWS data set

Step 1: Find an Open Data on AWS data set

To find an Open Data on AWS data set on AWS Data Exchange

- 1. Sign in to the AWS Management Console and open the AWS Data Exchange console.
- 2. In the left navigation pane, for **Subscribed with AWS Marketplace**, choose **Browse catalog**.
- 3. For **Refine results**, do one of the following:
 - a. For the Affiliated programs filter, select one or both of the following options:
 - AWS Open Data Sponsorship Program
 - Amazon Sustainability Data Initiative
 - b. (Optional) For the **Contract type** filter, select **Open Data Licenses** to see all publicly available affiliated and non-affiliated data sets.

For more information, see **Browse the catalog**.

4. (Optional) Enter a term or phrase in the search bar, and then choose **Search**.

A list of Open Data on AWS data sets that match the search terms appears.

5. Select an Open Data on AWS data set and view its details page.

The information on the details page includes a description, resources on AWS, usage examples, and links.

- a. (Optional) In the **Provided by** information, choose the link to the provider's information to view more information about the provider.
- b. (Optional) For **Labels**, choose a label to view a list of similar products.

Step 2: Use an Open Data on AWS data set

To use an Open Data on AWS data set

- 1. On the product details page, choose the **Resources on AWS** tab.
- 2. Copy the Amazon Resource Name (ARN) that is displayed under **Amazon Resource Name** (ARN).
- 3. For AWS CLI Access (No AWS account required), choose the AWS CLI link.
 - The AWS Command Line Interface (AWS CLI) documentation opens.
- 4. Read the documentation to learn how to use the AWS CLI to make calls to your AWS services from the command line.

For more information, see the AWS Command Line Interface User Guide.

Getting started with Open Data on AWS data sets without an AWS account

The following topics describe how you can find and use Open Data on AWS data sets without an AWS account. The process has the following steps:

Steps

- Step 1: Find an Open Data on AWS data set
- Step 2: Use an Open Data on AWS data set

Step 1: Find an Open Data on AWS data set

To find an Open Data on AWS data set

1. Go to AWS Data Exchange and then choose Browse 3,000+ third-party data sets.

The AWS Marketplace catalog appears, with the AWS Data Exchange delivery method and the AWS Open Data Sponsorships Program and Amazon Sustainability Data Initiative affiliated programs selected.

- 2. (Optional) For the **Contract type** filter, select **Open Data Licenses** to see all publicly available affiliated and non-affiliated data sets.
- 3. (Optional) Enter a term or phrase in the **Search** bar.

A list of Open Data on AWS data sets that match the search terms appears.

4. Select an Open Data data set and view its details page.

The information on the details page includes a description, resources on AWS, usage examples, and links.

- a. (Optional) In the **Provided by** information, choose the link to the provider's information to view more information about the provider.
- b. (Optional) For **Labels**, choose a label to view a list of similar products.

Step 2: Use an Open Data on AWS data set

To use an Open Data on AWS data set

- 1. On the product details page, choose the **Resources on AWS** tab.
- 2. Copy the Amazon Resource Name (ARN) that is displayed under **Amazon Resource Name** (ARN).
- 3. For AWS CLI Access (No AWS account required), choose the AWS CLI link.

The AWS Command Line Interface (AWS CLI) documentation opens.

Read the documentation to learn how to use the AWS CLI to make calls to your AWS services from the command line.

For more information, see the AWS Command Line Interface User Guide.

Data in AWS Data Exchange

Data is organized in AWS Data Exchange using three building blocks:

- Assets A piece of data
- Revisions A container for one or more assets
- Data sets A series of one or more revisions

These three building blocks form the foundation of the product that you manage using the AWS Data Exchange console or the AWS Data Exchange API.

To create, view, update, or delete data sets, you can use the AWS Data Exchange console, the AWS Command Line Interface (AWS CLI), your own REST client, or one of the AWS SDKs. For more information about programmatically managing AWS Data Exchange data sets, see the <u>AWS Data Exchange API Reference</u>.

Assets

Assets are the data in AWS Data Exchange.

The type of asset defines how the data is delivered to the receiver or subscriber through the data sets, data grants, or products that contain them.

An asset can be any of the following:

- A file stored on your local computer
- A file stored as an object in Amazon Simple Storage Service (Amazon S3)
- · A REST API created in Amazon API Gateway
- An Amazon Redshift data set
- An AWS Lake Formation data permission (Preview)
- An Amazon S3 data access data set

Asset structure

Assets have the following parameters:

Assets 14

- DataSetId The ID of the data set that contains this asset.
- RevisionId The ID of the revision that contains this asset.
- Id A unique ID generated when the asset is created.
- Arn A unique identifier for an AWS resource name.
- CreatedAt and UpdatedAt Date and timestamps for the creation and last update of the asset.
- AssetDetails Information about the asset.
- AssetType Either a snapshot of an Amazon S3 object, an Amazon API Gateway API, an Amazon Redshift data set, or an Amazon S3 data set.

Example asset resource

```
{
    "Name": "automation/cloudformation.yaml",
    "Arn": "arn:aws:dataexchange:us-east-1::data-sets/29EXAMPLE24b82c6858af3cEXAMPLEcf/
revisions/bbEXAMPLE74c02f4745c660EXAMPLE20/assets/baEXAMPLE660c9fe7267966EXAMPLEf5",
    "Id": "baEXAMPLE660c9fe7267966EXAMPLEf5",
    "CreatedAt": "2019-10-17T21:31:29.833Z",
    "UpdatedAt": "2019-10-17T21:31:29.833Z",
    "AssetType": "S3_SNAPSHOT",
    "RevisionId": "bbEXAMPLE74c02f4745c660EXAMPLE20",
    "DataSetId": "29EXAMPLE24b82c6858af3cEXAMPLEcf",
    "AssetDetails": {
        "S3SnapshotAsset": {
            "Size": 9423
        }
    }
}
```

Asset types

Types

- Files data set
- API assets
- Amazon Redshift datashare assets
- AWS Lake Formation data permission (Preview)
- Amazon S3 data access

Asset types 15

Files data set

Using Files, subscribers can access a copy of the data set as an entitled data set and export the assets.

A data set owner can both import and export Files using the AWS Data Exchange console, programmatically through the AWS CLI, their own REST application, or one of the AWS SDKs. For more information, about importing Amazon S3 assets. see Importing AWS Data Exchange assets from an S3 bucket. For more information about exporting assets, see Exporting AWS Data Exchange assets to an S3 bucket.

API assets

With API assets, data recipients or subscribers can view the API and download the API specification as an entitled data set. You can also make API calls to AWS Data Exchange-managed endpoints, which are then proxied through to API-owner endpoints.

A data set owner who has an existing Amazon API Gateway API can add an API asset using the AWS Data Exchange console, programmatically through the AWS CLI, or one of the AWS SDKs. For more information about importing API assets, see Importing AWS Data Exchange assets from an Amazon API Gateway API.

Note

Currently, the SendApiAsset operation is not supported for the following SDKs:

- SDK for .NET
- AWS SDK for C++
- SDK for Java 2.x

Data set owners who do not have an existing Amazon API Gateway API must create one before adding an API asset to their product. For more information, see <u>Developing a REST API in API Gateway in the Amazon API Gateway Developer Guide</u>.

Amazon Redshift datashare assets

With Amazon Redshift datashare assets, recipients can get read-only access to query the data in Amazon Redshift without extracting, transforming, and loading data.

Asset types 16

For more information about importing Amazon Redshift datashare assets, see <u>Importing AWS Data</u> Exchange assets from an AWS Data Exchange datashare for Amazon Redshift.

AWS Lake Formation data permission (Preview)

With AWS Lake Formation data permission assets, recipients or subscribers can access and query all databases, tables, or columns associated with the tags specified.

Data set owners must create and tag their data before importing the tags as part of an AWS Data Exchange asset. For more information about importing Lake Formation data permission assets, see Importing AWS Data Exchange assets from AWS Lake Formation (Preview).

Amazon S3 data access

With Amazon S3 data access assets, recipients or subscribers can directly access and use the provider's data without creating or managing data copies. Data set owners can set up AWS Data Exchange for Amazon S3 on top of their existing Amazon S3 buckets to share direct access to an entire S3 bucket or specific prefixes and Amazon S3 objects.

Revisions

A revision is a *container* for one or more assets.

You use revisions to update data in Amazon S3. For example, you can group a collection of .csv files or a single .csv file and a dictionary to create a revision. As new data is available, you create revisions and add assets. After you create and finalize the revision using the AWS Data Exchange console, that revision will be immediately available to subscribers. For more information, see Publishing a new product in AWS Data Exchange.

Keep the following in mind:

- To be finalized, a revision must contain at least one asset.
- It is your responsibility to ensure that the assets are correct before you finalize your revision.
- A finalized revision published to at least one data grant or product can't be unfinalized or changed in any way. (Except through the revoke revision process)
- After the revision is finalized, it is automatically published to your data grants or products.

Revisions 17

Revision structure

Revisions have the following parameters:

- DataSetId The ID of the data set that contains this revision.
- Comment A comment about the revision. This field can be 128 characters long.
- Finalized Either true or false. Used to indicate whether the revision is finalized.
- Id The unique identifier for the revision generated when it's created.
- Arn A unique identifier for an AWS resource name.
- CreatedAt Date and timestamp for the creation of the revision. Entitled revisions are created at the time of publishing.
- UpdatedAt Date and timestamp for the last update of the revision.
- Revoked A status indicating that subscribers' access to the revision was revoked.
- RevokedAt Date and timestamp indicating when subscriber access to the revision was revoked.
- RevocationComment A required comment to inform subscribers of the reason their access to the revision was revoked. The minimum required character length is 10. This field can be between 10 and 512 characters long.
- SourceID The revision ID of the owned revision corresponding to the entitled revision being viewed. This parameter is returned when a revision owner is viewing the entitled copy of its owned revision.

Example revision resource

```
{
    "UpdatedAt": "2019-10-11T14:13:31.749Z",
    "DataSetId": "1EXAMPLE404460dc9b005a0d9EXAMPLE2f",
    "Comment": "initial data revision",
    "Finalized": true,
    "Id": "e5EXAMPLE224f879066f9999EXAMPLE42",
    "Arn": "arn:aws:dataexchange:us-east-1:123456789012:data-sets/1EXAMPLE404460dc9b005a0d9EXAMPLE2f/revisions/e5EXAMPLE224f879066f9999EXAMPLE42",
    "CreatedAt": "2019-10-11T14:11:58.064Z"
}
```

Revision structure 18

Data sets

A data set in AWS Data Exchange is a collection of data that can change over time.

When recipients or subscribers access a Files data set, they're accessing a specific revision in the data set. This structure enables providers to change the data available in data sets over time without having to worry about changes to historical data.

When recipients or subscribers access an API data set, they're accessing a data set that contains API assets, which enable subscribers to make API calls to AWS Data Exchange-managed endpoints, which are then proxied through to provider endpoints.

When recipients or subscribers access an Amazon Redshift data set, they're accessing an AWS Data Exchange datashare for Amazon Redshift. This datashare gives subscribers read-only access to the schemas, tables, views, and user-defined functions that the data owner has added to the datashares.

When recipients or subscribers access an AWS Lake Formation data permission data set, they're accessing the databases, tables, and/or columns tagged with an LF-tag specified by the data set owner.

When recipients or subscribers access an Amazon S3 data access data set, they're granted readonly access to shared Amazon S3 objects hosted in the provider's Amazon S3 buckets. Recipients or subscribers can use this data directly with other AWS services.

To create, view, update, or delete data sets, providers can use the AWS Data Exchange console, AWS CLI, your own REST client, or one of the AWS SDKs. For more information about programmatically managing AWS Data Exchange data sets, see the <u>AWS Data Exchange API</u> Reference.

Topics

- Owned data sets
- Entitled data sets
- Data set types
- Amazon S3 data access data set
- AWS Lake Formation data set (Preview)
- AWS Regions and data sets
- Data set structure

Data sets 19

Data set best practices

Owned data sets

A data set is owned by the account that created it. Owned data sets can be identified using the origin parameter, which is set to OWNED.

Entitled data sets

Entitled data sets are a read-only view of a sender's owned data sets. Entitled data sets are created at time of data grant creation or product publishing and are made available to recipients or subscribers who have an active data grant or subscription to the product. Entitled data sets can be identified using the origin parameter, which is set to ENTITLED.

As a recipient, you can view and interact with your entitled data sets using the AWS Data Exchange API or in the AWS Data Exchange console.

As a data set owner, you also have access to the entitled data set view that your recipients or subscribers see. You can do so using the AWS Data Exchange API, or by choosing the data set name in the data grant or product page in the AWS Data Exchange console.

Data set types

The following data set types are supported in AWS Data Exchange:

- the section called "Files data set"
- API data set
- Amazon Redshift data set
- the section called "Amazon S3 data access data set"
- AWS Lake Formation data set (Preview)

Files data set

A Files data set is a data set that contains flat files permitted by Amazon S3.

As a recipient or subscriber, you can export data either locally (download to your computer) or to your Amazon S3 bucket.

Owned data sets 20

As a data set owner, you can import any type of flat file from your Amazon S3 bucket and add it to the data set.

API data set

An API data set is a data set that contains API assets. API assets enable recipients or subscribers to make API calls to AWS Data Exchange-managed endpoints, which are then proxied through to data set owner endpoints.

As a data set owner, you create an API in Amazon API Gateway and add it to the data set to license access to your API upon data grant creation or subscription.

Amazon Redshift data set

An Amazon Redshift data set includes AWS Data Exchange datashares for Amazon Redshift. When you subscribe to a data set with datashares, you are added as a consumer of the datashare. This gives you read-only access to the schemas, tables, views, and user-defined functions the data set owner has added to the datashares.

As a data set owner, you can create a database from the datashare in Amazon Redshift and then query live data without extracting, transforming, and loading files. You are automatically granted access to the datashare when your data grant or subscription is activated and lose access after your either of these expire.

As a data set owner, you create a datashare in Amazon Redshift and add it to the data set to license access to your datashare upon data grant creation or subscription.

Amazon S3 data access data set

With AWS Data Exchange for Amazon S3 data access, data recipients or subscribers can access third-party data files directly from data set owners' Amazon S3 buckets.

When you subscribe to an AWS Data Exchange for Amazon S3 data access product, AWS Data Exchange automatically does the following:

- Provisions an Amazon S3 access point. Amazon S3 Access Point is a feature of Amazon S3 that simplifies data sharing to an Amazon S3 bucket.
- Updates the S3 Access Point resource policies to grant you read-only access.

Amazon S3 data access data set

With AWS Data Exchange for Amazon S3, data set owners can share direct access to an entire Amazon S3 bucket or specific prefixes and Amazon S3 objects. In addition, AWS Data Exchange can be used to automatically manage data grants, subscriptions, entitlements, billing, and payments.

AWS Lake Formation data set (Preview)

An AWS Lake Formation data set is a data set that contains AWS Lake Formation data permission assets.

As a data recipient or subscriber, you can manage the data made available to you in your AWS Lake Formation. After creating resource links in your AWS Lake Formation, you can query the data using analytics services like Amazon Athena.

As a data set owner, you tag your data using LF-tags in AWS Lake Formation and import those tags as assets when creating your data set.

AWS Regions and data sets

Your data sets can be in any supported AWS Region, but all data sets in a single data grant or product must be in the same AWS Region.

Data set structure

Data sets have the following parameters:

- Name The name of the data set. This value can be up to 256 characters long.
- Description A description for the data set. This value can be up to 16,348 characters long.
- AssetType Defines the type of assets the data set contains.
- Origin A property that defines the data set as Owned by the account (for providers) or Entitled to the account (for subscribers).
- Id An ID that uniquely identifies the data set. Data set IDs are generated at data set creation. Entitled data sets have a different ID than the original owned data set.
- Arn A unique identifier for an AWS resource name.
- CreatedAt and UpdatedAt Date and timestamps for the creation and last update of the data set.



Note

As a data set owner, you can change some properties for owned data sets, like the **Name** or **Description**. Updating properties in an owned data set won't update the properties in the corresponding entitled data set.

Example data set resource

```
{
    "Origin": "OWNED",
    "AssetType": "S3_SNAPSHOT",
    "Name": "MyDataSetName",
    "CreatedAt": "2019-09-09T19:31:49.704Z",
    "UpdatedAt": "2019-09-09T19:31:49.704Z",
    "Id": "fEXAMPLE1fd9a5c8b0d2e6fEXAMPLEe1",
    "Arn": "arn:aws:dataexchange:us-east-2:123456789109:data-
sets/fEXAMPLE1fd9a5c8b0d2e6fEXAMPLEe1",
    "Description": "This is my data set's description that describes the contents of
 the data set."
}
```

Data set best practices

As a data set owner, when you create and update data sets, keep the following best practices in mind:

- The name of the data set is visible in the data grant or product details in the catalog. We recommend that you choose a concise, descriptive name so customers easily understand the content of the data set.
- The description is visible to recipients or subscribers who have an active data grant or subscription to the product. We recommend that you include coverage information and the features and benefits of the data set.

Tags

You can add tags to your owned data sets and their revisions. When you use tagging, you can also use tag-based access control in AWS Identity and Access Management (IAM) policies to control access to these data sets and revisions.

Data set best practices 23 Entitled data sets can't be tagged. Tags of owned data sets and their revisions are not propagated to their corresponding entitled versions. Specifically, recipients or subscribers, who have read-only access to entitled data sets and revisions, won't see the tags of the original owned data set.



Note

Currently, assets and jobs don't support tagging.

Tags

Creating data grants on AWS Data Exchange

At a high level, this is how to create a data grant on AWS Data Exchange:

- 1. **Create an AWS account** You must sign up for AWS and create a user before you can create data grants. For more information, see Setting up.
- 2. Create a data set, a revision, and import assets You can create data sets through the AWS Data Exchange console or API. Then, you can create revisions in the data set, and add assets to that revision.
- 3. Create a data grant To create a data grant, you must provide a data grant name and description, select the data set you wish to include in the data grant, specify the AWS account ID of the recipient you with to share the data grant with, and optionally set an end date on which the data grant should expire. For more information, see the following topics.
- 4. **Publish a new revision** You can update dynamic data sets over time by creating a new revision using the AWS Data Exchange API or console. These revisions can then be published to active data grants.



Before creating a data grant on AWS Data Exchange, review the information on Setting up.

The following topics explains more about how to publish a new data product on AWS Data Exchange.

Topics

- Programmatic access
- Creating a data grant on AWS Data Exchange containing file-based data
- Creating a data grant on AWS Data Exchange containing APIs
- Create a data grant on AWS Data Exchange containing Amazon Redshift data sets
- Creating a data grant on AWS Data Exchange containing Amazon S3 data access
- Creating a data grant on AWS Data Exchange containing AWS Lake Formation data permission data sets (Preview)

Programmatic access

AWS Data Exchange also offers programmatic access to its resources using the following API:

AWS Data Exchange API – Use these API operations to create, view, update, and delete data sets
and revisions. You can also use these API operations to import and export assets to and from
those revisions. For more information, see the AWS Data Exchange API Reference.

Creating a data grant on AWS Data Exchange containing filebased data

The following topics describe the process of creating a data set and a new data grant containing file-based data on AWS Data Exchange by using the AWS Data Exchange console. The process has the following steps:

Steps

- Step 1: Create assets
- Step 2: Create a data set
- Step 3: Create a revision
- Step 4: Import assets to a revision
- Step 5: Create a new data grant

Step 1: Create assets

Assets are the data in AWS Data Exchange. For more information, see Assets.

Before you create a new file-based data grant, you must:

1. Create your files.

AWS Data Exchange supports all file types.

2. Store your files as objects in Amazon Simple Storage Service (Amazon S3) or on your local computer.

For more information about storing files in Amazon S3, see the Amazon S3 User Guide.

Programmatic access 26

Step 2: Create a data set

Data sets in AWS Data Exchange are dynamic and are versioned using revisions, with each revision containing at least one asset. For more information, see Data in AWS Data Exchange.

To create a data set

- 1. Open your web browser and sign in to the <u>AWS Data Exchange console</u>.
- 2. In the left side navigation pane, under My data, choose Owned data sets.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- 4. In Select data set type, choose Files.
- In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see
 Data set best practices.
- 6. (Optional) Under Add tags optional, add tags.
- 7. Choose **Create data set**.

Step 3: Create a revision

In the following procedure, you create a revision after you've created a data set in the AWS Data Exchange console. For more information, see <u>Revisions</u>.

To create a revision

- 1. On the **Data set overview** section of the data set details page:
 - a. (Optional) Choose **Edit name** to edit information about your data set.
 - b. (Optional) Choose **Delete** to delete the data set.
- 2. In the **Revisions** section, choose **Create revision**.
- Under Define revision, provide an optional comment for your revision that describes the purpose of the revision.
- 4. (Optional) Under **Add tags optional**, add tags associated with the resource.
- 5. Choose Create revision.
- 6. Review, edit, or delete your changes from the previous step.

Step 2: Create a data set 27

Step 4: Import assets to a revision

In the following procedure, you import data assets, and then finalize the revision in the AWS Data Exchange console. For more information, see Assets.

To import assets to the revision

- Under the Jobs section of the data set details page, choose either Import from Amazon S3 or Upload (to upload from your computer), depending on where the data assets for the data set are currently stored.
- 2. Follow the prompts, depending on your selection. A job is started to import your asset into your data set.
- 3. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed.**
- 4. If you have more data to add, repeat Step 1.
- 5. In **Revision overview**, review your revision and its assets.
- 6. Choose Finalize revision.

You have successfully finalized a revision for a data set.

You can edit or delete a revision before you add it to a product.

Topics

- Edit a revision
- Delete a revision

Edit a revision

To edit the revision after you've finalized it

1. In **Revision overview**, choose **De-finalize**.

You see a message that the revision is no longer in the finalized state.

- 2. To edit the revision, from **Revision overview**, choose **Actions**, **Edit**.
- 3. Make your changes, and then choose **Update**.
- 4. Review your changes, and then choose **Finalize**.

Delete a revision

To delete the revision after you've finalized it

- In Revision overview, choose Delete. 1.
- Type **Delete** in the **Delete revision** dialog box, and then choose **Delete**. 2.



Marning

This deletes the revision and all of its assets. This action cannot be undone.

Step 5: Create a new data grant

After you've created at least one data set and finalized a revision with assets, you're ready to use that data set as a part of a data grant.

To create a new data grant

- In the left navigation pane of the AWS Data Exchange console, under Exchanged data grants, choose Sent data grants.
- 2. From **Sent data grants**, choose **Create data grant** to open the **Define data grant** wizard.
- 3. In the **Select owned data set** section, select the check box next to the data set you want to add.



Note

The data set you choose must have a finalized revision. Data sets without finalized revisions can't be added to data grants.

Unlike with data sets included in data products which are shared on AWS Marketplace, data sets added to data grants have no revision access rules, meaning a recipient of a data grant, once the data grant is approved, will have access to all finalized revisions of a given data set (including historical revisions finalized prior to the data grant creation).

In the Grant overview section, enter information the recipient will see about your data grant, including the **Data grant name** and **Data grant description**.

Choose Next.

For more information, see Product best practices in AWS Data Exchange.

- In the Recipient access information section, under AWS account ID, enter the AWS account ID of the recipient account who should receive the data grant.
- 7. Under **Access end date**, select a specific end date for when the data grant should expire or, if the grant should exist in perpetuity, select **No end date**.
- 8. Choose **Next**.
- 9. In the **Review and send** section, review your data grant information.
- 10. If you're sure that you want to create the data grant and send it to the chosen recipient, choose **Create and send data grant**.

You've now completed the manual portion of creating a data grant. The data grant will show on the **Sent data grants** tab on the **Sent data grants** page showing its status as **Pending acceptance** until the recipient account accepts it.

Creating a data grant on AWS Data Exchange containing APIs

The following topics describe the process of creating a REST API data set and adding it to a data grant that contains APIs on AWS Data Exchange. You can complete the process by using either the AWS Data Exchange console or the AWS Command Line Interface.

After you have set up your Amazon API Gateway REST API, you can create a new API data set in AWS Data Exchange. You can then create a revision, and add API assets.

Creating a data grant with an API asset allows recipient requests to an AWS Data Exchange endpoint to proxy through to your API Gateway API.

The process has the following steps:

Steps

- Prerequisites
- Step 1: Update the API resource policy
- Step 2: Create an API data set
- Step 3: Create a revision

Containing APIs 30

- Step 4: Add API assets to a revision
- Step 5: Create a new data grant containing APIs

Prerequisites

Before you can publish a product containing APIs, you must meet the following prerequisites:

- Before you can use any AWS service, including AWS Data Exchange, you must sign up for AWS and create an administrative user. For more information, see Getting started in the AWS IAM Identity Center User Guide.
- Your REST API must be on Amazon API Gateway with an integration that uses an appropriate request and response model for accessing your data, such as Amazon DynamoDB or AWS Lambda. For more information, see Developing a REST API in API Gateway and Working with REST APIs in the Amazon API Gateway Developer Guide.



(i) Note

Only public API Gateway APIs are supported.

 Your API Gateway REST API must be able to authenticate and authorize calls from the AWS Data Exchange service principal. Every request from AWS Data Exchange to your API uses the Signature Version 4 (SigV4) protocol signed with AWS Data Exchange credentials. AWS Data Exchange works with custom domains and domain key mappings.



Note

AWS Data Exchange doesn't support Amazon Cognito, No-Auth, and AWS Lambda authorizers.

- If your API Gateway REST API uses a custom identity system for authentication and authorization, configure it to use IAM authentication and import an OpenAPI schema describing your API. AWS Data Exchange will invoke your API Gateway REST API with its own service credentials and include subscriber information such as account ID.
- Your API Gateway REST API is responsible for integrating with your backend. To do this, do one of the following:
 - · Attach a long-lived authentication token to every request that comes through your API Gateway REST API that the backend can verify.

Prerequisites 31 Use API Gateway to invoke a Lambda function that can generate credentials and invoke your API.

Your API is invoked per the API integration request specification.

For more information, see the following topics:

Topics

- API data set security
- API integration request specification
- Header forwarding

API data set security

AWS Data Exchange encrypts traffic end to end using Transport Layer Security (TLS) 1.2. All metadata is encrypted at rest. AWS Data Exchange will not store subscriber requests or the responses from your backend.

API integration request specification

An API on AWS Data Exchange passes through all headers (except for the headers listed in <u>Header forwarding</u>), body, http method, path, and query strings as-is from the customer request and appends the following headers.

```
// These headers help prevent Confused Deputy attacks. They enable the SourceAccount
// and SourceArn variables in IAM policies.
'x-amz-source-account': ACCOUNT_ID,
'x-amz-source-arn': `arn:aws:dataexchange:${REGION}:${OWNER_ACCOUNT_ID}:data-sets/
${DATA_SET_ID}/revisions/${REVISION_ID}/assets/${ASSET_ID}`,

// These headers identify the API Asset in Data Exchange.
'x-amzn-dataexchange-asset-id': ASSET_ID,
'x-amzn-dataexchange-data-set-id': DATA_SET_ID,
'x-amzn-dataexchange-revision-id': REVISION_ID,

// This header identifies the Data Exchange Product.
'x-amzn-dataexchange-product-id': PRODUCT_ID,

// This header identifies the caller of Data Exchange. It will contain subscriber
```

Prerequisites 32

```
// information.
'x-amzn-dataexchange-requester-account-id': REQUESTER_ACCOUNT_ID,

// Providers can attach custom metadata in the form of key/value pairs
// to a particular subscription. We will send these key/value pairs as stringified
// JSON.
'x-amz-dataexchange-subscription-metadata': STRINGIFIED_METADATA,
```

Header forwarding

AWS Data Exchange removes any headers related to authentication or namespaced to Amazon prior to forwarding it to a data owner backend. Specifically, AWS Data Exchange removes:

- Authentication header
- Any headers that begin with x-amz

The host header will be overwritten as a consequence of the proxying.

Step 1: Update the API resource policy

If you have an Amazon API Gateway REST API that meets the <u>Prerequisites</u>, you must update your API resource policy to grant AWS Data Exchange the ability to invoke your API when a subscriber makes a request to get your API's schema.

To update your API resource policy

1. Add the following policy to your API's resource policy:

```
{
"Effect": "Allow",
"Principal": {"Service": "dataexchange.amazonaws.com"},
"Action": "execute-api:Invoke",
"Resource": "*",
"Condition": {"StringEquals": {"aws:SourceAccount": "<account-id>"}}
}
```

2. Replace account-id with the account that will be creating the API data set.

The account with the API Gateway resource does not need to be in the same account that is creating the data set.

This policy restricts these permissions to calls made by the AWS Data Exchange service principal and requires that only your account can authorize AWS Data Exchange to integrate with your API.



Note

If you have a resource policy that explicitly denies AWS Data Exchange from doing this invocation, you must remove or limit this deny.

You're now ready to create an API data set.

Step 2: Create an API data set

Data sets in AWS Data Exchange are dynamic and are versioned using revisions, with each revision containing at least one asset. For more information, see Data in AWS Data Exchange.

You use either the AWS Data Exchange console or the AWS Command Line Interface to create an API data set:

- Creating an API data set (console)
- Creating an API data set (AWS CLI)

Creating an API data set (console)

To create an API data set (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. On the left side navigation pane, under My data, choose Owned data sets.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- In Select data set type, choose Amazon API Gateway API. 4.
- In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see Data set best practices.
- 6. (Optional) Under Add tags - optional, add tags.
- Choose Create. 7.

You are now ready to create a revision.

Creating an API data set (AWS CLI)

To create an API data set (CLI)

1. Use the create-data-set command to create an API data set:

```
$ AWS dataexchange create-data-set \
-\\-asset-type API_GATEWAY_API \
-\\-description 'Data Set Description' \
-\\-name 'Data Set Name'

{
    "Arn": "arn:aws:dataexchange:us-east-1:123456789012:data-sets/$DATA_SET_ID",
    "AssetType": "API_GATEWAY_API",
    "CreatedAt": "2021-09-11T00:16:46.349000+00:00",
    "Description": "Data Set Description",
    "Id": "$DATA_SET_ID",
    "Name": "Data Set Name",
    "Origin": "OWNED",
    "UpdatedAt": "2021-09-11T00:16:46.349000+00:00"
}
```

Note the new Asset Type of API_GATEWAY_API.

You are now ready to create a revision.

Step 3: Create a revision

In the following procedure, you create a revision after you've created a data set. For more information, see Revisions.

You use either the AWS Data Exchange console or the AWS Command Line Interface to create a revision:

- Creating a revision (console)
- Creating a revision (AWS CLI)

Step 3: Create a revision 35

Creating a revision (console)

To create a revision (console)

- 1. On the **Data set overview** section of the data set details page:
 - a. (Optional) Choose **Edit name** to edit information about your data set.
 - b. (Optional) Choose **Delete** to delete the data set.
- 2. On the **Revisions** section, choose **Create revision**.
- 3. Under **Define revision**, provide an optional comment for your revision that describes the purpose of the revision.
- 4. (Optional) Under **Add tags optional**, add tags associated with the resource.
- 5. Choose Create revision.
- 6. Review, edit, or delete your changes from the previous step.

You are now ready to add API assets to the revision.

Creating a revision (AWS CLI)

To create a revision (AWS CLI)

1. Use the create-revision command to create a revision:

```
$ AWS dataexchange create-revision \
-\\-data-set-id $DATA_SET_ID \
-\\-comment 'First Atlas Revision'
{
   "Arn": "arn:aws:dataexchange:us-east-1:123456789012:data-sets/$DATA_SET_ID/
   revisions/$REVISION_ID",
   "Comment": "First Atlas Revision",
   "CreatedAt": "2021-09-11T00:18:49.160000+00:00",
   "DataSetId": "$DATA_SET_ID",
   "Finalized": false,
   "Id": "$REVISION_ID",
   "UpdatedAt": "2021-09-11T00:18:49.160000+00:00"
}
```

2. Add the API assets to the revision.

Step 3: Create a revision 36



You will need to know the ID of the API Gateway REST API you want to import as well as the stage.

Step 4: Add API assets to a revision

API assets contain the information subscribers need to make calls to your API. For more information, see Assets.

In the following procedure, you import data assets, and then finalize the revision.

You use either the AWS Data Exchange console or the AWS CLI to add assets to a revision:

- Adding API assets to a revision (console)
- Adding API assets to a revision (AWS CLI)

Adding API assets to a revision (console)

To add assets to the revision (console)

- Under the API assets section of the data set details page, choose Add API stage.
- Under Select API stage, for Amazon API Gateway API, enter an API in the input box or choose 2. one of the following from the drop-down list:
 - API in another AWS account this is a cross account API that you have been given permission to access.
 - In this AWS account this is an API in your AWS account.
 - If you chose API in another AWS account, enter the API ID and the API Stage name in the a. input boxes.
 - If you chose In this AWS account, choose the API Stage name from the drop-down list



You can create a new API stage by choosing **Create new** and following the steps in the Create new API on Amazon API Gateway modal. Once the new stage has been created, repeat Step 2.

- Under Advanced configuration optional, you can choose to Connect existing Amazon API Gateway usage plan to use the throttling and quota limits as defined in the existing usage plan, and enter the API key.
- Under **Document API for subscribers**, provide details about the API that the recipients will see 4. after they accept the data grant.
 - For **API** name, enter a name that recipients can use to identify the API asset.

Note

If an In this AWS account was selected, the API name is automatically populated, which you can modify if necessary.

If a API in another AWS account was selected, the API name is populated with a default name, which you should modify to so the recipient can easily understand what it is.

- For **OpenAPI 3.0 specification**, either:
 - i. Enter or copy and paste the OpenAPI 3.0 specification file.
 - ii. Choose Import from .JSON file, and then select the .json file from your local computer to import.

The imported specification appears in the box.

iii. Choose Import from Amazon API Gateway, and then choose a specification to import.

The imported specification appears in the box.

For **Additional documentation - optional**, enter any additional information that is useful c. for the subscriber to know about your API. Markdown is supported.



You can't edit the OpenAPI specification and additional documentation after you add this asset to a revision.

If you want to update this information, and the revision is not finalized, you can replace the asset.

If you want to update this information, and the revision is finalized, you can create a new revision with the updated asset.

5. Choose **Add API stage**.

A job is started to import your asset (in this case, the API) into your data set.



Note

If you do not have an API on Amazon API Gateway, you will be prompted to create one.

- After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**. 6.
- If you have more APIs to add, repeat Step 2.
- 8. Under **Revision overview**, review your revision and its assets.
- Choose Finalize.

You have successfully finalized a revision for a data set.

You can edit a revision or delete a revision before adding it to a data grant.

You are now ready to Create a new data grant containing APIs.

Adding API assets to a revision (AWS CLI)

You can add API assets by running an IMPORT_ASSET_FROM_API_GATEWAY_API job.

To add API assets to a revision (AWS CLI):

Use the create-job command to add API assets to the revision:

```
$ AWS dataexchange create-job \
  -\\-type IMPORT_ASSET_FROM_API_GATEWAY_API \
```

```
-\\-details '{"ImportAssetFromApiGatewayApi":
{"DataSetId":"$DATA_SET_ID","RevisionId":"$REVISION_ID","ApiId":"$API_ID","Stage":"$API_STA
{
    "Arn": "arn:aws:dataexchange:us-east-1:123456789012:jobs/$JOB_ID",
    "CreatedAt": "2021-09-11T00:38:19.875000+00:00",
    "Details": {
        "ImportAssetFromApiGatewayApi": {
            "ApiId": "$API_ID",
            "DataSetId": "$DATA_SET_ID",
            "ProtocolType": "REST",
            "RevisionId": "$REVISION_ID",
            "Stage": "$API_STAGE"
        }
    },
    "Id": "$JOB_ID",
    "State": "WAITING",
    "Type": "IMPORT_ASSET_FROM_API_GATEWAY_API",
    "UpdatedAt": "2021-09-11T00:38:19.875000+00:00"
}
$ AWS dataexchange start-job -\\-job-id $JOB_ID
$ AWS dataexchange get-job -\\-job-id $JOB_ID
{
    "Arn": "arn:aws:dataexchange:us-east-1:0123456789012:jobs/$JOB_ID",
    "CreatedAt": "2021-09-11T00:38:19.875000+00:00",
    "Details": {
        "ImportAssetFromApiGatewayApi": {
            "ApiId": "$API_ID",
            "DataSetId": "$DATA_SET_ID",
            "ProtocolType": "REST",
            "RevisionId": "$REVISION_ID",
            "Stage": "$API_STAGE"
            "ApiEndpoint": "string",
            "ApiKey": "string",
            "ApiName": "string",
            "ApiDescription": "string",
            "ApiSpecificationDownloadUrl": "string",
            "ApiSpecificationDownloadUrlExpiresAt": "string"
        }
    },
    "Id": "$JOB_ID",
    "State": "COMPLETED",
    "Type": "IMPORT_ASSET_FROM_API_GATEWAY_API",
    "UpdatedAt": "2021-09-11T00:38:52.538000+00:00"
```

}

2. Use the list-revision-assets command to confirm that the new asset was created properly:

```
$ AWS dataexchange list-revision-assets \
  -\\-data-set-id $DATA_SET_ID \
  -\\-revision-id $REVISION_ID
    "Assets": [
    {
        "Arn": "arn:aws:dataexchange:us-east-1:123456789012:data-sets/$DATA_SET_ID/
revisions/$REVISION_ID/assets/$ASSET_ID",
        "AssetDetails": {
            "ApiGatewayApiAsset": {
                "ApiEndpoint": "https://$API_ID.execute-api.us-
east-1.amazonaws.com/$API_STAGE",
                "ApiId": "$API_ID",
                "ProtocolType": "REST",
                "Stage": "$API_STAGE"
            }
        },
        "AssetType": "API_GATEWAY_API",
        "CreatedAt": "2021-09-11T00:38:52.457000+00:00",
        "DataSetId": "$DATA_SET_ID",
        "Id": "$ASSET_ID",
        "Name": "$ASSET_ID/$API_STAGE",
        "RevisionId": "$REVISION_ID",
        "UpdatedAt": "2021-09-11T00:38:52.457000+00:00"
    }
    ]
}
```

You're now ready to create a new data grant containing APIs.

Edit a revision

To edit the revision after you've finalized it

On the Revision overview, choose De-finalize.

You see a message that the revision is no longer in the finalized state.

- To edit the revision, from **Revision overview**, choose **Actions**, **Edit**. 2.
- 3. Make your changes, and then choose **Update**.
- Review your changes and then choose **Finalize**. 4.

Delete a revision

To delete the revision after you've finalized it

- On the Revision overview, choose Delete.
- 2. Type **Delete** in the **Delete revision** dialog box, and then choose **Delete**.

Marning

This deletes the revision and all of its assets. This action can't be undone.

Step 5: Create a new data grant containing APIs

After you've created at least one data set and finalized a revision with assets, you're ready to publish that data set as a part of a data grant.

To create a new data grant

- In the left navigation pane of the AWS Data Exchange console, under **Exchanged data grants**, choose **Sent data grants**.
- From **Sent data grants**, choose **Create data grant** to open the **Define data grant** wizard.
- 3. In the **Select owned data** set section, select the check box next to the data set you want to add.



Note

The data set you choose must have a finalized revision. Data sets without finalized revisions can't be added to data grants.

Unlike with data sets included in data products which are shared on AWS Marketplace, data sets added to data grants have no revision access rules, meaning a recipient of a data grant, once the data grant is approved, will have access to all finalized revisions

of a given data set (including historical revisions finalized prior to the data grant creation).

- 4. In the **Grant overview** section, enter information the recipient will see regarding your data grant, including the **Data grant name**, and **Data grant description**.
- 5. Choose Next.
- In the Recipient access information section, under AWS account ID, enter the AWS account ID of the recipient account who should receive the data grant.
- 7. Also, in the **Recipient access information** section, under **Access end date**, choose whether the data grant should run in perpetuity, selecting **No end date**, or if it should have an end date, selecting **Specific end date**, and choosing the desired end date.
- 8. Choose **Next**.
- 9. In the **Review and send** section, review your data grant information.
- 10. If you're sure that you want to create the data grant and send it to the chosen recipient, choose **Create and send data grant**.

You've now completed the manual portion of creating a data grant. The data grant appears on the **Sent data grants** tab on the **Sent data grants** page, with a status of **Pending acceptance** until the recipient account accepts it.

Create a data grant on AWS Data Exchange containing Amazon Redshift data sets

An Amazon Redshift data set contains AWS Data Exchange datashares for Amazon Redshift. When customers subscribe to a product containing datashares, they are granted read-only access to the tables, views, schemas, and user-defined functions that a data owner adds to the datashare.

As a data owner, you create an AWS Data Exchange for Amazon Redshift datashare in your cluster. Then, you add to the datashare the schemas, tables, views, and user-defined functions that you want the recipient to access. You then import the datashare to AWS Data Exchange, create a data set, add it to a data grant. Recipients are granted access to the datashare upon acceptance of the data grant request.

After you have set up your Amazon Redshift datashare in Amazon Redshift, you can create a new Amazon Redshift data set in AWS Data Exchange. You can then create a revision, and add Amazon

Redshift datashare assets. This allows requests to the AWS Data Exchange endpoint to proxy through to your Amazon Redshift datashare. You can then add this data set to a data grant.

The following topics describe the process of creating an Amazon Redshift data set and a data grant containing it using the AWS Data Exchange console. The process has the following steps:

Steps

- Step 1: Create an Amazon Redshift datashare asset
- Step 2: Create an Amazon Redshift data set
- Step 3: Create a revision
- Step 4: Add Amazon Redshift datashare assets to a revision
- Step 5: Create a new data grant

Step 1: Create an Amazon Redshift datashare asset

Assets are the data in AWS Data Exchange. For more information, see Assets.

To create an Amazon Redshift datashare asset

Create a datashare within your Amazon Redshift cluster.

For more information about how to create a datashare, see Working with AWS Data Exchange datashares as a producer in the Amazon Redshift Database Developer Guide.



Note

We recommend setting your datashare as publicly accessible. If you do not, customers with publicly accessible clusters will not be able to consume your data.

Step 2: Create an Amazon Redshift data set.

Step 2: Create an Amazon Redshift data set

An Amazon Redshift data set includes AWS Data Exchange datashares for Amazon Redshift. For more information, see Amazon Redshift data set.

To create an Amazon Redshift data set

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. On the left side navigation pane, under My data, choose Owned data sets.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- 4. In Select data set type, choose Amazon Redshift datashare.
- 5. In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see Data set best practices.
- 6. Under Add tags optional, add tags.
- 7. Choose **Create**.

Step 3: Create a revision

In the following procedure, you create a revision after you've created a data set in the AWS Data Exchange console. For more information, see Revisions.

To create a revision

- 1. On the **Data set overview** section of the data set details page:
 - a. (Optional) Choose **Edit name** to edit information about your data set.
 - b. (Optional) Choose **Delete** to delete the data set.
- 2. On the **Revisions** section, choose **Create revision**.
- 3. Under **Define revision**, provide an optional comment for your revision that describes the purpose of the revision.
- 4. Under **Add tags optional**, add tags associated with the resource.
- Choose Create.
- 6. Review, edit, or delete your changes from the previous step.

Step 4: Add Amazon Redshift datashare assets to a revision

In the following procedure, you add Amazon Redshift datashare assets to a revision, and then finalize the revision in the AWS Data Exchange console. For more information, see Assets.

Step 3: Create a revision 45

To add assets to the revision

- Under the AWS Data Exchange datashares for Amazon Redshift section of the data set details page, choose **Add datashares**.
- Under AWS Data Exchange datashares for Amazon Redshift, select the datashares and then choose Add datashare(s).



(i) Note

You can add up to 20 datashares to a revision.

A job is started to import your assets into your revision.

- After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**. 3.
- 4. If you have more data to add, repeat Step 1.
- 5. Under **Revision overview**, review your revision and its assets.
- Choose Finalize. 6.

You have successfully finalized a revision for a data set.

You can edit or delete a revision before you add it to a data grant.

Step 5: Create a new data grant

After you've created at least one data set and finalized a revision with assets, you're ready to use that data set as a part of a data grant.

To create a new data grant

- From the left navigation pane of the AWS Data Exchange console, under Exchanged data grants, choose Sent data grants.
- From **Sent data grants**, choose **Create data grant** to open the **Define data grant** wizard. 2.
- In the Select owned data set section, select the check box next to the data set you want to add.



The data sets you choose must have a finalized revision. Data sets without finalized revisions won't be added to data grants.

Unlike with data sets included in data products which are shared on AWS Marketplace, data sets added to data grants have no revision access rules, meaning a recipient of a data grant, once the data grant is approved, will have access to all finalized revisions of a given data set (including historical revisions finalized prior to the data grant creation).

- In the **Grant overview** section, enter information the recipient will see regarding your data grant, including the **Data grant name**, and **Data grant description**.
- Choose Next.

For more information, see Product best practices in AWS Data Exchange.

- In the **Recipient access information** section, under **AWS account ID**, enter the AWS account ID of the data grant receiver account.
- In the **Recipient access information** section, under **Access end date**, choose whether the data grant should run in perpetuity, selecting **No end date**, or if it should have an end date, selecting **Specific end date**, and choosing the desired end date.
- 8. Choose Next.
- In the **Review and send** section, review your data grant information.
- 10. If you're sure that you want to create the data grant and send it to the chosen recipient, choose Create and send data grant.

You've now completed the manual portion of creating a data grant. The data grant will show on the **Sent data grants** tab on the **Sent data grants** page showing its status as **Pending acceptance** until the recipient account accepts it.

Creating a data grant on AWS Data Exchange containing Amazon S3 data access

With AWS Data Exchange for Amazon S3, data owners can share direct access to Amazon S3 buckets or specific prefixes and Amazon S3 objects. Data owners also use AWS Data Exchange to automatically manage entitlements through data grants.

As a data owner, you can share direct access to an entire Amazon S3 bucket or specific prefixes and Amazon S3 objects without creating or managing copies. These shared Amazon S3 objects can be server-side encrypted with customer managed keys stored in AWS Key Management Service (AWS KMS) or with AWS managed keys (SSE-S3). For more information about monitoring your KMS keys and understanding encryption contexts, see the section called "Key management for Amazon S3 data access". When a receiver gains access to your data products, AWS Data Exchange automatically provisions an Amazon S3 access point and updates its resource policies on your behalf to grant recipients read-only access. Recipients can use the Amazon S3 access point aliases in places where they use Amazon S3 bucket names to access data in Amazon S3.

When the subscription ends, the receiver's permissions are revoked.

Before you can create a data grant containing Amazon S3 data access, you must meet the following prerequisites:

Prerequisites

- Confirm that the Amazon S3 buckets hosting the data are configured with the Amazon S3
 bucket owner enforced setting turned on ACLs Disabled. For more information, see <u>Controlling</u>
 ownership of objects and disabling ACLs for your bucket in the Amazon Simple Storage Service
 User Guide.
- Your shared objects must be in the Amazon S3 Standard Storage class, or be managed using Amazon S3 Intelligent Tiering, for recievers to access them successfully. If they're in other storage classes, or if you have enabled Intelligent Tiering with Deep Archive, your receivers will get errors because they won't have permission to RestoreObject.
- Confirm that the Amazon S3 buckets hosting the data has encryption disabled or encrypted with Amazon S3 managed keys (SSE-S3) or customer managed keys stored in AWS Key Management Service (AWS KMS).
- If you're using customer managed keys, you must have the following:

IAM permissions to kms: CreateGrant on the KMS keys. You can access these permissions
through the key policy, IAM credentials, or through an AWS KMS grant on the KMS key. For
more information about key management and understanding how AWS Data Exchange uses
AWS KMS grants, see Creating AWS KMS grants.

To provide access, add permissions to your users, groups, or roles:

• Users and groups in AWS IAM Identity Center:

Create a permission set. Follow the instructions in <u>Create a permission set</u> in the AWS IAM Identity Center User Guide.

• Users managed in IAM through an identity provider:

Create a role for identity federation. Follow the instructions in <u>Create a role for a third-party</u> identity provider (federation) in the *IAM User Guide*.

- IAM users:
 - Create a role that your user can assume. Follow the instructions in <u>Create a role for an IAM</u> user in the *IAM User Guide*.
 - (Not recommended) Attach a policy directly to a user or add a user to a user group. Follow the instructions in Adding permissions to a user (console) in the *IAM User Guide*.

Users need programmatic access if they want to interact with AWS outside of the AWS Management Console. The way to grant programmatic access depends on the type of user that's accessing AWS.

To grant users programmatic access, choose one of the following options.

Which user needs programmatic access?	То	Ву
Workforce identity (Users managed in IAM Identity Center)	Use temporary credentia Is to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Configuring the AWS CLI to use AWS IAM Identity Center in the AWS Command Line Interface User Guide. • For AWS SDKs, tools, and AWS APIs, see IAM Identity Center authentic ation in the AWS SDKs and Tools Reference Guide.
IAM	Use temporary credentia ls to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions in Using temporary credentia ls with AWS resources in the IAM User Guide.

Which user needs programmatic access?	То	Ву
IAM	(Not recommended) Use long-term credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Authenticating using IAM user credentials in the AWS Command Line Interface User Guide. • For AWS SDKs and tools, see Authenticate using long-term credentials in the AWS SDKs and Tools Reference Guide. • For AWS APIs, see Managing access keys for IAM users in the IAM User Guide.

Following is an example JSON policy that shows how you could add to the key policy of the KMS key.

```
{
    "Sid": "AllowCreateGrantPermission",
    "Effect": "Allow",
    "Principal": {
"AWS": "<IAM identity who will call Dataexchange API>"
    },
    "Action": "kms:CreateGrant",
    "Resource": "*"
}
```

The following policy shows an example policy addition for the IAM identity that is used.

JSON

Note

Cross account KMS keys are also permitted if the kms: CreateGrant permission on the KMS keys are obtained through the earlier step. If another account owns the key, you must have permissions on the key policy and your IAM credentials as detailed in the above examples.

- 2. Make sure to use KMS keys to encrypt existing and new objects in the Amazon S3 bucket using the Amazon S3 bucket key feature. For more details, see Configuring S3 Bucket Keys in the Amazon Simple Storage Service User Guide.
 - For new objects added to your Amazon S3 bucket, you can set up Amazon S3 bucket key
 encryption by default. If existing objects have been encrypted without using the Amazon
 S3bucket key feature, these objects must be migrated to use the Amazon S3 bucket key for
 encryption.

To enable the Amazon S3 bucket key for existing objects, use the copy operation. For more information, see Configuring an Amazon S3 bucket key at the object level using batch operations.

- AWS managed KMS keys or AWS owned keys aren't supported. You can migrate from an unsupported encryption scheme to the ones currently supported. For more information, see Changing your Amazon S3 encryption at the AWS Storage Blog.
- 3. Set the Amazon S3 buckets hosting the data to trust AWS Data Exchange owned access points. You must update these Amazon S3 bucket policies to give AWS Data Exchange permissions to create Amazon S3 access points and grant or remove subscribers' access on your behalf. If the policy statement is missing, you must edit the bucket policy to add the Amazon S3 locations to your data set.

An example policy is shown below. Replace <Bucket ARN> with the appropriate value.

JSON

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "AWS": "*"
            },
            "Action": [
                "s3:GetObject",
                "s3:ListBucket"
            ],
            "Resource": [
               "arn:aws:s3:::BucketName",
               "arn:aws:s3:::/*"
            ],
            "Condition": {
                "StringEquals": {
                     "s3:DataAccessPointAccount": [
                         "337040091392",
                         "504002150500",
                         "366362662752",
                         "330489627928",
                         "291973504423",
                         "461002523379",
                         "036905324694",
                         "540564263739",
                         "675969394711",
                         "108584782536",
```

```
"844053218156"

]
}
}
}
}
```

You can delegate data sharing through AWS Data Exchange to an entire Amazon S3 bucket. However, you can scope delegation to the specific prefixes and objects of the bucket that you want to share in the data set. Following is an example of a scoped policy. Replace <Bucket ARN> and "mybucket/folder1/*" with your own information.

JSON

```
{
 "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DelegateToAdxGetObjectsInFolder1",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::mybucket/folder1/*"
      ],
      "Condition": {
        "StringEquals": {
          "s3:DataAccessPointAccount": [
            "337040091392",
            "504002150500",
            "366362662752",
            "330489627928",
            "291973504423",
            "461002523379",
            "036905324694",
```

```
"540564263739",
            "675969394711",
            "108584782536",
            "844053218156"
          ]
        }
      }
    },
      "Sid": "DelegateToAdxListObjectsInFolder1",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::mybucket",
      "Condition": {
        "StringLike": {
          "s3:prefix": [
            "folder1/*"
          ]
        },
        "StringEquals": {
          "s3:DataAccessPointAccount": [
            "337040091392",
            "504002150500",
            "366362662752",
            "330489627928",
            "291973504423",
            "461002523379",
            "036905324694",
            "540564263739",
            "675969394711",
            "108584782536",
            "844053218156"
          ]
        }
      }
    }
  ]
}
```

Similarly, to scope access to only a single file, a data owner can use the following policy.

JSON

```
"Version": "2012-10-17",
  "Statement": [
      "Sid": "DelegateToAdxGetMyFile",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::mybucket/folder1/myfile"
      ],
      "Condition": {
        "StringEquals": {
          "s3:DataAccessPointAccount": [
            "337040091392",
            "504002150500",
            "366362662752",
            "330489627928",
            "291973504423",
            "461002523379",
            "036905324694",
            "540564263739",
            "675969394711",
            "108584782536",
            "844053218156"
          ]
        }
      }
    }
 ]
}
```

The following topics describe the process of creating an Amazon S3 data set and a data grant with Amazon S3 data sets using the AWS Data Exchange console. The process has the following steps:

Steps

- Step 1: Create an Amazon S3 data set
- Step 2: Configure Amazon S3 data access
- Step 3: Review and finalize the data set
- Step 4: Create a new data grant

Step 1: Create an Amazon S3 data set

To create an Amazon S3 data set

- 1. On the left side navigation pane, under My data, choose Owned data sets.
- 2. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- 3. In Select data set type, choose Amazon S3 data access.
- 4. In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see Data set best practices.
- 5. (Optional) Under Add tags optional, add tags.
- 6. Choose **Create data set** and continue.

Step 2: Configure Amazon S3 data access

Choose the Amazon S3 buckets or Amazon S3 bucket locations that you want to make available to recipients. You can select an entire Amazon S3 bucket, or specify up to five prefixes or objects within an Amazon S3 bucket. To add more Amazon S3 buckets, you must create another Amazon S3 data share

To configure shared Amazon S3 data access

- 1. On the Configure Amazon S3 data access page, select Choose Amazon S3 locations.
- In Choose Amazon S3 locations, enter your Amazon S3 bucket name in the search bar or select your Amazon S3 bucket, prefixes, or Amazon S3 files and choose Add selected. Then, choose Add locations.



We recommend choosing a top-level folder where a majority of objects and prefixes are stored so data owners don't need to reconfigure which prefixes or objects to share.

- In **Configuration details**, choose your **Requester Pays** configuration. There are two options: 3.
 - Enable Requester Pays (recommended) Requesters will pay for all requests and transfers in the Amazon S3 bucket. We recommend this option because it helps protect against unintended costs from receiver requests and transfers.
 - Disable Requester Pays You pay for receiver requests and transfers in the Amazon S3 bucket.

For more information about Requester Pays, see Objects in Requester Pays Buckets in the Amazon Simple Storage Service User Guide.

- Select the **Bucket Policy** that best suits your needs. Choose **General** to use one bucket policy for your entire Amazon S3 bucket. This is a one-time configuration and additional configuration isn't needed to share prefixes or objects in the future. Choose **Specific** to use a bucket policy that is specific to the selected Amazon S3 locations. Your shared Amazon S3 bucket needs a bucket policy in place to create an Amazon S3 data access data set successfully and can't have ACLs enabled.
 - To disable ACLs, navigate to your bucket permissions and set **Object Ownership** to **Bucket** a. owner enforced.
 - To add a bucket policy, copy the bucket statement to your clipboard. In the Amazon S3 console, from the Amazon S3 permissions tab, choose Edit in the bucket policy section, paste the bucket policy into the statement, and **Save changes**.
- If the Amazon S3 bucket contains objects encrypted using AWS KMS customer managed keys, you must share all such KMS keys with AWS Data Exchange. For information about required prerequisites when using KMS keys to encrypt objects in your Amazon S3 bucket, see the section called "Containing Amazon S3 data access". To share these KMS keys with AWS Data Exchange, do the following:
 - From the Configure Amazon S3 data access page, in Customer managed KMS keys, select Choose from your AWS KMS keys or Enter AWS KMS key ARN and select all AWS KMS keys currently being used to encrypt the Amazon S3 shared locations. AWS

Data Exchange uses these KMS keys to create grants for recipients to access your shared locations. For more information, see Grants in AWS KMS.



Note

AWS KMS has a limit of 50,000 grants per KMS key including pre-existing grants.

6. Review your Amazon S3 locations, selected KMS keys, and configuration details, and choose Save and continue.

Step 3: Review and finalize the data set

Review and finalize your newly created data set. If you wish to create and add another Amazon S3 data access to share access to additional Amazon S3 buckets, prefixes, objects, choose Add another Amazon S3 data access.



Note

We recommend this when needing to share access to data hosted in a different Amazon S3 bucket than the one previously picked in the initial Amazon S3 data access.

If you would like to make changes prior to publishing, you can save the data set as a draft by choosing **Save draft**. Then, choose **Finalize data set** to add it to your data grant.

Step 4: Create a new data grant

After you've created at least one data set and finalized a revision with assets, you're ready to use that data set as a part of a data grant.

To create a new data grant

- In the left navigation pane of the AWS Data Exchange console, under Exchanged data grants, choose Sent data grants.
- From **Sent data grants**, choose **Create data grant** to open the **Define data grant** wizard.
- 3. In the **Select owned data set** section, select the check box next to the data set you want to add.



The data set you choose must have a finalized revision. Data sets without finalized revisions can't be added to data grants.

Unlike with data sets included in data products which are shared on AWS Marketplace, data sets added to data grants have no revision access rules, meaning a recipient of a data grant, once the data grant is approved, will have access to all finalized revisions of a given data set (including historical revisions finalized prior to the data grant creation).

- In the **Grant overview** section, enter information the recipient will see about your data grant, including the **Data grant name** and **Data grant description**.
- Choose **Next**.

For more information, see Product best practices in AWS Data Exchange.

- In the **Recipient access information** section, under **AWS account ID**, enter the AWS account ID of the recipient account who should receive the data grant. .
- Under Access end date, select a specific end date for when the data grant should expire or, if the grant should exist in perpetuity, select **No end date**.
- Choose **Next**.
- In the **Review and send** section, review your data grant information.
- 10. If you're sure that you want to create the data grant and send it to the chosen recipient, choose Create and send data grant.

You've now completed the manual portion of creating a data grant. The data grant will show on the **Sent data grants** tab on the **Sent data grants** page showing its status as **Pending acceptance** until the recipient account accepts it.

Creating a data grant on AWS Data Exchange containing AWS Lake Formation data permission data sets (Preview)

If you're interested in creating data grants containing AWS Lake Formation data permission data sets during this Preview, contact AWS Support.

An AWS Lake Formation data permission data set contains a set of LF-tags and permissions for data managed by AWS Lake Formation. When customers accept data grants containing Lake Formation data permissions, they are granted read-only access to the databases, tables, and columns associated with the LF-tags added to the data set.

As a data owner, you start by creating LF-tags in AWS Lake Formation and associating those tags with the data you want to make available to recipients. For more information about tagging your resources in Lake Formation, see Lake Formation Developer Guide. Then you import those LF-tags and a set of data permissions into AWS Data Exchange as an asset. Recipients are granted access to the data associated with those LF-tags upon acceptance of the data grant.

The following topics describe the process of creating a data grant containing AWS Lake Formation data permissions. The process has the following steps:

Steps

- Step 1: Create an AWS Lake Formation data set (Preview)
- Step 2: Create an AWS Lake Formation data permission (Preview)
- Step 3: Review and finalize
- Step 4: Create a revision
- Step 5:Create a new data grant containing AWS Lake Formation data sets (Preview)
- Considerations when creating data grants containing an AWS Lake Formation data permission data set (Preview)

Step 1: Create an AWS Lake Formation data set (Preview)

To create an AWS Lake Formation data set

- 1. Open your web browser and sign in to the <u>AWS Data Exchange console</u>.
- 2. From the left navigation pane, under My data, choose Products.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- 4. In Select data set type, choose AWS Lake Formation data permission.
- 5. In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see the section called "Data set best practices".
- Under Add tags optional, choose Add new tag.

Choose Create data set and continue. 7.

Step 2: Create an AWS Lake Formation data permission (Preview)

AWS Data Exchange uses LF-Tags to grant data permissions. Choose the LF-Tags that are associated with the data you want to share to grant recipients permissions to the data.

To create AWS Lake Formation data permission

- 1. On the **Create Lake Formation data permission** page, choose **Add LF-Tag**.
- 2. Enter the **Key** and choose your LF-Tag **Values**.
- Choose **Preview resource(s)** to view how your LF-Tags are interpreted. 3.
 - From Preview resource(s), select your Associated data catalog resource(s).



Note

Make sure to revoke IAMAllowedPrincipals group on the following resources. For more information, see Revoking IAM role temporary security credentials in the IAM User Guide.

- Review the interpretation of the LF-Tag expression in the dialog box below and **Permissions** associated with the data set.
- For **Service access**, select your existing service role that allows AWS Data Exchange to assume the role and access, grant, and revoke entitlements to Lake Formation data permissions on your behalf. Then choose **Create Lake Formation data permission**. For more information about creating a role for an AWS service, see Creating a role to delegate permissions to an AWS service.

Step 3: Review and finalize

After creating your AWS Lake Formation data permission (Preview), you can Review and finalize your data set.

To review and finalize

Review your **Data set details** and **Tags** in **Step 1** for accuracy.

Review your LF-Tag expression(s), Add another Lake Formation data permission (optional), **Associated data catalog resource(s)**, and job details.



Note

Job are deleted 90 days after they're created.

Choose Finalize.

Step 4: Create a revision

To create a revision

- From the **Owned data sets** section, choose the data set for which you want to add a revision.
- Choose the **Revisions** tab. 2.
- In the **Revisions** section, choose **Create revision**.
- On the Revise Lake Formation data permission page, choose Add LF-Tag. 4.
- Review the **Permissions** for **Database** and **Table**. 5.
- From Service access, select an existing service role and then choose Create Lake Formation data permission.

Step 5:Create a new data grant containing AWS Lake Formation data sets (Preview)

After you've created at least one data set and finalized a revision with assets, you're ready to create a data grant with an AWS Lake Formation data permission data set.

To create a new data grant

- In the left navigation pane of the AWS Data Exchange console, under Exchanged data grants, choose Sent data grants.
- From **Sent data grants**, choose **Create data grant** to open the **Define data grant** wizard.
- 3. In the **Select owned data set** section, select the check box next to the data set you want to add.

Step 4: Create a revision 63



The data set you choose must have a finalized revision. Data sets without finalized revisions can't be added to data grants.

Unlike with data sets included in data products which are shared on AWS Marketplace, data sets added to data grants have no revision access rules, meaning a recipient of a data grant, once the data grant is approved, will have access to all finalized revisions of a given data set (including historical revisions finalized prior to the data grant creation).

- In the **Grant overview** section, enter information the recipient will see about your data grant, including the **Data grant name** and **Data grant description**.
- Choose Next.

For more information, see Product best practices in AWS Data Exchange.

- In the **Recipient access information** section, under **AWS account ID**, enter the AWS account ID of the recipient account who should receive the data grant. .
- Under Access end date, select a specific end date for when the data grant should expire or, if the grant should exist in perpetuity, select **No end date**.
- Choose **Next**.
- In the **Review and send** section, review your data grant information.
- 10. If you're sure that you want to create the data grant and send it to the chosen recipient, choose Create and send data grant.

You've now completed the manual portion of creating a data grant. The data grant will show on the **Sent data grants** tab on the **Sent data grants** page showing its status as **Pending acceptance** until the recipient account accepts it.

Considerations when creating data grants containing an AWS Lake Formation data permission data set (Preview)

To ensure an optimal receiver experience, we strongly advise against making any of the following modifications to any permissions where your product contains AWS Data Exchange for Lake Formation data sets (Preview).

- We recommend not deleting or modifying IAM roles passed to AWS Data Exchange in active data grants containing AWS Lake Formation data sets. If you delete or modify such IAM roles, the following issues occur:
 - AWS accounts that have access to the Lake Formation data permissions might retain access indefinitely.
 - AWS accounts that are the receivers of your data grant but have not yet received access to the Lake Formation data permissions will fail to receive access.

AWS Data Exchange will not be liable for any IAM roles that you delete or modify.

- We recommend that you don't revoke granted AWS Lake Formation data permissions from IAM
 roles passed to AWS Data Exchange in data grants containing AWS Lake Formation data sets. If
 you revoke granted data permissions from such IAM roles, the following issues occur:
 - AWS accounts that have access to the Lake Formation data permissions might retain access indefinitely.
 - AWS accounts that subscribe to your product but have not yet received access to the Lake Formation data permissions will fail to receive access.
- We recommend not revoking granted AWS Lake Formation data permissions from AWS accounts
 with active data grants containing AWS Lake Formation data sets. If you revoke granted data
 permissions from AWS accounts which are the receivers of your data grant, those accounts will
 lose access, creating a poor customer experience.
- We recommend setting the cross account version in your AWS Glue Data Catalog to version 3
 when creating data grants containing AWS Lake Formation data sets. If you downgrade the cross
 account version of your Data Lake Catalog while having active data grants containing AWS Lake
 Formation data sets, the AWS accounts that are the receivers of your data grant, but have not yet
 received access to the Lake Formation data permissions, may fail to get access to the data.

Accepting data grants and accessing data on AWS Data Exchange

The following steps describe the process of accepting a data grant on AWS Data Exchange using the AWS Data Exchange console:

Accepting a data grant

- 1. You must sign up for an AWS account and create a user before you can accept a data grant. For more information see the section called "Sign up for an AWS account".
- 2. In the left navigation pane of the AWS Data Exchange console, under **Exchanged data grants**, choose **Received data grants**.
- 3. Any data grants in which your AWS account is the receiver of will appear in the table under the tab of **Pending data grants** showing the pending data grant details with the status of **Pending acceptance**.
- 4. To accept a data grant, select the check box next to the data grant you wish to approve, and choose **Accept data grant**.
- 5. When the acceptance of the data grant has completed processing, the data grant will appear under the **Accepted and expired data grants** tab showing the data grant details with the status of **Accepted**.
- 6. After the acceptance of the data grant, choose the data grant name from the **Entitled data** sets table to access the data. You can also navigate to the **Entitled data** page from **My data** to view your data grant and to view all data sets shared with your account.
- 7. Next, use the included data sets. You can take any of the following actions depending on the type of data set you have access to:
 - a. Export the associated files to your Amazon Simple Storage Service (Amazon S3) or locally through a signed URL.
 - b. Call the Amazon API Gateway API.
 - c. Query the Amazon Redshift data share.
 - d. Access the Amazon S3 data.
 - e. Query the AWS Lake Formation data lake (Preview).



When you accept a data grant, you agree that your use of the underlying data set remains subject to the AWS Customer Agreement or other agreement with AWS governing your use of such services.

Related topics

- Access an AWS Data Exchange data set after accepting a data grant
- Access an AWS Data Exchange data set containing file-based data
- Access an AWS Data Exchange data set containing APIs
- Access an AWS Data Exchange data set containing Amazon Redshift data sets
- Access an AWS Data Exchange data set containing Amazon S3 data access
- Access an AWS Data Exchange data set containing AWS Lake Formation data sets (Preview)

Access an AWS Data Exchange data set after accepting a data grant

The following topics describe the process of accessing a data set on AWS Data Exchange using the AWS Data Exchange console.

Topics

- Access an AWS Data Exchange data set containing file-based data
- Access an AWS Data Exchange data set containing APIs
- Access an AWS Data Exchange data set containing Amazon Redshift data sets
- Access an AWS Data Exchange data set containing Amazon S3 data access
- Access an AWS Data Exchange data set containing AWS Lake Formation data sets (Preview)

Access an AWS Data Exchange data set containing file-based data

The following topics describe the process of accessing a data set containing file-based data stored as files on AWS Data Exchange. To complete the process, use the AWS Data Exchange console.

Related topics 67 After you successfully accept a data grant, you will have access to the data set include in it.

To view the data sets, revisions, and assets

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, under **My data**, choose **Entitled data sets**.
- 3. Choose a data set.
- 4. View the Data set overview, Auto-export destinations (Amazon S3 data sets only), the **Revisions**, and the **Description** of the data set.

(Optional) Exporting data

After your data grant is active, you can set up your Amazon S3 bucket to receive assets that you export. You can export the associated assets to Amazon S3 or you can use jobs with a signed URL.

If you want to export or download your data at a later time, including getting new revisions, see the section called "Exporting assets".



Important

We recommend that you consider Amazon S3 security features when exporting data to Amazon S3. For more information about general guidelines and best practices, see Security best practices for Amazon S3 in the Amazon Simple Storage Service User Guide. For more information about how to export data, see the section called "Exporting assets" and the section called "Exporting revisions".

Access an AWS Data Exchange data set containing APIs

The following topics describe the process of accessing a data set containing APIs on AWS Data Exchange using the AWS Data Exchange console.

Viewing an API

To view an API

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- In the left side navigation pane, under **My data**, choose **Entitled data sets**. 2.

- 3. Choose a data set.
- 4. Under the **Revisions** tab, choose a revision.
- 5. Under API assets, choose the API.
- 6. View the **Asset overview**.
- 7. Follow the guidance in the **Integration notes** to call the API.

Downloading the API specification

To download the API specification

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, under My data, choose Entitled data sets.
- 3. Choose a data set.
- 4. Under the **Revisions** tab, choose a revision.
- 5. Under API assets, choose the API.
- 6. On the **OpenAPI 3.0** specification, choose **Download API specification**.

The specification is downloaded onto your local computer. You can then export the asset to a third- party tool for SDK generation.

Making an API call (console)

You can call a single endpoint in the AWS Data Exchange console.

To make an API call from the console

- Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, under **My data**, choose **Entitled data sets**.
- 3. Choose a data set.
- 4. Under the **Revisions** tab, choose a revision.
- 5. Under **API assets**, choose the API.
- 6. For **Integration notes**:
 - a. Choose **Copy** to use the **Base URL**.
 - b. Choose **Copy** to use the **Code structure**.

c. Follow the information provided in the specification documentation to call the API.

Making an API call (AWS CLI)

To make an API call (AWS CLI)

Use the send-api-asset command to call the API.

```
$ AWS dataexchange send-api-asset \
--asset-id $ASSET_ID \
--data-set-id $DATA_SET_ID \
--revision-id $REVISION_ID \
--body "..." \
{
   "headers": {
   ...
},
   "body": "..."
}
```

Access an AWS Data Exchange data set containing Amazon Redshift data sets

Overview for recipients

An Amazon Redshift data set is a data set that contains AWS Data Exchange datashares for Amazon Redshift. Datashares give you read-only access to the tables, views, schemas, and user-defined functions that a data owner adds to the datashare.

As a recipient, after you accept a data grant, you get access to query the data in Amazon Redshift without extracting, transforming, and loading data. You lose access to the datashare after your data grant expires.



It might take a few minutes to access the datashares after you accept the data grant.

After accepting a data grant, you can do the following:

• Query data without extracting, transforming, or loading data.

Access the latest data as soon as the data owner updates it.

For more information, see <u>Managing AWS Data Exchange datashares</u> in the *Amazon Redshift Database Developer Guide*.

Access an AWS Data Exchange data set containing Amazon S3 data access

Overview for recipients

AWS Data Exchange for Amazon S3 allows recipients to access third-party data files directly from data owners' Amazon S3 buckets.

As a recipient, after you are entitled to an AWS Data Exchange for Amazon S3 data set, you can start your data analysis with AWS services such as Amazon Athena, SageMaker AI Feature Store, or Amazon EMR directly using the data owner's data in their Amazon S3 buckets.

Consider the following:

- Data owners have the option to enable Requester Pays, an Amazon S3 feature, on the Amazon S3 bucket hosting the data offered. If enabled, recipients pay to read, use, transfer, export, or copy data into theirAmazon S3 buckets. For more information, see <u>Using Requester Pays buckets</u> for storage transfers and usage in the Amazon Simple Storage Service User Guide.
- When you accept a data grant to an AWS Data Exchange for Amazon S3 data product, AWS Data
 Exchange automatically provisions an Amazon S3 access point and updates its resource policies
 to grant you read-only access. Amazon S3 access points is a feature of Amazon S3 that simplifies
 data sharing to an Amazon S3 bucket. For more information, see Managing data access with
 Amazon S3 access points in the Amazon Simple Storage Service User Guide.
- Before you use the Amazon S3 access point Amazon Resource Name (ARN) or alias to access the shared data, you must update your IAM permissions. You can verify that the current role and its associated policy allows GetObject and ListBucket calls to the provider's Amazon S3 bucket and the Amazon S3 access point provided by AWS Data Exchange.

The following sections describe the complete process of accessing an AWS Data Exchange for Amazon S3 data set after accepting a data grant by using the AWS Data Exchange console.

You can run gueries to analyze the data in-place without setting up your own Amazon S3 buckets, copying data files into Amazon S3 buckets, or paying associated storage fees. You access the same Amazon S3 objects that the data owner maintains allowing you to use the most current data available.

With a data grant, you can do the following:

- Analyze data without setting up individual Amazon S3 buckets, copying files, or paying storage fees.
- Access the latest provider data as soon as the data owner updates it.

To view the data sets, revisions, and assets

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- In the left side navigation pane, under My data, choose Entitled data sets. 2.
- 3. On the **Entitled data sets** page, choose a data set.
- View the **Data set overview**. 4.

Note

The data provided is stored in the data owner's Amazon S3 bucket. When accessing this data, you'll be responsible for the cost of the request and the data downloaded from the owner's Amazon S3 bucket, unless the owner specifies otherwise.

- Before getting started, your role must have IAM permissions to use your entitled Amazon S3 5. data access. On the Data set overview page, on the Amazon S3 data access tab, select Verify **IAM permissions** to determine if your role has the correct permissions to access your data.
- If you have the necessary IAM permissions, choose **Next** on the **IAM Policy** prompt displayed. If you don't have the needed permissions, follow the prompt to embed the JSON policy in the user or role.
- 7. Review your **Shared locations** to view the Amazon S3 bucket or prefixes and objects shared by the data owner. Review the data access information for Amazon S3 access point information to determine if the data owner enabled Requester Pays.
- Choose Browse shared Amazon S3 locations to view and explore the data owner's Amazon S3 bucket, prefixes, and objects shared.

- 9. Use the Access Point alias anywhere you use Amazon S3 bucket names to access your entitled data programmatically. For more information, see <u>Using access points with compatible</u>

 Amazon S3 operations in the *Amazon Simple Storage Service User Guide*.
- 10. (Optional) When you gain an entitlement to an Amazon S3 data access data set that contains data encrypted with a data owner's AWS KMS key, you can view the KMS key ARN in your console. AWS Data Exchange creates an AWS KMS grant on the key for you, so you can access the encrypted data. You must obtain kms:DecryptIAM permission on the AWS KMS key to read encrypted data from the Amazon S3 Access Point from which you've gained entitlement. You can choose between the following IAM policy statements:
 - a. IAM policy allowing users to decrypt or encrypt data with any KMS key.

JSON

```
{
    "Version": "2012-10-17",
    "Statement": [{
        "Effect": "Allow",
        "Action": ["kms:Decrypt"],
        "Resource": ["*"]
    }
]
}
```

b. IAM policy allowing users to specify the exact KMS key ARNs visible in the recipient console.

JSON

```
]
           }
     ]
}
```

AWS KMS grants can take up to 5 minutes for the operation to achieve eventual consistency. You might not have access to the Amazon S3 data access data set until this is complete. For more information, see Grants in AWS KMS in the AWS KMS key Management Service Developer Guide.

Access an AWS Data Exchange data set containing AWS Lake Formation data sets (Preview)

Overview for recipients

An AWS Lake Formation data set is a data set that contains AWS Lake Formation data permission assets.

As a recipient, you can accept a data grant containing AWS Lake Formation data sets. Once you're entitled to an AWS Data Exchange for AWS Lake Formation data set, you can guery, transform, and share access to the data within your AWS account using AWS Lake Formation, or across your AWS organization using AWS License Manager.

After you accept a data grant containing an AWS Lake Formation data set, you can use Lake Formation compatible query engines, like Amazon Athena, to query your data.

After acceptance of the data grant is complete, you must do the following:

- Accept the AWS Resource Access Manager (AWS RAM) share within 12 hours after you accept the data grant. You can accept the AWS RAM share from your entitled data sets page for your AWS Lake Formation data permission data set on the AWS Data Exchange console. You only need to accept an AWS RAM share once per provider. For more information about accepting a resource share invitation from AWS RAM, see Accepting a resource share invitation from AWS RAM.
- Navigate to AWS Lake Formation and create resource links from the new shared resources. 2.

Navigate to Amazon Athena or another AWS Lake Formation compatible query engine to query your data.

Sharing an AWS Data Exchange data grant license in an organization

When you accept a data grant, you receive a license that allows you to share the underlying data set under the following conditions:

- The data grant sender allows you to share the underlying data set.
- Your AWS account belongs to an organization. For more information about AWS Organizations, see the AWS Organizations User Guide.



Note

You can only share access with accounts in your organization.

The following topics explain how to share licenses across accounts.

Topics

- Prerequisites for license sharing
- Viewing your AWS Data Exchange licenses
- Sharing your AWS Data Exchange licenses

Prerequisites for license sharing

Before you can share licenses, you must complete the following setup tasks:

- In the AWS Data Exchange console, use the **Data Grant settings** page to enable integration with AWS Organizations.
- Give AWS Data Exchange permission to read information about accounts in your organization and manage licenses on your behalf so that it can create the associated license grants when you share your licenses. For more information, see Using service-linked roles for AWS Data Exchange, in this guide.

Viewing your AWS Data Exchange licenses

The following topics explain how to view your AWS Data Exchange licenses.

Topics

- Viewing all licenses
- · Viewing a single license

Viewing all licenses

You can use the AWS License Manager console to view all of the licenses for AWS Data Exchange data grants that you have access to.

To view all licenses

- 1. Sign in to the <u>AWS Management Console</u>.
- 2. Open the AWS License Manager console.
- 3. In the left side navigation pane, choose **Granted licenses**.
- 4. View all the licenses for your accepted data grants.

Viewing a single license

You can use the AWS License Manager console to view a single license for an AWS Data Exchange data grant.

To view a single license

- 1. Sign in to the AWS Data Exchange console.
- 2. Under Exchanged data grants, choose Received data grants.
- 3. Choose a data grant.
- 4. On the next page, choose **View license** or **Distribute with License Manager**. What you see varies, depending on the data grant's distribution permissions.
- 5. View the details on the **License detail** page.

Viewing your licenses 76

Sharing your AWS Data Exchange licenses

Overview for receivers

You can manage and share your AWS Data Exchange licenses with other accounts in your organization by using AWS License Manager.

For more details about using AWS License Manager with AWS managed licenses, see <u>Granted</u> <u>licenses and Seller issued licenses</u> in the *AWS License Manager User Guide*.

Sharing your licenses 77

Subscribing to AWS Data Exchange data products on AWS Data Exchange

At a high level, this is how to subscribe to AWS Data Exchange data products available through AWS Marketplace:

- 1. Create an AWS account You must sign up for AWS and create a user before you can subscribe to data products. For more information, see *Setting up*.
- 2. Browse the public catalog Products are published to the AWS Marketplace catalog. You can find products and review the associated public or custom offers and product details. If the provider has issued a private offer to your account, the product is available on the **My product offers** page of the AWS Data Exchange console.
- 3. **Submit a request for a subscription** You must submit a request to subscribe. The request form requires additional information about your identity and intended use case. For more information, see Subscription verification for subscribers in AWS Data Exchange.
- 4. **Subscriber subscribes to the product** If you subscribe to a paid product, you are billed on your AWS bill. You get access to the entitled data set.
- 5. **Uses the included data sets** You have access to the product data sets according to the terms of the data subscription agreement. You can take any of the following actions depending on the type of data set you have access to:
 - Export the associated files to your Amazon Simple Storage Service (Amazon S3) or locally through a signed URL.
 - Call the Amazon API Gateway API.
 - Query the Amazon Redshift data share.
 - Access the provider's Amazon S3 data.
 - Query the provider's AWS Lake Formation data lake (Preview).

For more information, see Jobs in AWS Data Exchange.

6. Request a data product recommendation – If you are not able to find a product in the catalog, you can use the Request data product page in the AWS Data Exchange console to request personalized recommendations from the AWS Data Exchange Data Discovery Team. For more information, see Request a recommendation for a data product.



When subscribing to data products from some non-US sellers, you might also receive a tax invoice from the seller. For more information, see Tax Help - AWS Marketplace Sellers.

Related topics

- Product subscriptions in AWS Data Exchange
- Getting started as a subscriber in AWS Data Exchange
- Subscribing to and accessing an AWS Data Exchange product
- Subscription verification for subscribers in AWS Data Exchange
- Sharing AWS Data Exchange license subscriptions in an organization
- Accepting Bring Your Own Subscription (BYOS) offers in AWS Data Exchange
- Accepting private products and offers in AWS Data Exchange
- AWS Data Exchange Heartbeat
- AWS Data Exchange for APIs (Test Product)
- Worldwide Event Attendance (Test Product) on AWS Data Exchange
- AWS Data Exchange for AWS Lake Formation (Test Product) (Preview)
- AWS Data Exchange for Amazon S3 (Test Product)
- AWS Data Exchange Provider-Generated Notifications (Test Product)
- Data in AWS Data Exchange

Product subscriptions in AWS Data Exchange

All AWS Data Exchange products are subscription-based. When you subscribe to a product, you agree to the product's offer terms, including the price, duration, payment schedule, data subscription agreement, and refund policy. When you subscribe to a product, you pay according to the payment schedule chosen by the provider for the duration that you subscribed to.



Important

The data subscription agreement (DSA) sets forth the provider's terms and conditions for the data product. The use of any data product subscribed to on AWS Data Exchange must

Related topics 79 also be in compliance with the AWS Customer Agreement or other agreement governing your use of AWS services.



Note

Data products that are part of the Open Data on AWS program are free for anyone to use and do not require a subscription. For more information, see Using Open Data on AWS data sets with AWS Data Exchange.

Each product's public offer terms can contain one or more price and duration combinations. When you subscribe to a product, you can choose the duration of the subscription. You can also choose whether you would like to enable auto-renewal for that subscription, if the provider has enabled it for the product.

If the data provider has indicated that the product contains any categories of sensitive or personal data, for example, mobile IDs, it will be displayed with the product details. For more information about the categories of sensitive data, see Sensitive categories of information in AWS Data Exchange.

If the data provider has indicated that the product contains protected health information (PHI) subject to the Health Insurance Portability and Accountability Act of 1996 (HIPAA), you may not export the product's data sets into your AWS account unless such AWS account is designated as a HIPAA account (as defined in the AWS Business Associate Addendum found in AWS Artifact).

After a subscription is processed and active, it appears on your AWS bill according to the payment schedule as part of your AWS Marketplace charges. For more information, see AWS Marketplace Paying for Products.

During the duration of your subscription, you can view and access all the product's data sets. You can also export the data sets' assets in jobs. For more information, see Jobs in AWS Data Exchange. Once a subscription has expired, you can no longer view or export the data sets.

Product subscriptions



For information about data sets and revisions, including details about what you have access to in your subscription, see Data sets and revisions.

If a provider decides to unpublish a product, you still have access to the data sets as long as your subscription is active. However, you cannot auto-renew the subscription when it expires.

You can view all of your active product subscriptions and auto-renewal status on the **Subscriptions** page of the AWS Data Exchange console. For more information, see Managing AWS Data Exchange subscriptions.

Important

If you enable auto-renew, and the product's offer terms have changed at the time of renewal, then the new product offer terms (including new price and new DSA) apply. This ensures that you keep access to the data regardless of potential changes to offer terms.

Visit the **Entitled data sets** page to find and access all of your entitled data sets in a specific AWS Region, based on your active subscriptions.

When you subscribe to a data product, we might share your contact information with the provider. For more information, see Security on AWS Marketplace in the AWS Marketplace Buyer Guide.

When you purchase a data product on AWS Data Exchange that has an upfront commitment, you will receive an invoice from Amazon Web Services (AWS) immediately. You can see charges for each data product by name in the Detail section of the invoice. You will receive separate bills for usage of AWS infrastructure and analytics services such as Amazon Simple Storage Service (Amazon S3) or Amazon Athena. For more information about AWS Billing and Cost Management, see Paying for products in the AWS Marketplace Buyer Guide.

When your subscription to an AWS Data Exchange Files data set ends, you retain access to any files that you already exported. Review your Data Subscription Agreement to verify if your agreement requires that you delete exported data when ending a subscription.

Product subscriptions

Data sets and revisions

Every product in AWS Data Exchange is made up of one or more data sets, each with one or more revisions. Data sets in AWS Data Exchange are typically different data, and revisions are newer or modified versions of the same data. For more information about data sets and revisions, see Data in AWS Data Exchange.

Each revision may contain all the data for the data set (updated for the revision), or just the new data since the previous revision. It is even possible that each revision has completely different data. What data to provide in each revision is up to the data provider.

When you subscribe to a product, you have access to all data sets in the product. When the data provider creates the offer, they give you access to 0 or more historical revisions, up to all historical revisions. They can also give you access to future revisions that are made available during your subscription period. The terms of the subscription are shown on the product details page in the AWS Data Exchange console.

After you subscribe to a product containing files, you can manually export each revision or asset individually, or you can select to automatically export new revisions to your Amazon S3 buckets (up to five buckets maximum) when the provider publishes new revisions. For more information, see Subscribing to and accessing an AWS Data Exchange product containing file-based data. For more information about how to export revisions, see Exporting revisions from AWS Data Exchange.

After you subscribe to a product containing an Amazon API Gateway API, you can view and invoke the data provider's API. For more information, see <u>Subscribing to and accessing an AWS Data Exchange product containing APIs</u>.

After you subscribe to a product containing Amazon Redshift data sets, you get access to query the data in Amazon Redshift. For more information, see <u>Subscribing to and accessing an AWS Data Exchange product containing Amazon Redshift data sets.</u>

After you subscribe to an Amazon S3 data access data set, you can view and directly use the provider's Amazon S3 objects. For more information, see the section called "Containing Amazon S3">the section called "Containing Amazon S3">the section called "Containing Amazon S3" data access".

After you subscribe to a product containing AWS Lake Formation data permission data sets (Preview), you can manage the data in AWS Lake Formation and query it with downstream services like Amazon Athena.

Data sets and revisions 82

Data dictionaries and samples

Some products have data sets that include data dictionaries and samples. To help you determine whether you want to subscribe to the product, you can view and download the data dictionaries and samples before you subscribe to it.

A *data dictionary* is a visual representation of the contents of a data set. It includes details about what columns are included and their meaning.

Samples are pieces of data that reflect the data that you would receive after you subscribe to the product. Samples can be any file type supported by Amazon S3.

For more information about how to discover a product that contains data dictionaries and samples, see Browse the catalog.

For more information about how to evaluate a product using data dictionaries and samples, see Evaluate products containing data dictionaries and samples.

Getting started as a subscriber in AWS Data Exchange

The following topics describe the complete process of becoming a data product subscriber on AWS Data Exchange using the AWS Data Exchange console. The process has the following steps:

Steps

- Step 1: Set up AWS Data Exchange
- Step 2: Browse the catalog
- Step 3: (Optional) Request a recommendation for a data product
- Step 4: (Optional) Evaluate products containing data dictionaries and samples
- Step 5: Subscribe to and access a product

Step 1: Set up AWS Data Exchange

Before you can use AWS Data Exchange, you must sign up for AWS and create a user. For more information, see Setting up AWS Data Exchange.

To set up AWS Data Exchange

1. Sign up for an AWS account. For more information, see Sign up for an AWS account.

Data dictionaries and samples 83

2. Create a user. For more information, see Create a user.

Step 2: Browse the catalog

You can find products and review the associated public or custom offers and product details on both AWS Marketplace and AWS Data Exchange.

If the provider has issued a private offer to your account, the product is available on the **My product offers page** of the AWS Data Exchange console. For more information, see <u>Subscribing to</u> AWS Data Exchange data products on AWS Data Exchange.

To browse the catalog

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, under **Discover data products**, choose **Browse catalog**.
- 3. Enter a term or phrase in the **Search** bar and then choose **Search**.
- 4. (Optional) Under **Browse catalog**, enter in a word or phrase and then choose **Search** to view results matching your query.
- 5. (Optional) Under **Refine results**, choose from one of the specific **Categories** to browse specific data products.
- (Optional) Under Refine results, use the Data set type filter and select from the following options to find products:
 - Files (Amazon S3 Objects) Products containing file-based data
 - Amazon Redshift Products containing Amazon Redshift datashares
 - API Products containing APIs
 - Access to Amazon S3 Products containing Amazon S3 data access
 - AWS Lake Formation Products containing AWS Lake Formation data permissions (Preview)
- 7. Select a product from the list of returned results, and review its product details page.

Step 3: (Optional) Request a recommendation for a data product

If you're unable to find a product in the catalog, you can request personalized recommendations from the AWS Data Exchange Data Discovery Team.

Step 2: Browse the catalog 84

To request a data product recommendation

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, under **Discover data products**, choose **Request data product**.
- 3. On the **Request data product** page, for **Details**, enter a **Data description**.
- 4. (Optional) Expand Additional details optional and complete the fields as directed.
 - a. Select one or more **Product categories**.
 - b. Enter an **Example data product URL**.
 - c. For Data set type, choose from Files (Amazon S3 Objects), Amazon API Gateway API, Amazon Redshift datashare, AWS Lake Formation data permissions (Preview) or Amazon S3 data access.
 - d. Enter specific details about the product you want including **Delivery cadence**, **Example** data product URL, Subscription start date, Subscription length, and Subscription budget.
 - e. If the **Data set type** you chose is **Amazon API Gateway API**, under **Subscription budget**, select **Including metered costs**.
- 5. For **Data providers**, choose from a list of **Existing providers** or enter the name of **Other providers**. Then indicate whether you have an existing relationship with the providers.
- 6. Choose Submit.

You should receive a response from the AWS Data Exchange Data Discovery Team within 2 business days.

Step 4: (Optional) Evaluate products containing data dictionaries and samples

A provider might include a data dictionary and samples of the data set with their product. To help you determine if the product's data set will meet your needs, you can view and download the data dictionary and samples before you subscribe. For more information, see Data dictionaries and samples.

You can perform the following actions to help with your evaluation of a product's data sets:

- View a data dictionary
- Download a data dictionary

- View and download all data dictionaries (for products containing multiple data sets)
- Preview a sample
- Download a sample

Viewing a data dictionary

A provider can add one data dictionary per data set that you can view.

To view a data dictionary

- 1. On the product detail page, choose the **Data dictionary and samples** tab.
- 2. View the data dictionary in one of the following ways:
 - Scroll down to the product Overview section to see the data dictionary under View data dictionaries.
 - Choose the **Data dictionaries and samples** tab, expand a data set row, choose the option button next to a data dictionary, and then choose **View all data dictionaries**.
- (Optional) Enter a keyword or phrase into the Search bar to search across all data sets and all tables.
- 4. (Optional) Modify your search and filters as necessary.

Downloading a data dictionary

A provider can add one data dictionary per data set that you can download.

To download a data dictionary

- 1. On the product detail page, choose the **Data dictionary and samples** tab.
- 2. Expand the data set row by choosing the expand icon (plus icon to the left of the data set name).
- 3. Choose the option button next to a data dictionary name.
- Choose Download.

The data dictionary file is downloaded to your computer.

Viewing and downloading all data dictionaries

If the product has multiple data sets, the provider might add a data dictionary for each data set. To evaluate all the data sets, you might want to view and download all data dictionaries.

To view and download all data dictionaries

- 1. On the product detail page, choose the **Data dictionary and samples** tab.
- 2. Choose View all data dictionaries.
- 3. In the **View data dictionaries** dialog box, choose the **Download (CSV)** to download the .csv file.

The .csv file is downloaded to your computer.

4. Choose **Close** to close the dialog box.

Previewing a sample

To preview a sample

- 1. On the product detail page, choose the **Data dictionary and samples** tab.
- 2. Expand the data set by choosing the expand icon (plus icon to the left of the data set name)
- 3. Choose the option button next to a sample name.
- 4. Choose **Preview sample (CSV only)** to preview the sample.
 - a. (Optional) In the preview dialog box, choose **Download** to download the .csv file.

The .csv file is downloaded to your computer.

b. Choose **Close** to close the dialog box.

Downloading a sample

To download a sample

- 1. On the product detail page, choose the **Data dictionary and samples** tab.
- 2. Expand the data set by choosing the expand icon (plus icon to the left of the data set name)
- 3. Choose the option button next to a sample name.
- Choose Download.

The sample is downloaded to your computer.

Step 5: Subscribe to and access a product

After you discover a product in the AWS Data Exchange catalog and determine that it meets your needs, you can subscribe to the product and then access the product.

If you subscribe to a paid product, you are billed on your AWS bill. You get access to the entitled data set. For more information, see <u>Subscribing to AWS Data Exchange data products on AWS Data Exchange</u>.

For more information about how to subscribe to products containing different types of data sets, see the following:

- Subscribing to and accessing an AWS Data Exchange product containing file-based data
- Subscribing to and accessing an AWS Data Exchange product containing APIs
- <u>Subscribing to and accessing an AWS Data Exchange product containing Amazon Redshift data</u> sets
- Subscribing to and accessing an AWS Data Exchange product containing Amazon S3 data access
- Subscribing to and accessing an AWS Data Exchange product containing AWS Lake Formation data sets (Preview)

Subscribing to and accessing an AWS Data Exchange product

The following topics describe the process of subscribing to and accessing a product on AWS Data Exchange using the AWS Data Exchange console.

Topics

- Subscribing to and accessing an AWS Data Exchange product containing file-based data
- Subscribing to and accessing an AWS Data Exchange product containing APIs
- Subscribing to and accessing an AWS Data Exchange product containing Amazon Redshift data sets
- Subscribing to and accessing an AWS Data Exchange product containing Amazon S3 data access
- <u>Subscribing to and accessing an AWS Data Exchange product containing AWS Lake Formation</u> data sets (Preview)

Viewing and downloading a data dictionary in AWS Data Exchange



Note

By subscribing to a product, you agree that your use of the product is subject to the provider's offer terms including pricing information and data subscription agreement (DSA). You also agree and acknowledge that AWS may share information about the transaction (including your payment terms and product usage metrics) with the respective seller, reseller, or underlying provider, as applicable, in accordance with the AWS Privacy Notice. AWS will issue invoices and collect payments from you on behalf of the provider through your AWS account. Your use of AWS services remains subject to the AWS Customer Agreement or other agreement with AWS governing your use of such services.

Subscribing to and accessing an AWS Data Exchange product containing file-based data

The following topics describe the complete process of subscribing to and accessing a product containing file-based data stored as files on AWS Data Exchange. To complete the process, use the AWS Data Exchange console.

For information about how to evaluate a product before subscribing, see Evaluate products containing data dictionaries and samples.

The process has the following steps:

Steps

- Step 1: Subscribing to a product containing the file-based data
- Step 2: Accessing a product containing file-based data

To practice subscribing to and accessing a product containing file-based data, see the AWS Data Exchange Heartbeat.

Step 1: Subscribing to a product containing the file-based data

If you subscribe to a paid product, you are billed on your AWS bill. You get access to all entitled data sets. For more information, see Subscribing to AWS Data Exchange data products on AWS Data Exchange.

To subscribe to a product containing the file-based data

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, under **Discover data products**, choose **Browse catalog**.
- 3. Under Refine results, use the Data set type filter and select Files (Amazon S3 Objects) to find products containing file-based data.
 - For more information, see Browse the catalog.
- Select a data product containing **Files (Amazon S3 Objects)**, and view its product detail page. 4.

The information on the product detail page includes a product description, the provider's contact information, and the details of the product's public offer. The public offer information includes price and durations, the data subscription agreement (DSA), and the refund policy. You can view the names of the data sets included in the product and the AWS Regions in which they are available. You can also continue to browse other product detail pages by choosing a product under Similar products.

If the provider has issued a custom offer to your account (for example, a private offer or Bring Your Own Subscription (BYOS) offer), you see those details, too.

- In the top right corner, choose **Continue to subscribe**. 5.
- Choose your preferred price and duration combination, choose whether to enable auto-6. renewal for the subscription, and review the offer details, including the DSA.



Note

Some products require subscription verification. For more information, see Subscription verification for subscribers in AWS Data Exchange.

7. Review the pricing information, choose the pricing offer, and then choose **Subscribe**.



If you subscribe to a paid product, you are prompted to confirm your decision to subscribe.

Under Data sets included with your subscription, view the listed Data sets. 8.

After the subscription finishes processing, you can choose a data set to access your entitled data or choose **View subscription** to view your subscription.

- (Optional) For **Set up exports optional**, select the check boxes for the data sets that contain the revisions that you want to export. Selecting a data set will prepare its most recently published revision to be exported.
 - Choose a **Simple** destination option to select an Amazon S3 bucket location or choose **Advanced** to configure an Amazon S3 key naming pattern. This choice determines where your revisions will be exported. For more information about using key patterns, see Key patterns when exporting asset revisions from AWS Data Exchange.
 - b. For **Auto-export future revisions**, choose whether to turn on or turn off automatic revision export:
 - On All future revisions will always be exported.
 - Off Only one export of the most recent revision will be exported.
 - Choose the **Encryption** options, and review the **Amazon S3 pricing**.



Note

If you choose to export using AWS Key Management Service (AWS KMS) encryption, make sure your account has the correct AWS Identity and Access Management (IAM) permissions to create and revoke grants on the AWS KMS key you choose. Without these permissions, automatic export will fail.

Choose **Export** to export the data to Amazon S3, or choose **Skip** if you prefer to wait d. and export or download later. For more information about how to export data after subscribing, see (Optional) Exporting data.



It can take a few minutes for your subscription to become active after you choose **Subscribe**. If you choose **Export** before the subscription is active, you are prompted to wait until it is complete.

After your subscription is active, your export will begin.

Navigating away from this page prior to your subscription becoming active will not prevent the subscription from processing. It will prevent your data export from occurring.

Step 2: Accessing a product containing file-based data

After you successfully subscribe to a product, you have access to the product data sets according to the terms of the data subscription agreement (DSA).

The following topic describes how to access a product containing file-based data.

Viewing data sets, revisions, and assets

To view the data sets, revisions, and assets

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, choose **Subscriptions**, and then choose your product.
- 3. View the data sets that are part of the product under **Entitled data sets**.
- Choose a data set.
- 5. View the **Data set overview**, **Auto-export job destinations** (Amazon S3 products only), the Revisions, and the Description of the data set.

For more information, see Data in AWS Data Exchange.

Choose a revision.

Revisions are listed from latest to oldest.

7. View the **Revision overview**, **Assets**, and the **Jobs** that have been performed.

For information about exporting file-based assets, see Exporting AWS Data Exchange assets to an S3 bucket as a subscriber (console).

(Optional) Exporting data

After your subscription is active, you can set up your Amazon S3 bucket to receive assets that you export.

You can export the associated assets to Amazon S3 or you can use jobs with a signed URL.

If you want to export or download your data at a later time, including getting new revisions, see Exporting AWS Data Exchange assets to an S3 bucket as a subscriber (console).



Important

We recommend that you consider Amazon S3 security features when exporting data to Amazon S3. For more information about general guidelines and best practices, see Security best practices for Amazon S3 in the Amazon Simple Storage Service User Guide.

For more information about how to export data, see Exporting assets from AWS Data Exchange and Exporting revisions from AWS Data Exchange.

Subscribing to and accessing an AWS Data Exchange product containing APIs

The following topics describe the complete process of subscribing to and accessing a product containing APIs on AWS Data Exchange by using the AWS Data Exchange console.

For information about how to evaluate a product before subscribing, see Evaluate products containing data dictionaries and samples.

The process has the following steps:

Steps

- Step 1: Subscribing to a product containing APIs
- Step 2: Accessing an API product

To practice subscribing to and accessing a product containing APIs, see the AWS Data Exchange for APIs (Test Product).

Step 1: Subscribing to a product containing APIs

If you subscribe to a paid product, you're billed on your AWS bill. You get access to all entitled data sets. For more information, see <u>Subscribing to AWS Data Exchange data products on AWS Data Exchange</u>.

A provider might include metered costs to their product containing APIs. If a provider decreases metered costs, the price decrease goes into effect immediately. If the provider increases metered costs, and you're an existing subscriber, the price increase goes into effect on the first day of the month, 90 days after the price increase was submitted OR upon renewal (whichever is sooner). An email message is sent to existing subscribers when the price change is submitted.

Example

For example, assume that a provider submits a metered cost price increase on May 10. Existing subscribers receive an email message about the price change. The price increase goes into effect on September 1.

To subscribe to a product containing APIs

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, under **Discover data products**, choose **Browse catalog**.
 - For more information, see Browse the catalog.
- Under Refine results, use the Data set type filter and select API to find products containing APIs.
 - For more information, see Browse the catalog.
- 4. Select a product containing APIs, and view its product detail page.

The information on the product detail page includes a product description, the provider's contact information, and the details of the product's public offer. The public offer information includes price and durations, metered costs (if included), the data subscription agreement (DSA), and the refund policy. You can view the names of the data sets included in the product and the AWS Regions in which they are available. You can also continue to browse other product detail pages by choosing a product under **Similar products**.

If the provider has issued a custom offer to your account (for example, a <u>private offer</u> or <u>Bring</u> <u>Your Own Subscription (BYOS) offer</u>), you see those details, too.

- Under **Public offer**, view the **API metered costs** (if included).
- (Optional) In the Metered cost calculator, choose Select metered cost and then enter the number of units to display an example of the cost.
- In the top right corner, choose **Continue to subscribe**. 5.
- Choose your preferred price and duration combination, choose whether to enable auto-6. renewal for the subscription, and review the offer details, including the DSA.



Some products require subscription verification. For more information, see Subscription verification for subscribers in AWS Data Exchange.

7. Review the pricing information, choose the pricing offer, and then choose **Subscribe**.



Note

If you subscribe to a paid product, you're prompted to confirm your decision to subscribe.

Under Data sets included with your subscription, view the listed Data sets. 8.

After the subscription finishes processing, you can choose a data set to access your entitled data or choose **View subscription** to view your subscription.

Step 2: Accessing an API product

The following topics provide details about how to access a product that includes API data sets:

Topics

- Viewing an API
- Downloading the API specification
- Making an API call (console)
- Making an API call (AWS CLI)

Viewing an API

To view an API

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, under My subscriptions, choose Entitled data.
- 3. Choose a data set.
- 4. Under the **Revisions** tab, choose a revision.
- 5. Under **API assets**, choose the API.
- 6. View the **Asset overview**.
- 7. Follow the guidance in the **Integration notes** to call the API.

Downloading the API specification

To download the API specification

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, under **My subscriptions**, choose **Entitled data**.
- 3. Choose a data set.
- 4. Under the **Revisions** tab, choose a revision.
- 5. Under API assets, choose the API.
- 6. On the OpenAPI 3.0 specification, choose Download API specification.

The specification is downloaded onto your local computer. You can then export the asset to a third-party tool for SDK generation.

Making an API call (console)

You can call a single endpoint in the AWS Data Exchange console.

To make an API call from the console

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, under **My subscriptions**, choose **Entitled data**.
- Choose a data set.
- 4. Under the **Revisions** tab, choose a revision.

- 5. Under API assets, choose the API.
- 6. For **Integration notes**:
 - a. Choose **Copy** to use the **Base URL**.
 - b. Choose **Copy** to use the **Code structure**.
 - c. Follow the information provided in the specification documentation to call the API.

Making an API call (AWS CLI)

To make an API call (AWS CLI)

• Use the send-api-asset command to call the API.

```
$ AWS dataexchange send-api-asset \
    --asset-id $ASSET_ID \
    --data-set-id $DATA_SET_ID \
    --revision-id $REVISION_ID \
    --body "..." \
{
        "headers": {
            ...
        },
        "body": "..."
}
```

Subscribing to and accessing an AWS Data Exchange product containing Amazon Redshift data sets

Overview for recipients

An Amazon Redshift data set is a data set that contains AWS Data Exchange datashares for Amazon Redshift. Datashares give you read-only access to the tables, views, schemas, and user-defined functions that a data provider adds to the datashare.

As a data subscriber, you can find and subscribe to products containing Amazon Redshift data sets. After your subscription starts, you get access to query the data in Amazon Redshift without extracting, transforming, and loading data. You lose access to a product's datashares after your subscription expires.

Consider the following:

• It might take a few minutes to access the datashares after your subscription starts.

The following sections describe the complete process of becoming an Amazon Redshift datashare product subscriber on AWS Data Exchange by using the AWS Data Exchange console.

For information about how to evaluate a product before subscribing, see <u>Evaluate products</u> containing data dictionaries and samples.

The process has the following steps:

Steps

- Step 1: Subscribing to products containing Amazon Redshift data sets
- Step 2: Accessing the AWS Data Exchange datashares for Amazon Redshift

To practice subscribing to and accessing a product containing Amazon Redshift data sets, see the Worldwide Event Attendance (Test Product) on AWS Data Exchange.

Step 1: Subscribing to products containing Amazon Redshift data sets

If you subscribe to a paid product, you're billed on your AWS bill. You get access to all data sets included in the product. For more information, see <u>Subscribing to AWS Data Exchange data</u> products on AWS Data Exchange.

To subscribe to a product containing Amazon Redshift data sets

- 1. Open and sign in to the <u>AWS Data Exchange console</u>.
- 2. In the left side navigation pane, under **Discover data products**, choose **Browse catalog**.
 - For more information, see **Browse the catalog**.
- Under Refine results, use the Data set type filter and select Amazon Redshift to find products containing Amazon Redshift datashares.
 - For more information, see **Browse** the catalog.
- 4. Select a product and view its product detail page.
 - The information on the product detail page includes a product description, the provider's contact information, and the details of the product's public offer. The public offer information

includes price and duration, the data subscription agreement (DSA), and the refund policy. You can view the names of the data sets included in the product and the AWS Regions in which they are available. You can also continue to browse other product detail pages by choosing a product under Similar products.

If the provider has issued a custom offer to your account (for example, a private offer or Bring Your Own Subscription (BYOS) offer), you see those details, too.

Important

Be sure to review the date, time, and duration of the cluster's maintenance window. During the maintenance window, you do not have access to the datashare.

- In the top right corner, choose **Continue to subscribe**. 5.
- 6. Review the **Product offer**, the **Subscription terms**, the **Data sets** that are included in the offer, and the **Support information**.
- Choose whether to enable Offer auto-renewal for the subscription



Note

Some products require subscription verification. For more information, see Subscription verification for subscribers in AWS Data Exchange.

Choose Subscribe. 8.



Note

If you subscribe to a paid product, you're prompted to confirm your decision to subscribe.

Under Data sets included with your subscription, view the listed Data sets.

After the subscription finishes processing, you can choose a data set to access your entitled data or choose **View subscription** to view your subscription.

Step 2: Accessing the AWS Data Exchange datashares for Amazon Redshift

You have access to the product's data sets according to the terms of the data subscription agreement (DSA). As a subscriber, your subscription to a product that includes AWS Data Exchange datashares for Amazon Redshift gives you read-only access to the tables, views, schemas, and functions within the datashare.

With a subscription, you can do the following:

- Query data without having to extract, transform, or load data.
- Access the latest provider data as soon as the provider updates it.

For more information, see Working with AWS Data Exchange datashares in the Amazon Redshift Database Developer Guide.



Note

You lose access to a product's datashares after your subscription expires.

For more information about how to subscribe to an Amazon Redshift data set, see Worldwide Event Attendance (Test Product) on AWS Data Exchange.

Subscribing to and accessing an AWS Data Exchange product containing Amazon S3 data access

AWS Data Exchange for Amazon S3 allows data subscribers to access third-party data files directly from data providers' Amazon S3 buckets.

As a data subscriber, after you are entitled to an AWS Data Exchange for Amazon S3 data set, you can start your data analysis with AWS services such as Amazon Athena, SageMaker AI Feature Store, or Amazon EMR directly using the provider's data in their Amazon S3 buckets.

Consider the following:

 Providers have the option to enable Requester Pays, an Amazon S3 feature, on the Amazon S3 bucket hosting the data offered. If enabled, subscribers pay to read, use, transfer, export, or copy data into their Amazon S3 buckets. For more information, see Using Requester Pays buckets for storage transfers and usage in the Amazon Simple Storage Service User Guide.

- When you subscribe to an AWS Data Exchange for Amazon S3 data product, AWS Data Exchange
 automatically provisions an Amazon S3 access point and updates its resource policies to grant
 you read-only access. Amazon S3 access points is a feature of Amazon S3 that simplifies data
 sharing to an Amazon S3 bucket. For more information, see Managing data access with Amazon
 S3 access points in the Amazon Simple Storage Service User Guide.
- Before you use the Amazon S3 access point Amazon Resource Name (ARN) or alias to access the shared data, you must update your IAM permissions. You can verify that the current role and its associated policy allows GetObject and ListBucket calls to the provider's Amazon S3 bucket and the Amazon S3 access point provided by AWS Data Exchange.

The following sections describe the complete process of becoming an AWS Data Exchange for Amazon S3 subscriber by using the AWS Data Exchange console.

The process has the following steps:

Steps

- Step 1: Subscribing to products containing Amazon S3 data access
- Step 2: Accessing a product containing Amazon S3 data access

Step 1: Subscribing to products containing Amazon S3 data access

If you subscribe to a paid product, you're billed on your AWS bill. You get access to all data sets included in the product. For more information, see <u>Subscribing to AWS Data Exchange data</u> products on AWS Data Exchange.

To subscribe to a product containing access to Amazon S3

- 1. Open and sign in to the <u>AWS Data Exchange console</u>.
- 2. In the left side navigation pane, under **Discover data products**, choose **Browse catalog**.
 - For more information, see **Browse the catalog**.
- 3. Under **Refine results**, use the **Data set type** filter and select **Access to Amazon S3** to find products containing access to Amazon S3 data.
 - For more information, see **Browse the catalog**.
- 4. Select a product and view its product detail page.

The information on the product detail page includes a product description, the provider's contact information, and the details of the product's public offer. The public offer information includes price and duration, the data subscription agreement (DSA), and the refund policy. You can view the names of the data sets included in the product and the AWS Regions in which they are available. You can also continue to browse other product detail pages by choosing a product under **Similar products**.

If the provider has issued a custom offer to your account (for example, a private offer or Bring Your Own Subscription (BYOS) offer), you see those details, too.

- In the top right corner, choose **Continue to subscribe**.
- Review the Product offer, the Subscription terms, the Data sets that are included in the offer, and the **Support information**.
- 7. Choose whether enable Offer auto-renewal for the subscription



Note

Some products require subscription verification. For more information, see Subscription verification for subscribers in AWS Data Exchange.

Choose Subscribe. 8.



Note

If you subscribe to a paid product, you're prompted to confirm your decision to subscribe.

Under Data sets included with your subscription, view the listed Data sets. 9.

After the subscription finishes processing, you can choose a data set to access your entitled data or choose **View subscription** to view your subscription.

Step 2: Accessing a product containing Amazon S3 data access

You can run queries to analyze the data in-place without setting up your own Amazon S3 buckets, copying data files into Amazon S3 buckets, or paying associated storage fees. You access the same Amazon S3 objects that the data provider maintains allowing you to use the most current data available.

With a subscription, you can do the following:

- Analyze data without setting up individual Amazon S3 buckets, copying files, or paying storage fees.
- Access the latest provider data as soon as the provider updates it.

To view the data sets, revisions, and assets

- Open and sign in to the AWS Data Exchange console. 1.
- 2. In the left navigation pane, under My subscriptions, choose Entitled data.
- 3. On the **Entitled data page**, expand a product and choose a data set.
- View the **Data set overview**. 4.



Note

The data provided is stored in the provider's Amazon S3 bucket. When accessing this data, you'll be responsible for the cost of the request and the data downloaded from the provider's Amazon S3 bucket, unless the provider specifies otherwise.

- 5. Before getting started, your role must have IAM permissions to use your entitled Amazon S3 data access. On the Data set overview page, on the Amazon S3 data access tab, select Verify **IAM permissions** to determine if your role has the correct permissions to access your data.
- If you have the necessary IAM permissions, choose **Next** on the **IAM Policy** prompt displayed. If you don't have the needed permissions, follow the prompt to embed the JSON policy in the user or role.
- Review your **Shared locations** to view the Amazon S3 bucket or prefixes and objects shared by the provider. Review the data access information for Amazon S3 Access Point information to determine if the provider enabled Requester Pays.
- Choose Browse shared Amazon S3 locations to view and explore the provider's Amazon S3 bucket, prefixes, and objects shared.
- Use the Access Point alias anywhere you use Amazon S3 bucket names to access your entitled data programmatically. For more information, see Using access points with compatible Amazon S3 operations in the Amazon Simple Storage Service User Guide.
- (Optional) When you gain an entitlement to an Amazon S3 data access data set that contains data encrypted with a provider's AWS KMS key, you can view the KMS key ARN in your

subscriber console. AWS Data Exchange creates an AWS KMS grant on the key for you, so you can access the encrypted data. You must obtain kms:Decrypt IAM permission on the KMS key to read encrypted data from the Amazon S3 Access Point from which you've gained entitlement. You can choose between the following IAM policy statements:

a. IAM policy allowing users to decrypt or encrypt data with any KMS key.

JSON

b. IAM policy allowing users to specify the exact KMS key ARNs visible in the subscriber console.

JSON

}] }



Note

AWS KMS grants can take up to 5 minutes for the operation to achieve eventual consistency. You might not have access to the Amazon S3 data access data set until this is complete. For more information, see Grant in AWS KMS in the AWS Key Management Service Developer Guide.

For more information about how to subscribe to an Amazon S3 data set, see the section called "Containing Amazon S3 data access".

Subscribing to and accessing an AWS Data Exchange product containing AWS Lake Formation data sets (Preview)

An AWS Lake Formation data set is a data set that contains AWS Lake Formation data permission assets.

As a data subscriber, you can find and subscribe to products containing AWS Lake Formation data sets. Once you're entitled to an AWS Data Exchange for AWS Lake Formation data set, you can query, transform, and share access to the data within your AWS account using AWS Lake Formation, or across your AWS organization using AWS License Manager.

Step 1: Subscribing to products containing AWS Lake Formation data sets

If you subscribe to a paid product, you're billed on your AWS bill. You get access to all data sets included in the product. For more information, see Subscribing to AWS Data Exchange data products on AWS Data Exchange.

To subscribe to a product containing AWS Lake Formation data sets

- Open and sign in to the AWS Data Exchange console. 1.
- 2. In the left side navigation pane, under **Discover data products**, choose **Browse catalog**.

For more information, see Browse the catalog.

- Under Refine results, use the Data set type filter and select AWS Lake Formation to find 3. products containing AWS Lake Formation data sets.
- 4. Select a product and view its product detail page.

The information on the product detail page includes a product description, the provider's contact information, and the details of the product's public offer. The public offer information includes price and duration, the data subscription agreement (DSA), and the refund policy. You can view the names of the data sets included in the product and the AWS Regions in which they're available. You can also continue browsing other product detail pages by choosing a product under Similar products.

If the provider has issued a custom offer to your account (for example, a private offer or Bring Your Own Subscription (BYOS) offer), you see those details, too.

- In the top right corner, choose **Continue to subscribe**. 5.
- Review the **Product offer**, the **Subscription terms**, the **Data sets** that are included in the offer, 6. and the **Support information**.
- 7. Choose whether to enable **Offer auto-renewal** for the subscription.



Note

Some products require subscription verification. For more information, see Subscription verification for subscribers in AWS Data Exchange.

- Choose **Subscribe**. If you subscribe to a paid product, you're prompted to confirm your 8. decision to subscribe.
- Under Data sets included with your subscription, view the listed Data sets.

After the subscription finishes processing, you can choose a data set to access your entitled data or choose **View subscription** to view your subscription.

Step 2: Accessing the AWS Data Exchange datashares for AWS Lake Formation

After you subscribe to a product containing AWS Lake Formation data sets, you can use Lake Formation compatible query engines, like Amazon Athena, to query your data.

After subscription completion, you must do the following:

- Accept the AWS Resource Access Manager (AWS RAM) share within 12 hours after you
 subscribe to the product. You can accept the AWS RAM share from your subscription page or
 the entitled data page for your AWS Lake Formation data permission data set on the AWS
 Data Exchange console. You only need to accept an AWS RAM share once per provider. For
 more information about accepting a resource share invitation from AWS RAM, see <u>Accepting a</u>
 resource share invitation from AWS RAM.
- 2. Navigate to AWS Lake Formation and create resource links from the new shared resources.
- 3. Navigate to Athena or another AWS Lake Formation compatible query engine to query your data.

Viewing and downloading a data dictionary in AWS Data Exchange

Providers can attach data dictionaries to all AWS Data Exchange products. The following procedures describe how to view and download a data dictionary.

For more information about data dictionaries and samples, see Data dictionaries and samples.

To view and download a data dictionary

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, under My subscriptions, choose Entitled data.
- 3. On the **Entitled data page**, expand a product and choose a data set.
- 4. On the data set details page, choose the **Data dictionary** tab.
 - View the data dictionary on the **Data dictionary** tab.
 - Download the data dictionary by choosing **Download** and then saving the file to your computer.

Subscription verification for subscribers in AWS Data Exchange

When you subscribe to AWS Data Exchange public products, you are required to submit additional information about your identity and your intended use case. The provider reviews this information before approving subscriptions.

Topics

- Completing a subscription request in AWS Data Exchange
- Reviewing your pending AWS Data Exchange subscription requests
- Email notifications to verify subscriptions in AWS Data Exchange

Completing a subscription request in AWS Data Exchange

AWS Data Exchange public products require subscription verification. After choosing **Continue to subscribe**, you must complete an additional form on the **Complete subscription request** page.

To complete a subscription request

- On the Complete subscription request page, review and choose the product offer (if more than one offer is available).
- 2. Review the **Subscription terms**, included **Data sets**, **Support information**, and **Refund policy**.
- 3. Choose if you want to renew the offer automatically when it expires.
- 4. On the **Subscription request form**, your AWS account ID will be added automatically. Complete the form by completing the following fields:
 - · Company name
 - Name
 - Email address
 - · Company location
 - Intended use case

In addition to your proposed use case, you may include additional comments that could help the provider evaluate your request.

5. Choose **Send subscription request to provider**.

After you submit your request, the provider has up to 45 days to approve or decline your request.

Reviewing your pending AWS Data Exchange subscription requests

Review your pending subscriptions for AWS Data Exchange products that require subscription verification.

To review your pending AWS Data Exchange subscription requests

- Open and sign in to the AWS Data Exchange console. 1.
- 2. Choose **Subscriptions**.
- 3. Choose **Subscription requests**.
- 4. Review the status of your pending subscription requests.

Each subscription request is uniquely identified by its ID. The ID is visible to both the provider and the subscriber. You can use the subscription request ID to identify the request in your communications with the provider.

After a provider approves your request, the subscription appears on the **Subscriptions** page.



Note

You can cancel a pending subscription request at any time as long as it hasn't expired or already been processed.

Email notifications to verify subscriptions in AWS Data Exchange

You receive an email notification to your AWS account email address when your request is approved, declined, or when it expires. Although most subscription request status changes result in an email notification, the delivery of these emails is on a best-effort basis.



Note

You will not receive email notifications for subscription request status changes that you have initiated yourself (for example, cancelling a subscription).

Sharing AWS Data Exchange license subscriptions in an organization

When you subscribe to AWS Data Exchange products, an agreement is created that grants you license to use those products. If your AWS account is a member of an organization, you can share that license for AWS Data Exchange products with the other accounts in that organization.

Email notifications 109



Note

For more information about AWS Organizations, see the AWS Organizations User Guide.

The following topics outline the process of sharing the licenses across accounts.

Topics

- Prerequisites for license sharing
- Step 1: View your licenses
- Step 2: Share your licenses

Prerequisites for license sharing

Before you can share licenses for data products, you must first set up license sharing for your organization. Complete the following tasks to set up license sharing for your organization:

- Give AWS Marketplace permission to manage licenses on your behalf so that it can create the associated license grants when you purchase or share your licenses. For more information, see Service-linked roles for AWS Marketplace in the AWS Marketplace Buyer Guide.
- Set up AWS License Manager for first use. For more information, see Getting started with AWS License Manager in the AWS License Manager User Guide.

Step 1: View your licenses

The following topics outline the process of viewing your licenses.

Topics

- Viewing all licenses
- Viewing a single license

Viewing all licenses

You can use the AWS License Manager console to view all of the licenses for AWS Data Exchange products that you purchased.

To view all licenses for your subscribed products

- 1. Sign in to the AWS Management Console.
- 2. Open the AWS License Manager console.
- 3. In the left navigation pane, choose **Granted licenses**.
- 4. View all the licenses for your subscribed products.

Viewing a single license

You can use the AWS License Manager console to view a single license for an AWS Data Exchange data grant.

To view a license for a single subscription

- 1. Sign in to the AWS Data Exchange console.
- 2. Under My subscriptions, choose Subscriptions.
- 3. Choose a subscription.
- 4. On the next page, choose **View license** or **Distribute with License Manager**. What you see varies, depending on the data grant's distribution permissions.
- 5. View the details on the **License detail** page.

Step 2: Share your licenses

You can manage and share your licenses with other accounts in your organization by using AWS License Manager.

For more details about using License Manager with AWS managed licenses, see <u>Granted licenses</u> and <u>Seller issued licenses</u> in the <u>AWS License Manager User Guide</u>.

Accepting Bring Your Own Subscription (BYOS) offers in AWS Data Exchange

As a subscriber, you might want to migrate your existing data subscriptions to AWS Data Exchange. Bring your own subscription (BYOS) functionality allows you to migrate and fulfill existing subscriptions with participating data providers at no additional cost.

Step 2: Share your licenses 111

With BYOS offers, any billing relationship between providers and subscribers continues. BYOS offers are not subject to fulfillment fees. As a subscriber, you receive an AWS Marketplace invoice for the subscription with no charge for a fulfilment fee.

Because the subscription lifecycle starts outside of AWS Data Exchange, the workflow for migrating the existing subscriptions to AWS Data Exchange using BYOS requires collaboration between the provider and subscriber.

Important

With BYOS offers, you're migrating a subscription that predates the availability of this product on AWS. AWS might verify your BYOS offer with the existing subscription agreement. If AWS cannot verify your BYOS offer, the offer and entitlements can be revoked without notice.

Before creating or accepting a BYOS offer on AWS Data Exchange, the provider and subscriber should perform the following steps together:

Prerequisites

- The provider and the subscriber contact each other about implementing a BYOS AWS Data Exchange solution.
- 2. The subscriber provides the AWS account ID that they want to use to subscribe to data products on AWS Data Exchange.

The subscriber accepts the BYOS offer as follows.

To accept a BYOS offer

- Sign in to the AWS Data Exchange console. 1.
- 2. In the left navigation pane, from **Discover data products**, choose **My product offers**.
- Select the offer to which you would like to subscribe. You can use the filter at the top of the 3. page to choose between All products, Private products, and Public products.
- Choose Continue to subscribe. 4.
- 5. Review the terms of the offer, the data subscription agreement, and the included data sets.

BYOS offers 112 If you accept the terms of the offer, review and accept the acknowledgement, and then choose Subscribe.

Accepting private products and offers in AWS Data Exchange

Data providers can provide a product to you in AWS Data Exchange that isn't available to the general public, or they can offer their product at terms that are different from the publicly available offer terms. A private offer can be different from the public offer in any dimension, including price, duration, payment schedule, data subscription agreement, or refund policy.

Note

Unlike Bring Your Own Subscription (BYOS) offers, private offers are not required to be based on an existing subscription that predates the product's availability on AWS Data Exchange.

The provider must create a custom offer for your AWS account ID to target the offer to you. If a private offer hasn't been extended to you, you can request one by contacting a provider using the contact information on the details page of the public offer.

As a subscriber, you can accept a private offer as follows.

To accept a private offer

- 1. Sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, from **Discover data products**, choose **My product offers**.
- Find the product offer you are looking for in the list. You can filter at the top of the page to choose between All products, Private products, or Public products.
- Select the offer to which you want to subscribe.
 - Under **Custom offers**, view the **API metered costs** (if included). a.
 - (Optional) In the Metered cost calculator, choose Select metered cost and enter the number of units to display an example of the cost.
- Choose Continue to subscribe. 5.
- 6. Review the terms of the offer, the payment schedule, the data subscription agreement, and the included data sets.

Private products and offers 113



Note

To accept a private offer with a multiple payment schedule, you must be on invoice billing terms. You can create a support ticket if you want to switch to invoice billing terms.

Private offers with a multiple payment schedule are not eligible for automatic renewal.

If you accept the terms of the offer, review and accept the acknowledgement, and then choose Subscribe.



Note

Your account is automatically invoiced according to the dates specified in the payment schedule.

Managing AWS Data Exchange subscriptions

The following topics describe how to manage your subscriptions in AWS Data Exchange.

Topics

- Viewing your AWS Data Exchange subscriptions
- Turning subscription auto-renewal on or off in AWS Data Exchange
- Unsubscribing from an AWS Data Exchange product

Viewing your AWS Data Exchange subscriptions

View your subscriptions through the AWS Data Exchange console.

To view your subscriptions

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, choose **Subscriptions**.
- 3. View the list of your subscriptions.

Managing subscriptions 114

Turning subscription auto-renewal on or off in AWS Data Exchange

Manage your subscription auto-renewals through the AWS Data Exchange console.

To turn subscription auto-renewal on or off

- 1. Open and sign in to the AWS Data Exchange console.
- In the left navigation pane, choose **Subscriptions**, and then choose your product. 2.
- 3. On the product detail page, under **Renewal terms**, turn on **Auto-renewal** or turn off **Auto**renewal.

A success message appears, confirming your updated renewal settings.

Unsubscribing from an AWS Data Exchange product

Use the AWS Data Exchange console to unsubscribe from a data product.



Note

If you require immediate removal of a subscription, contact AWS Data Exchange Customer Support by using the AWS Support Center Console.

To unsubscribe from a product

- 1. Open and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, under My subscriptions, choose Subscriptions.
- 3. Select the subscription from which you want to unsubscribe.
- Under Renewal terms, turn off the Auto-renewal enabled option. 4.
- Do not export any more data, and let the subscription run its course. 5.



(i) Note

For paid products, consult the provider's refund policy. Contact the provider for any exceptions.

Products for learning about interacting with AWS Data Exchange

AWS Data Exchange offers the following products that help subscribers understand how to subscribe to and interact with an AWS Data Exchange data product.

Topics

- AWS Data Exchange Heartbeat
- AWS Data Exchange for APIs (Test Product)
- Worldwide Event Attendance (Test Product) on AWS Data Exchange
- AWS Data Exchange for AWS Lake Formation (Test Product) (Preview)
- AWS Data Exchange for Amazon S3 (Test Product)
- AWS Data Exchange Provider-Generated Notifications (Test Product)

AWS Data Exchange Heartbeat

AWS Data Exchange Heartbeat (Test product) is a free product that subscribers can use to understand how to interact with an AWS Data Exchange product subscription. You can use it for testing purposes and to get familiar with the AWS Data Exchange API and concepts.

AWS Data Exchange Heartbeat contains a single data set named **Heartbeat**. Approximately every 15 minutes, a new revision is published to this data set.

Example content of a revision

Each new revision contains two assets:

- Epoch asset
- Manifest asset

Epoch asset

Each AWS Data Exchange Heartbeat revision contains a JSON file Amazon Simple Storage Service (Amazon S3) object that contains a single array. The array's name is TimestampsSinceLastRevision, and its value is a list of each UNIX Epoch second that has elapsed since the last revision.

The name of the asset is in the form Epoch{start}-{end}.json where {start} and {end} represent the Epoch seconds corresponding to the period of time covered by the revision.

Manifest asset

Each AWS Data Exchange Heartbeat revision contains a JSON file S3 object that contains metadata about the revision and the schema of the Epoch asset JSON file. The name of the asset is in the form Manifest{start}-{end}.json where {start} and {end} represent the Epoch seconds corresponding to the period of time covered by the revision. The following example shows the content of a manifest file.

```
{
        "manifestSchemaVersion":"1.0",
        "schema":"{
                \"type\":\"object\",
                \"properties\":{
                    \"TimestampsSinceLastRevision\":{
                         \"type\":\"array\",
                         \"description\":\"List of epoch timestamps in seconds.\",
                         \"items\":{
                             \"type\":\"number\",
                             \"description\":\"Epoch timestamp in seconds.\"
                          }
                     }
                 }
        }",
        "startTimestamp":1554898111,
        "endTimestamp":1554905311,
        "numberOfTimestamps":7201
}
```

The following topic describes how to subscribe to AWS Data Exchange Heartbeat on AWS Data Exchange.

Topics

• Subscribing to AWS Data Exchange Heartbeat on AWS Data Exchange

Subscribing to AWS Data Exchange Heartbeat on AWS Data Exchange

The following procedure shows how to browse the AWS Data Exchange catalog to find and subscribe to AWS Data Exchange Heartbeat.

AWS Data Exchange Heartbeat 117

To find and subscribe to AWS Data Exchange Heartbeat

- Open and sign in to the AWS Data Exchange console. 1.
- 2. From the left navigation pane, under **Discover data products**, choose **Browse catalog**.
- From the search bar, enter AWS Data Exchange Heartbeat and press Enter. 3.
- 4. Choose the AWS Data Exchange Heartbeat product to view its details page.
 - (Optional) To view the data dictionary, scroll down to the product **Overview** section to see the data dictionary under **Data dictionaries**.
 - (Optional) To download the data dictionary, choose the **Data dictionary and samples** tab, choose the option button next to **Data dictionary**, and then choose **Download**.
 - (Optional) To download the sample, choose the option button next to the sample name (Heartbeat manifest sample.json), and then choose **Download**.
- In the top right corner, choose **Continue to subscribe**. 5.
- 6. Choose your preferred price and duration combination, choose whether to enable autorenewal for the subscription, and review the offer details, including the data subscription agreement.



Note

AWS Data Exchange Heartbeat doesn't require subscription verification, but some products do. For more information, see Subscription verification for subscribers in AWS Data Exchange.

7. Review the pricing information, choose the pricing offer, and then choose **Subscribe**.



Note

AWS Data Exchange Heartbeat is a free product. If you subscribe to a paid product, you are prompted to confirm your decision to subscribe.

On the **Set up your first export** page, select the check boxes for the data sets containing the 8. revisions you would like to export. Selecting a data set will prepare its most recently published revision to be exported.

- Choose an Amazon S3 bucket location or configure an Amazon S3 key naming pattern. This will determine where your revisions will be exported. For more information about using key patterns, see Key patterns when exporting asset revisions from AWS Data Exchange.
- 10. Choose **Export** to export the data to Amazon S3, or choose **Skip** if you'd rather wait and export or download later.

Note

It can take a few minutes for your subscription to become active after you choose **Subscribe.** If you choose **Export** before the subscription is active, you are prompted to wait until it is complete. After your subscription is active, your export will begin. Navigating away from this page prior to your subscription becoming active will not prevent the subscription from processing. It will prevent your data export from occurring.

AWS Data Exchange for APIs (Test Product)

AWS Data Exchange for APIs (Test Product) is a free product that is made available to subscribers to understand how to interact with an AWS Data Exchange product containing API data sets. You can use this product for testing purposes and to learn how to make API calls to providers in order to retrieve API-based data.

AWS Data Exchange for APIs (Test Product) contains an API data set named AWS Data Exchange for APIs (Test Product) that is in the US East (N. Virginia) Region.

Topics

- Subscribing to AWS Data Exchange for APIs (Test Product) on AWS Data Exchange
- Viewing the AWS Data Exchange API
- Downloading the AWS Data Exchange API specification
- Making an AWS Data Exchange API call

Subscribing to AWS Data Exchange for APIs (Test Product) on AWS Data Exchange

The following procedure shows how to browse the AWS Data Exchange catalog to find and subscribe to AWS Data Exchange for APIs (Test Product).

AWS Data Exchange for APIs 119

To find and subscribe to AWS Data Exchange for APIs (Test Product)

- Open and sign in to the AWS Data Exchange console. 1.
- 2. From the left navigation pane, under **Discover data products**, choose **Browse catalog**.
- From the search bar, enter AWS Data Exchange for APIs and press Enter. 3.
- 4. Choose the AWS Data Exchange for APIs (Test Product) and view its details page.
 - (Optional) To download the sample, choose the **Data dictionary and samples** tab, choose the option button next to the sample name (ADX for APIs sample.json), and then choose Download.
- In the top right corner, choose **Continue to subscribe**. 5.
- Choose the **Product offer**. 6.



Note

AWS Data Exchange for APIs (Test Product) is a free product.

- Review the **Subscription terms**, **Data sets**, and **Support information**. 7.
- Choose whether to enable **Offer auto-renewal** for the subscription. 8.



Note

AWS Data Exchange for APIs (Test Product) doesn't require subscription verification, but some products do. For more information, see Subscription verification for subscribers in AWS Data Exchange.

Choose Subscribe. 9.

It can take a few minutes for your subscription to become active after you choose **Subscribe**. Navigating away from this page before your subscription becomes active will not prevent the subscription from processing.

Viewing the AWS Data Exchange API

You can view the API with AWS Data Exchange for APIs (Test Product) using the following steps.

AWS Data Exchange for APIs 120

To view the API

- 1. Open and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under My subscriptions, choose Entitled data.
- Choose the product titled AWS Data Exchange for APIs (Test Product) and then choose the AWS Data Exchange for APIs data set.
- 4. Under the **Revisions** tab, choose a revision.
- 5. Under API assets, choose the API.
- View the Asset overview.
- 7. Follow the guidance in the **Integration notes** to call the API.

Downloading the AWS Data Exchange API specification

You can download the API specification with AWS Data Exchange for APIs (Test Product) using the following steps.

To download the API specification

- 1. Open and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under My subscriptions, choose Entitled data.
- 3. Choose the product titled **AWS Data Exchange for APIs (Test Product)** and then choose the **AWS Data Exchange for APIs** data set.
- 4. Under the **Revisions** tab, choose a revision.
- 5. Under **API assets**, choose the API.
- 6. On the **OpenAPI 3.0 specification**, choose **Download API specification**.

The specification is downloaded onto your local computer. You can then export the asset to a third-party tool for SDK generation.

Making an AWS Data Exchange API call

You can call a single endpoint in the AWS Data Exchange console.

To make an API call from the console

1. Open and sign in to the AWS Data Exchange console.

AWS Data Exchange for APIs 121

- 2. From the left navigation pane, under My subscriptions, choose Entitled data.
- 3. Choose the product titled AWS Data Exchange for APIs (Test Product) and then choose the AWS Data Exchange for APIs data set.
- 4. Under the **Revisions** tab, choose the revision.
- 5. Under API assets, choose the API.

You will see the sample **Code structure** and **OpenApi 3.0 specification** to structure your API request, which you can use in the AWS Command Line Interface to call the API.

- 6. Under **Integration notes**, choose **Copy** to copy the **Code structure** and then paste it into the AWS CLI.
- 7. Replace the sample values with the parameter key-value pairs you need using the information in the specification documentation.

Following is a sample API request for AWS Data Exchange for APIs (Test Product).

```
aws dataexchange send-api-asset \
    --data-set-id 8d494cba5e4720e5f6072e280daf70a8 \
    --revision-id b655d5be3da04fcbdca21a5a2932d789 \
    --asset-id 8550cfab16b444a794402f2c3f11eae1 \
    --method POST \
    --path "someresource" \
    --query-string-parameters 'param1=value1, param2=value2' \
    --request-headers 'header=header_value' \
    --body "{\"body_param\":\"body_param_value\"}"
```

Worldwide Event Attendance (Test Product) on AWS Data Exchange

Worldwide Event Attendance (Test Product) is a free product that helps subscribers understand how to subscribe to and interact with an AWS Data Exchange product containing Amazon Redshift data sets. You can use this product for testing purposes and to learn how to query, analyze, and build applications within minutes.

Worldwide Event Attendance (Test Product) contains an Amazon Redshift data set named **Worldwide Event Data (Test Data)** that is in the US East (N. Virginia) AWS Region.

You use the AWS Data Exchange console to find and subscribe to Worldwide Event Attendance (Test Product). Then, you can use either the Amazon Redshift console or SQL commands to query the datashare.

Worldwide Event Attendance 122

Topics

- Subscribing to Worldwide Event Attendance (Test Product) on AWS Data Exchange
- Querying Worldwide Event Attendance (Test Product) data with an Amazon Redshift cluster (console)
- Querying Worldwide Event Attendance (Test Product) data on Amazon Redshift (SQL)

Subscribing to Worldwide Event Attendance (Test Product) on AWS Data Exchange

The following procedure shows how to browse the AWS Data Exchange catalog to find and subscribe to Worldwide Event Attendance (Test Product).

To find and subscribe to Worldwide Event Attendance (Test Product)

- 1. Open and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Discover data products**, choose **Browse catalog**.
- From the search bar, enter Worldwide Event Attendance (Test Product) and press Enter.
- 4. Choose the Worldwide Event Attendance (Test Product) to view its details page.
 - a. (Optional) To view the data dictionary, scroll down to the product **Overview** section to see the data dictionary under **Data dictionaries**.
 - b. (Optional) To download the data dictionary, choose the **Data dictionary and samples** tab, choose the option button next to **Data dictionary**, and then choose **Download**.
 - c. (Optional) To preview the sample, choose the option button next to the sample name (Worldwide Event Attendance Sample.csv), and then choose Preview sample (CSV only).
 - d. (Optional) To download the sample, choose the option button next to the sample name (Worldwide Event Attendance Sample.csv), and then choose Download.
 - If you are previewing the sample, you can also choose **Download** in the sample preview dialog box.
- 5. In the top right corner, choose **Continue to subscribe**.
- 6. Choose your preferred price and duration combination, choose whether to enable autorenewal for the subscription, and review the offer details, including the data subscription agreement.

Worldwide Event Attendance 123



Note

Worldwide Event Attendance (Test Product) doesn't require subscription verification, but some products do. For more information, see Subscription verification for subscribers in AWS Data Exchange.

7. Review the pricing information, choose the pricing offer, and then choose **Subscribe**.



Note

Worldwide Event Attendance (Test Product) is a free product. If you subscribe to a paid product, you are prompted to confirm your decision to subscribe.

It can take a few minutes for your subscription to become active after you choose **Subscribe**.

Navigating away from this page prior to your subscription becoming active will not prevent the subscription from processing.

Querying Worldwide Event Attendance (Test Product) data with an Amazon Redshift cluster (console)

The following procedure shows how to set up and query the datashare using the Amazon Redshift console.

To query Worldwide Event Attendance (Test Product) data on Amazon Redshift (console)

- 1. Open and sign in to the Amazon Redshift console.
- 2. Choose **Clusters**, and choose your existing RA3 cluster.
- 3. Choose the **Datashares** tab.
- 4. Select the datashare you want to create the database from.
- 5. Under Subscriptions to AWS Data Exchange datashares, choose Create database from datashare.
- 6. In Create database from datashare, enter the Database name for your new database, and then choose Create.
- 7. Choose the **Marketplace** icon on the navigation pane, and open the **Query editor**.
- 8. Under **Resources**, select a database and a schema.

Worldwide Event Attendance 124 9. Run the following SQL query.

```
select * from database.schema.table
```

Querying Worldwide Event Attendance (Test Product) data on Amazon Redshift (SQL)

The following procedure shows how to set up and query the datashare using the SQL commands.

To query Worldwide Event Attendance (Test Product) data on Amazon Redshift (SQL)

1. To find the datashare, run the following command.

```
SHOW DATASHARES [ LIKE 'namepattern' ]
```

This command lists all datashares, including the one from Worldwide Event Attendance (Test Product), in addition to the provider's account_id and namespace. For more information, see Show Datashares in the Amazon Redshift Database Developer Guide.

2. Run the following command to create a database from the datashare.

```
CREATE DATABASE database_name
```

```
FROM DATASHARE datashare_name OF ACCOUNT account_id NAMESPACE namespace_guid
```

For more information, see Create Database in the Amazon Redshift Database Developer Guide.

3. Run the following SQL query.

```
select * from database.schema.table
```

AWS Data Exchange for AWS Lake Formation (Test Product) (Preview)

AWS Data Exchange for AWS Lake Formation (Test Product) is a free product that helps subscribers understand how to subscribe to and interact with an AWS Data Exchange product containing AWS Lake Formation data sets. You can use this product for testing purposes and learn how to query, analyze, and share data internally within minutes.

Topics

- Subscribing to AWS Data Exchange for AWS Lake Formation (Test Product) on AWS Data Exchange (Preview)
- Setting up and querying AWS Data Exchange for Lake Formation (Test Product) (Preview)

Subscribing to AWS Data Exchange for AWS Lake Formation (Test Product) on AWS Data Exchange (Preview)

The following procedure shows how to browse the AWS Data Exchange catalog to find and subscribe to AWS Data Exchange for Lake Formation (Test Product). (Preview)

To subscribe to AWS Data Exchange for Lake Formation (Preview)

- 1. Open and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Discover data products**, choose **Browse catalog**.
- 3. From the search bar, enter AWS Data Exchange for Lake Formation (Test product) and choose Enter.
- 4. Choose AWS Data Exchange for Lake Formation (Test Product) to view its details page.
 - a. (Optional) To view the data dictionary, scroll down to the product **Overview** section to see the data dictionary under **Data dictionaries**.
 - b. (Optional) To download the data dictionary, choose the **Data dictionary and samples** tab, choose the option button next to **Data dictionary**, and then choose **Download**.
 - c. (Optional) To preview the sample, choose the option button next to the sample name (AWS Data Exchange for Lake Formation (Test Product)), and then choose Preview sample (CSV only).
 - d. (Optional) To download the sample, choose the option button next to the sample name (AWS Data Exchange for Lake Formation (Test Product)), and then choose Download.
- 5. If you are previewing the sample, you can also choose **Download** in the sample preview dialog box.
- 6. In the top right corner, choose **Continue to subscribe**.
- 7. Choose your preferred price and duration combination and review the offer details, including the data subscription agreement.



Note

AWS Data Exchange for Lake Formation (Test Product) doesn't require subscription verification, but some products do. For more information, see Subscription verification for subscribers in AWS Data Exchange.

Review the pricing information, choose the pricing offer, and then choose **Subscribe**. 8.



Note

AWS Data Exchange for Lake Formation (Test product) is a free product. If you subscribe to a paid product, you are prompted to confirm your decision to subscribe.

It can take a few minutes for your subscription to become active after you choose **Subscribe**.

Navigating away from this page prior to your subscription becoming active will not prevent the subscription from processing.

Setting up and querying AWS Data Exchange for Lake Formation (Test Product) (Preview)

The following procedure shows how to set up and query a Lake Formation data permission set (Preview) using the AWS Management Console.

To enable querying on the AWS Data Exchange for Lake Formation (Test Product) data set (Preview)

- Open and sign in to the AWS Data Exchange console. 1.
- From the left navigation pane under My subscriptions, choose Entitled data. 2.
- 3. From the list of **Products**, choose **AWS Data Exchange for Lake Formation (Test Product)** (Preview).
- Choose **Accept** to accept the AWS RAM share.



Note

You must accept the AWS RAM share within 12 hours of subscribing to the data product. If your AWS RAM share invitation expires, select Request invitation and allow several business days for a new share to be sent. You only need to accept the AWS RAM share once for each provider that you license Lake Formation data sets from.

- 5. Open the Lake Formation console.
- 6. Sign in as a principal who has the Lake Formation CREATE_TABLE or CREATE_DATABASE permission, as well as the glue:CreateTable or glue:CreateDatabase AWS Identity and Access Management (IAM) permission.
- 7. In the navigation pane, choose **Tables**, and then choose **Create table**.
- 8. On the **Create table** page, choose **Resource Link**, and then provide the following information:
 - **Resource link name** Enter a name that adheres to the same rules as a table name. The name can be the same as the name of the target shared table.
 - **Database** The database in the local Data Catalog must contain the resource link.
 - Shared table Select one of the tables shared through AWS Data Exchange for Lake
 Formation (Test product). All of the table names shared through that product begin with
 adxlf_test, or enter a local (owned) or shared table name.

The list contains all of the tables shared with your account. The database and owner account ID are listed with each table. If you don't see a table that you know was shared with your account, check the following:

- If you aren't a data lake administrator, confirm with your administrator that you were granted Lake Formation permissions on the table.
- If you're a data lake administrator and your account is not the same AWS organization as
 the granting account, confirm that you've accepted the AWS Resource Access Manager
 (AWS RAM) resource share invitation for the table. For more information, see <u>Accepting a</u>
 resource share invitation from AWS RAM.
- Shared table's database If you selected a shared table from the list, this field is populated with the shared table's database in the external account. If you didn't select a shared table, enter a local database for a resource link to a local table, or the shared table's database in the external account.
- Shared table owner If you selected a shared table from the list, this field is populated with the shared table's owner account ID. If you didn't select a shared table, enter your AWS account ID for a resource link to a local table, or the ID of the AWS account that shared the table.

To query the AWS Data Exchange for Lake Formation (Test Product) data set (Preview) with Amazon Athena (Console)

- 1. Sign in to the Amazon Athena console with a role that has permissions for Amazon Athena.
- 2. In the Amazon Athena guery editor, choose the resource link that you created previously.
- 3. Choose the additional menu options icon next to source_data and choose Preview table.
- 4. Choose **Run query**.

To allow querying on the AWS Data Exchange for Lake Formation (Test Product) data set (Preview) (AWS CLI)

 To retrieve a list of all invitations available to your AWS account, enter the following command. The AWS CLI query parameter lets you restrict the output to only those invitations shared from AWS Data Exchange.

```
$ AWS ram get-resource-share-invitations
--region us-east-1
--query 'resourceShareInvitations[?
senderAccountId==147854383891]'
```

2. Find the invitations for the AWS Data Exchange for Lake Formation data set. Then, note the resourceShareInvitationArn in the output to use in the following command to accept the invitation.

```
$ AWS ram accept-resource-share-invitation --region us-east-1 --
resource-share-invitation-arn [resourceShareInvitationArn]
```

If successful, the response shows that the status has changed from **PENDING** to **ACCEPTED**.

3. Create a resource link to one of the tables shared through the AWS Data Exchange for Lake Formation data set with the following command:

```
aws glue create-table --database-name
[local_database_to_store_resource_link] --table-
input '{"Name":"resource_link_name","TargetTable":
{"CatalogId":"[account_owning_original_table]","DatabaseName":"[shared_db_in_
```



Note

To create resource links, use the Lake Formation CREATE_TABLE or CREATE_DATABASE permission, as well as the glue: CreateTable or glue:CreateDatabase IAM permission.

AWS Data Exchange for Amazon S3 (Test Product)

AWS Data Exchange for Amazon S3 (Test Product) is a product that helps subscribers understand how to subscribe to and interact with an AWS Data Exchange product. In this tutorial, the product contains Amazon Simple Storage Service (Amazon S3) data access data sets. You can use this product for testing purposes and to learn how to query and analyze data directly from a data provider's Amazon S3 bucket.

You can run queries to analyze the data in-place without setting up your own Amazon S3 buckets, copying data files into Amazon S3 buckets, or paying associated storage fees.

Topics

- Subscribing to AWS Data Exchange for Amazon S3 (Test Product)
- Setting up and querying AWS Data Exchange for Amazon S3 (Test Product)

Subscribing to AWS Data Exchange for Amazon S3 (Test Product)

The following procedure shows how to browse the AWS Data Exchange catalog to find and subscribe to AWS Data Exchange for Amazon S3 (Test Product).

To find and subscribe to AWS Data Exchange for Amazon S3 (Test Product)

- 1. Open and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Discover data products**, choose **Browse catalog**.
- From the search bar, enter AWS Data Exchange for Amazon S3 (Test Product) and choose Enter.
- Choose **AWS Data Exchange for Amazon S3 (Test Product)** to view its details page.
 - (Optional) To view the data dictionary, scroll down to the product **Overview** section to see the data dictionary under **Data dictionaries**.

- (Optional) To download the data dictionary, choose the **Data dictionary and samples** tab, choose the option button next to **Data dictionary**, and then choose **Download**.
- (Optional) To preview the sample, choose the option button next to the sample name of Blockchain Transactions (Test Data), and then choose Preview sample (CSV only).
- (Optional) To download the sample, choose the option button next to the sample name of **Blockchain Transactions (Test Data)**, and then choose **Download**.
 - If you're previewing the sample, you can also choose **Download** in the sample preview dialog box.
- 5. In the top right corner, choose **Continue to subscribe**.
- 6. Choose your preferred price and duration combination. Choose whether to enable autorenewal for the subscription, and review the offer details, including the data subscription agreement.



Note

AWS Data Exchange for Amazon S3 (Test Product) doesn't require subscription verification, but some products do. For more information, see Subscription verification for subscribers in AWS Data Exchange.

Review the pricing information, choose the pricing offer, and then choose **Subscribe**. 7.



Note

AWS Data Exchange for Amazon S3 (Test Product) is a free product. If you subscribe to a paid product, you are prompted to confirm your decision to subscribe.

It can take a few minutes for your subscription to become active after you choose **Subscribe**.

Navigating away from this page before your subscription becomes active will not prevent the subscription from processing.

Setting up and querying AWS Data Exchange for Amazon S3 (Test Product)

The following procedure shows how to set up and query an Amazon S3 data access data set using the AWS Command Line Interface (AWS CLI). Before querying, you must obtain the appropriate AWS Identity and Access Management (IAM) permissions to attach policies to your user. To access

data in a provider's bucket directly through the Amazon S3 delivery method, embed the following JSON policy to the user or role.

To set up AWS Data Exchange for Amazon S3 (Test Product)

- Open and sign in to the AWS Data Exchange console. 1.
- From the left navigation pane under My subscriptions, choose Entitled data in the AWS Region that hosts the data set. For the purposes of this tutorial, the Region is **us-east-1**.
- 3. From the list of **Products**, choose **AWS Data Exchange for Amazon S3 (Test Product)** and then choose the **Blockchain Transactions (Test Data)** data set.
- 4. Choose **Verify IAM permissions**.



Note

If you don't have the correct permissions, you'll receive a notification detailing how to create and attach the IAM policy to your user or role. In the following example, replace each user input placeholder with your own information.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": [
        "arn:aws:s3:us-east-1:111122223333:accesspoint/my-access-point",
        "arn:aws:s3:::aws-data-exchange-s3-data-access-btc-demo-us-east-1"
      ]
   },
      "Effect": "Allow",
      "Action": "s3:GetObject",
      "Resource": [
        "arn:aws:s3:us-east-1:111122223333:accesspoint/my-access-point/
object/*",
        "arn:aws:s3:::aws-data-exchange-s3-data-access-btc-demo-us-east-1/*"
```

```
]
]
]
}
```

To allow querying on the AWS Data Exchange for Amazon S3 (Test Product) data access data set using the AWS CLI

- Open AWS CloudShell in us-east-1.
- 2. Choose the copy button next to the access point alias to copy and paste the code inside. After the command is added in AWS CloudShell with the correct access point alias, you can see the list of Amazon S3 objects included in this product.

Important

When a provider has enabled Requester Pays, the subscriber pays for the data transfer and the request. The provider pays for the data storage. For more information, see Using Requester Pays buckets for storage transfers and usage in the Amazon Simple Storage Service User Guide.

3. (Optional) You can also copy an object to your local system using the following command.

```
aws s3api get-object --bucket <Access point alias> --key 'v1.0/ btc/transactions/date=2022-11-27/part-00000-03a88dba-27dd-4f59-a890-70a3d2c7ad26-c000.snappy.parquet' AWS_btc.snappy.parquet --request-payer requester
```

AWS Data Exchange Provider-Generated Notifications (Test Product)

AWS Data Exchange Provider-Generated Notifications (Test Product) is a free product that helps subscribers understand how to subscribe to and interact with an AWS Data Exchange product using provider-generated notifications.

Providers use this feature to notify you of important events related to their data sets. You'll receive these events in a consistent, structured format using Amazon EventBridge, that you can use to build automated workflows. Provider-generated notifications also supports the delivery of human-readable notification to emails and chat programs using <u>AWS User Notifications</u>.

Topics

- Subscribing to AWS Data Exchange for Provider-Generated Notifications (Test Product)
- Configuring AWS Data Exchange provider-generated notifications using Amazon EventBridge

Subscribing to AWS Data Exchange for Provider-Generated Notifications (Test **Product)**

The following procedure shows how to subscribe to AWS Data Exchange Provider-Generated Notifications (Test Product).

To find and subscribe to AWS Data Exchange for Provider-Generated Notifications (Test **Product)**

- 1. Open and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Discover data products**, choose **Browse catalog**.
- 3. From the search bar, enter AWS Data Exchange Provider-Generated Notifications (Test Product) and choose Enter.
- Choose AWS Data Exchange for Provider-Generated Notifications (Test Product) to view its details page.
- In the top right corner, choose **Continue to subscribe**.
- Choose your preferred price and duration combination. Choose whether to enable autorenewal for the subscription, and review the offer details, including the data subscription agreement.



Note

AWS Data Exchange for Provider-Generated Notifications (Test Product) doesn't require subscription verification, but some products do. For more information, see Subscription verification for subscribers in AWS Data Exchange.

Review the pricing information, choose the pricing offer, and then choose **Subscribe**. 7.



Note

AWS Data Exchange for Provider-Generated Notifications (Test Product) is a free product. If you subscribe to a paid product, you are prompted to confirm your decision to subscribe.

It can take a few minutes for your subscription to become active after you choose **Subscribe**.

Navigating away from this page before your subscription becomes active will not prevent the subscription from processing.

Configuring AWS Data Exchange provider-generated notifications using Amazon **EventBridge**

AWS Data Exchange delivers provider-generated notifications using Amazon EventBridge. Your role must be able to create Amazon EventBridge rules, a target, and must be able to subscribe to a data product.

AWS Data Exchange events are published to your default Amazon EventBridge event bus in the same AWS Region as where your data set is located. Use the following steps to create an Amazon EventBridge rule for provider-generated notifications:

To create an Amazon EventBridge rule for provider-generated notifications

- Create a target for the Amazon EventBridge rule. For a simple Lambda function in Python do the following:
 - Navigate to the AWS Lambda console. a.
 - Choose **Create function** and select **Author from scratch**. b.
 - Provide a function name and select **Python 3.10** as the **runtime**. Choose **Create function**. c.
 - d. Enter the following code for **lambda_function.py**:

```
import json
def lambda_handler(event, context):
   print(" ".join(["Event of type", event["detail-type"], "received!"]))
    print(" ".join(["Details", json.dumps(event["detail"])]))
```

```
return {"statusCode": 200, "body": json.dumps("Hello from Lambda!")
}
```

- 2. Navigate to the EventBridge console.
- 3. Navigate to the **Rules** and select the default event bus.
- 4. Choose **Create rule** and provide the **Name** and optional **Description**. Make sure the **Rule** type is **Rule with an event pattern**.
- Choose Next.
 - Make sure the Event source is AWS events or EventBridge partner events. Under Creation method, select Custom pattern (JSON editor). Under Event pattern, enter the following JSON:

```
{
    "source": ["aws.dataexchange"],
    "detail-type": ["Data Set Update Delayed", "Data Updated in Data Set",
    "Deprecation Planned for Data Set", "Schema Change Planned for Data Set"]
}
```

- 6. Choose **Next**.
 - a. For **Target 1**, select **AWS service** and choose **Lambda function**.
 - b. For the **function**, select the function created in Step 1. Complete the creation of the rule.

This Lambda function will be triggered any time a provider-generated notification is delivered. From the **Monitor** tab in the Lambda console, you can view recent invocations of the function.

Providing AWS Data Exchange data products on AWS Marketplace

At a high level, this is how to list AWS Data Exchange data products on AWS Marketplace:

- 1. **Potential provider registers to be a provider** Registering allows you to list products on AWS Data Exchange and make them available on AWS Marketplace. For more information, see Step 2: Register to be a provider.
- 2. The data is eligible to be published on AWS Data Exchange You're limited to distributing data sets that meet the legal eligibility requirements set forth in the Terms and Conditions for AWS Marketplace Sellers. For more information about the types of permitted data, see Publishing guidelines for AWS Data Exchange.
- 3. **Provider creates a data set, a revision, and imports assets** You can create data sets through the AWS Data Exchange console or API. Then, you can create revisions in the data set, and add assets to that revision. For more information, see Data in AWS Data Exchange.
- 4. **Provider creates a product and its offer** To create a product, you must provide product details, include one or more data sets, and optionally provide public offer details. For more information, see Publishing a new product in AWS Data Exchange.
 - **Products containing Files (Amazon S3 Objects)** When an owned data set containing Amazon S3 objects is published in a product, AWS Data Exchange creates a copy of the data set. Subscribers can access that copy of the data set as an entitled data set.
 - Products containing Amazon API Gateway APIs When an owned data set containing
 Amazon API Gateway APIs is published in a product, AWS Data Exchange allows requests
 to the AWS Data Exchange endpoint to proxy through to your Amazon API Gateway API.
 Subscribers can view the API and download the API specification as an entitled data set.
 Subscribers can also call the API through the AWS Data Exchange console.
 - Products containing Amazon Redshift data sets When an owned data set containing
 Amazon Redshift data sets is published in a product, AWS Data Exchange allows requests
 to the AWS Data Exchange endpoint to proxy through to your Amazon Redshift datashare.
 Subscribers can have read-only access to the tables, views, schemas, and user-defined
 functions that you've added to the datashare.
 - **Products containing Amazon S3 data access** When an owned data set containing Amazon S3 data access is published in a product, AWS Data Exchange allows subscribers the same Amazon S3 objects that the data provider maintains. This approach provides the most current

data available. Providers share direct access to an Amazon S3 bucket or specific prefix and Amazon S3 objects and use AWS Data Exchange to manage subscriptions, entitlements, billing, and payment.

- Products containing AWS Lake Formation data sets (Preview) When an owned data set
 containing Lake Formation data permission data sets is published in a product, AWS Data
 Exchange grants read-only access to the data associated with the LF-tags you included in
 the data set. Subscribers can subscribe to the databases, tables, or columns that you share
 with them with downstream query services integrated with Lake Formation such as Amazon
 Athena and Redshift Spectrum.
- 5. **(Optional) Provider enables subscription verification** If you enable subscription verification, subscribers must request a subscription to your product. This gives you an opportunity to review potential subscribers before they access your data sets. For more information, see <u>Subscription</u> verification for providers in AWS Data Exchange.
- 6. **(Optional) Provider creates custom offers for the product** In addition to a public offer, you can create custom offers, including private and Bring Your Own Subscription (BYOS) offers, for select customers. For more information, see Creating an offer for AWS Data Exchange products.
- 7. **(Optional) Provider publishes new revision** You can update dynamic data sets over time by creating a new revision using the AWS Data Exchange API or console. These revisions can then be published. For more information, see Revisions or Updating products in AWS Data Exchange.
- 8. Provider reviews reports through the AWS Marketplace Management Portal Reports are available to all registered AWS Marketplace sellers and are released on a regular cadence (daily, weekly, or monthly). For more information, see AWS Data Exchange provider financials on AWS Marketplace.
- 9. **Provider receives funds distributed by AWS Marketplace** For more information, see <u>AWS Data Exchange provider financials on AWS Marketplace</u>.

Extended Provider Program (EPP)

The Extended Provider Program (EPP) is a program for qualified data providers to publish data products containing sensitive categories of personal information and/or personal information that is not otherwise publicly available.

Providers seeking to participate in the EPP must complete an additional review process by the AWS Data Exchange team. For more information about eligibility for the Extended Provider Program, contact Support.

For more information about publishing guidelines for data providers who are enrolled in the EPP, see Publishing guidelines for AWS Data Exchange.

Programmatic access

If you're using AWS Data Exchange programmatically, there are two different sets of resources with two different APIs:

- AWS Data Exchange API Use these API operations to create, view, update, and delete data sets and revisions. You can also use these API operations to import and export assets to and from those revisions. For more information, see the AWS Data Exchange API Reference.
- AWS Marketplace Catalog API Used by providers to view and update products on AWS Data Exchange and AWS Marketplace. For more information, see the <u>AWS Marketplace Catalog API</u> Reference.

Before you become a data product provider on AWS Data Exchange, review the following topic:

Setting up AWS Data Exchange

After you review this topic, you're ready to get started.

Related topics

- Publishing guidelines for AWS Data Exchange
- Product best practices in AWS Data Exchange
- Getting started as a provider in AWS Data Exchange
- Publishing a new product in AWS Data Exchange
- Product description templates in AWS Data Exchange
- Updating products in AWS Data Exchange
- Creating an offer for AWS Data Exchange products
- Provider-generated notifications in AWS Data Exchange
- Data in AWS Data Exchange

Programmatic access 139

Getting started as a provider in AWS Data Exchange

The following topics describe the complete process of becoming a data product provider on AWS Data Exchange using the AWS Data Exchange console. The process has the following steps:

Steps

- Step 1: Confirm your eligibility
- Step 2: Register to be a provider
- Step 3: Confirm eligibility of your data

Step 1: Confirm your eligibility

Before you can register, you must meet the following requirements to confirm your eligibility.

Requirements for publishing data products

Whether you charge for your AWS Data Exchange data product, you're selling that product on AWS Marketplace. To create and offer data products, you must:

- Have a defined customer support process and support organization.
- Provide a means to keep data regularly updated and free of vulnerabilities.
- Follow best practices and guidelines when marketing your product.
- Be an AWS customer in good standing and meet the requirements in the terms and conditions for AWS Marketplace sellers and for AWS Data Exchange providers.
- Be a permanent resident or citizen in an <u>eligible jurisdiction</u>, or a business entity organized or incorporated in one of those areas.
- To provide data products, you must also request on-boarding through the <u>Create case</u> wizard for Support. The AWS Data Exchange team will contact you to complete the qualification and registration process.

Additionally, if you want to offer products and charge for them, you must provide the following information:

• You must provide tax and bank account information. For US-based entities, a W-9 form and a banking account from a US-based bank are required.

Getting started as a provider 140

• Non-US sellers are required to provide a W-8 form, value-added tax (VAT) or goods and services tax (GST) registration number, and US bank information. If you don't have a US bank account, you can register for a virtual US bank account from Hyperwallet.

Eligible jurisdictions for AWS Data Exchange products

To provide data products on AWS Data Exchange, you must be a permanent resident or citizen in one of the following countries or SARs, or a business entity organized or incorporated therein:

- Australia¹
- Bahrain¹²
- European Union (EU) member state¹
- Hong Kong SAR
- Israel¹²
- Japan²³
- New Zealand¹
- Norway¹²
- Qatar
- Switzerland¹²
- United Arab Emirates (UAE)12
- United Kingdom (UK)¹
- United States (US)

For more information about VAT, invoicing, and your tax obligations as a provider, see <u>AWS</u> Marketplace Sellers on Amazon Web Service Tax Help.

¹ Providers of paid products in these countries must provide VAT registration information in country of establishment.

² If you, as a provider, are located in the same country as the subscriber, you may be responsible for tax invoicing, collections, and remittances. Please consult with your tax advisor.

³ Providers based in Japan have an obligation to self-account for the Japan Consumption Tax (JCT) on the listing fee charges.

Step 2: Register to be a provider

To use AWS Data Exchange as a provider, you must be a registered seller on AWS Marketplace and be qualified by the AWS Data Exchange team. When you register an account as an AWS Marketplace seller, the account is the seller of record for your products and is used for reporting and disbursement. All products and their public offers are discoverable on AWS Data Exchange and AWS Marketplace.

If your AWS Data Exchange qualification and registration process is complete and you want to upgrade from publishing free products to paid products, contact the AWS Marketplace Seller Operations team.

Important

You can't change the AWS account that you use to list a product on AWS Marketplace. Only data sets owned by that account can be included in products published by that account. Only AWS accounts that are registered to provide data products on AWS Marketplace and AWS Data Exchange can publish products.

To register as a provider for AWS Data Exchange and AWS Marketplace

- 1. From your web browser, open the AWS Marketplace Management Portal.
- 2. Choose **Register now** to open the registration wizard.
- 3. Confirm your company or full name, and review the Terms and Conditions. If you agree to them, choose I have read and agree to these terms.
- 4. On the **Account Settings** page, choose **Add** to add a public profile.
- 5. (Optional) If you want to submit paid products to AWS Marketplace or AWS Data Exchange, you must provide your tax and banking information. On the **Account Settings** page, from the **Provide tax and banking information** tab, choose **Start** to complete the tax and banking wizard. This submits your tax and banking information in the AWS Marketplace Management Portal.



Note

We strongly recommend that you sign and submit the tax form electronically. Otherwise, you must print, complete the signature section, and mail a hard copy of the tax form to the address provided in the tax information interview. This delays the registration process.

6. In addition to being a registered AWS Marketplace seller, you must submit an AWS Data Exchange qualification request. Create a case with the <u>AWS Marketplace Seller Operations</u> team. They will contact you to complete the qualification and registration process.

Step 3: Confirm eligibility of your data

To confirm the eligibility of your data, review the Publishing guidelines for AWS Data Exchange.

If you have questions about the eligibility of your data set, contact the <u>AWS Marketplace Seller</u> Operations team.

You can create your product after you've reviewed the publishing guidelines for data products on AWS Data Exchange, and you've confirmed that your data set can be listed.

Publishing guidelines for AWS Data Exchange

The following guidelines outline restrictions for listing products on AWS Data Exchange. As a provider, you're responsible for complying with these guidelines and the <u>Terms and Conditions for AWS Marketplace Sellers</u> and the <u>AWS Customer Agreement</u>. AWS may update these guidelines from time to time. AWS removes any product that breaches these guidelines and may suspend the provider from future use of the service.

In addition to accepting and following the guidelines under the Terms and Conditions for AWS Marketplace Sellers, providers must abide by the following publishing guidelines for data products.

AWS Data Exchange publishing guidelines for data products

- Your data products may not contain any illegal content, viruses, malware, or any other material that is harmful to others.
- 2. Your data products may not include any information that can be used to trace or associate a device or an identifiable person with a *Sensitive Location*. A *Sensitive Location* includes the following: any location offering cancer treatment, treatment for HIV/AIDS, fertility or abortion clinics, mental health treatment facilities, and emergency room trauma centers; places of religious worship; correctional facilities; dependency or addiction treatment centers; domestic abuse or rape crisis centers; places that may be used to infer an LGBTQ+ identification or

other sexual orientation; military bases; temporary places of assembly such as political rallies, marches, or protests, during the times that such rallies, marches or protests take place; places primarily intended to be occupied by children under 16; places that may be used to infer engagement with explicit sexual content, material, or acts; places that may be used to infer refugee or immigrant status, such as refugee or immigration centers and immigration services; welfare or homeless shelters; halfway houses, credit repair, debt services, bankruptcy services, or payday lending institutions.

In addition, unless you're a qualified data provider under AWS Data Exchange's the section called "Extended Provider Program (EPP)", your data products may not include information that can be used to identify any person, unless that information is *Publicly Available Information*. *Publicly Available Information* means information: (1) that is lawfully made available through federal, state, local government records, open court records, or public company filings; or (2) that is lawfully made available to the general public by the data subject.

- 3. The following categories of information must be aggregated or anonymized so that no person in your data product can be identified: biometric or genetic data, health, racial or ethnic origin, political opinions, religious or philosophical beliefs, sex or sexual orientation, trade union membership, personal payment or financial information (for example, credit history), Sensitive Locations, or other similar categories of sensitive information.
 - Some examples of data sets that can be included on AWS Data Exchange Historic stock prices for public companies, names of judges and their court opinions, and aggregated or anonymized research findings from pharmaceutical drug studies.
 - For HCLS use-cases, data that has been de-identified through Expert Determination or Safe Harbor methods in compliance with HIPAA de-identification guidelines.
 - Some examples of data sets that are prohibited on AWS Data Exchange Lists of names organized by race, geo-location data that can be used to identify a person, and protected health information under the Health Insurance Portability and Accountability Act of 1996 (HIPAA).
- 4. You should carefully consider how subscribers may and may not use your data products, and you should clearly include this information in your Data Subscription Agreement (DSA).
- 5. Product listing descriptions must be accurate, contain valid contact information, and note if any data has been aggregated or anonymized.
- 6. You may not use AWS Data Exchange to provide or otherwise make accessible or available Bulk U.S. Sensitive Personal Data, including, deidentified, key coded, or anonymized data, or U.S. Government-related Data to Countries of Concern or Covered Persons as each is defined in the

Publishing guidelines 144

- U.S. Department of Justice Final Rule implementing Executive Order 14117, Preventing Access to Americans' Bulk Sensitive Personal Data and United States Government-Related Data by Countries of Concern, available here.
- 7. You may not use AWS Data Exchange to promote any other products or solutions not listed on AWS Marketplace, except for products or solutions that are not compatible with AWS Marketplace.
- 8. You are limited to distributing data sets that meet the legal eligibility requirements set forth in the Terms and Conditions for AWS Marketplace Sellers. If you breach these terms in any way, the prohibited product is removed from AWS Data Exchange and you might be suspended from the service.
- 9. If you're listing an API data set in a product:
 - You must first integrate your API with Amazon API Gateway. For more information about how
 to integrate your REST API with API Gateway, see <u>Working with REST APIs</u> in the *API Gateway*Developer Guide.
 - You must respond to support-related questions from subscribers about the data product in 1 business day. Not following this guideline may result in your products being removed from AWS Data Exchange.

10Logos, DSAs, and other attachments added to your product might be stored separately from where your actual data products sits.

Note

Providers who are enrolled in the Extended Provider Program are subject to the restrictions set forth in the Extended Provider Program Addendum to the Terms and Conditions for AWS Marketplace Providers which are supplemental to guidelines 2 and 3 above. For more information, see Extended Provider Program (EPP).

If you have questions about the eligibility of your data set:

Contact the AWS Marketplace Seller Operations team.

After you've reviewed the publishing guidelines for data products on AWS Data Exchange, and you've confirmed that your data set can be listed, you can create your product.

Publishing guidelines 145

Publishing a new product in AWS Data Exchange

The following topics describe the process of publishing a new product on AWS Data Exchange by using the AWS Data Exchange console.

Topics

- Publishing a product in AWS Data Exchange containing file-based data
- Publishing a product in AWS Data Exchange containing APIs
- Publishing a product in AWS Data Exchange containing Amazon Redshift data sets
- Publishing a product in AWS Data Exchange containing Amazon S3 data access
- <u>Publishing a product in AWS Data Exchange containing AWS Lake Formation data permission</u> <u>data sets (Preview)</u>

The following video explains more about how to publish a new data product on AWS Data Exchange.

Publishing a product in AWS Data Exchange containing file-based data

The following topics describe the process of creating a data set and publishing a new product in AWS Data Exchange containing file-based data on AWS Data Exchange by using the AWS Data Exchange console. The process has the following steps:

Steps

- Step 1: Create assets
- Step 2: Create a data set
- Step 3: Create a revision
- Step 4: Import assets to a revision
- Step 5: Publish a new product
- Step 6: (Optional) Copy a product

Step 1: Create assets

Assets are the data in AWS Data Exchange. For more information, see Assets.

Before you create and publish a new file-based data product, you must:

Publishing a new product 146

1. Create your files.

AWS Data Exchange supports all file types.

2. Store your files as objects in Amazon Simple Storage Service (Amazon S3) or on your local computer.

For more information about storing files in Amazon S3, see the Amazon S3 User Guide.

Step 2: Create a data set

Data sets in AWS Data Exchange are dynamic and are versioned using revisions, with each revision containing at least one asset. For more information, see Data in AWS Data Exchange.

To create a data set

- 1. Open your web browser and sign in to the <u>AWS Data Exchange console</u>.
- 2. In the left side navigation pane, under **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- 4. In Select data set type, choose Files.
- 5. In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see Data set best practices.
- 6. (Optional) Under Add tags optional, add tags.
- 7. Choose **Create data set**.

Step 3: Create a revision

In the following procedure, you create a revision after you've created a data set in the AWS Data Exchange console. For more information, see Revisions.

To create a revision

- 1. On the **Data set overview** section of the data set details page:
 - a. (Optional) Choose **Edit name** to edit information about your data set.
 - b. (Optional) Choose **Delete** to delete the data set.
- 2. In the **Revisions** section, choose **Create revision**.

- 3. Under **Define revision**, provide an optional comment for your revision that describes the purpose of the revision.
- 4. (Optional) Under **Add tags optional**, add tags associated with the resource.
- 5. Choose Create revision.
- 6. Review, edit, or delete your changes from the previous step.

Step 4: Import assets to a revision

In the following procedure, you import data assets, and then finalize the revision in the AWS Data Exchange console. For more information, see Assets.

To import assets to the revision

- Under the Jobs section of the data set details page, choose either Import from Amazon S3 or Upload (to upload from your computer), depending on where the data assets for the data set are currently stored.
- 2. Follow the prompts, depending on your selection. A job is started to import your asset into your data set.
- 3. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed.**
- 4. If you have more data to add, repeat Step 1.
- 5. In **Revision overview**, review your revision and its assets.
- 6. Choose Finalize revision.

You have successfully finalized a revision for a data set.

You can edit or delete a revision before you add it to a product.

Topics

- · Edit a revision
- · Delete a revision

Edit a revision

To edit the revision after you've finalized it

1. In Revision overview, choose De-finalize.

You see a message that the revision is no longer in the finalized state.

- 2. To edit the revision, from **Revision overview**, choose **Actions**, **Edit**.
- 3. Make your changes, and then choose **Update**.
- 4. Review your changes, and then choose **Finalize**.

Delete a revision

To delete the revision after you've finalized it

- In Revision overview, choose Delete.
- 2. Type **Delete** in the **Delete revision** dialog box, and then choose **Delete**.

Marning

This deletes the revision and all of its assets. This action cannot be undone.

Step 5: Publish a new product

After you've created at least one data set and finalized a revision with assets, you're ready to publish that data set as a part of a product. For more information, see Product best practices in AWS Data Exchange. Make sure that you have all required details about your product and offer.

To publish a new product

- In the left navigation pane of the <u>AWS Data Exchange console</u>, under <u>Publish data</u>, choose <u>Products</u>.
- 2. From **Products**, choose **Publish new product** to open the **Publish new product** wizard.
- 3. In the **Product visibility** section, choose your product's **Product visibility options** and **Sensitive information** configuration, and then choose **Next**. For more information, see Product visibility in AWS Data Exchange and Sensitive categories of information in AWS Data Exchange.
- 4. In the **Add data** section, under **Owned data sets**, select the check boxes next to the data sets you want to add, and then choose **Add selected**.



Note

The data sets you choose must have a finalized revision. Data sets without finalized revisions can't be added.

- Go to **Selected data sets** to review your selections.
 - You can review the **Name** of the data set, the **Type** of data set, and the timestamp of when the data set was **Last updated**.
- Go to **Select revision access rules**, choose the revision access rules that you want to set for data sets included in this product, and then choose **Next**. For more details, see Revision access rules in AWS Data Exchange.
- In the **Define product** section, under **Product overview**, enter information about your product, including the Product name, Product logo, Support contact information, and Product categories.

For more information, see Product best practices in AWS Data Exchange.

- (Optional) In the **Define product** section, under **Data dictionaries and samples optional**, choose a data set by selecting the option button next to the data set name and then choose Edit.
 - In the **Edit** dialog box, under **Upload data dictionary**, choose **Add file** to upload a new data dictionary.
 - You can choose one data dictionary, in .csv format, with a maximum size of 1 MB.
 - Choose a saved data dictionary from your computer, and then choose **Open**.

The data dictionary .csv file appears on the **Edit** dialog box.



Note

Your data dictionary must conform to the AWS Data Exchange data dictionary template. If you don't have a saved data dictionary to upload, you can choose either the blank data dictionary template link or the example data dictionary link in the AWS Data Exchange console.

- Choose **Data dictionary preview** to preview it. c.
- d. Under Samples - optional, choose Upload samples, choose a sample from your computer, and then choose **Open**.

The samples appear on the **Edit** dialog box.



Note

You can upload up to 10 samples with a maximum size of 50 MB. Samples in .csv format can be previewed.

- Enter a description for each sample that will be visible on the product detail page.
- f. Choose Save.
- Under **Product definition**, enter a **Short description** and a **Long description** of your product.

If you want to use a template for your long description, select **Apply template**, choose your template type, and then fill out the template with your specific product details.

- Choose Next. 8.
- Configure your offer. 9.
 - If you're creating a public offer, in the **Add public offer** section, configure your offer. All AWS Data Exchange products with visibility set to **Public** require a public offer.
 - 1. Choose your **Pricing and access duration** options for the subscription.
 - 2. Choose your US sales tax settings, data subscription agreement (DSA), and refund policy.
 - 3. (Optional) Set **Subscription verification**, which enables you to control who can subscribe to this product. For more information, see Subscription verification for providers in AWS Data Exchange.
 - 4. Choose your Offer auto-renewal option. For more information, see Creating an offer for AWS Data Exchange products.
 - Choose Next.
 - If you're creating a private offer, configure the offer details in the **Add custom offer** section.
 - 1. In the **Subscriber account information** section, add at least one subscriber account to which you want to extend the offer.

- 3. Choose the **Offer expiration date** by which the subscriber must accept the offer.
- 4. Choose your US sales tax settings, data subscription agreement (DSA), and refund policy.
- 5. Choose your **Offer auto-renewal** option. For more information, see Creating an offer for AWS Data Exchange products.
- 6. Choose **Next**.
- 10. In the Review & publish section, review your product information and then expand the **Product page preview** to see how it will look after it's published.
- 11. If you're sure that you want to make the product and public offer visible and available to everyone, choose **Publish**.

You've now completed the manual portion of publishing a data product with a public offer. AWS Data Exchange prepares and publishes your product. On the **Product overview** page, the status of your product is **Awaiting approval** and then changes to **Published** after it's published.

Step 6: (Optional) Copy a product

After you have created your first product, you can copy its details and public offers to create a new product.



Note

You can copy a public, private, published, or unpublished product. Custom offers associated with the product will not be copied, but public offers will be copied.

To copy a product

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- From **Products**, choose the button next to the product you want to copy. 3.
- 4. Select the **Actions** dropdown, and then choose **Create copy**.
- Continue through the **Publish a new product** workflow, with details already filled in, based on 5. the product you chose in Step 3. For more information, see Step 5: Publish a new product.

Publishing a product in AWS Data Exchange containing APIs

The following topics describe the process of creating a REST API data set and publishing a new product that contains APIs on AWS Data Exchange. You can complete the process by using either the AWS Data Exchange console or the AWS Command Line Interface.

After you have set up your Amazon API Gateway REST API, you can create a new API data set in AWS Data Exchange. You can then create a revision, and add API assets.

Creating and publishing an API asset allows subscriber requests to an AWS Data Exchange endpoint to proxy through to your API Gateway API. You can then add this data set to a product and add pricing. Then, subscribers can view your product and subscribe to it in the AWS Marketplace catalog and the AWS Data Exchange catalog.

AWS Data Exchange features are available including revision access rules, private products, private offers, and subscription verification.

You can choose only contract-based pricing, metered cost pricing (where the contract pricing is \$0), or a combination of metered and contract pricing.

You can choose standard metered costs, or you can specify a custom metered cost. There are three types of standard metered costs available:

- Per API request
- Per successful API request
- Per unit of data transferred in bytes



Note

Metered costs apply to all API data sets in a product. Therefore, if you want to charge different prices for the same dimension for different API data sets, we recommend that you create these data sets in different products.

The process has the following steps:

Steps

Prerequisites

- Step 1: Update the API resource policy
- Step 2: Create an API data set
- Step 3: Create a revision
- Step 4: Add API assets to a revision
- Step 5: Publish a new product containing APIs
- Step 6: (Optional) Copy a product

Prerequisites

Before you can publish a product containing APIs, you must meet the following prerequisites:

- Before you can use any AWS service, including AWS Data Exchange, you must sign up for AWS and create an administrative user. For more information, see Getting started in the AWS IAM Identity Center User Guide.
- To create products on AWS Data Exchange, you must register your AWS account as an AWS Marketplace Seller. Use this account to create your data sets. The account with the API Gateway resource doesn't need to be in the same account that is creating the data sets.
- Your REST API must be on Amazon API Gateway with an integration that uses an appropriate request and response model for accessing your data, such as Amazon DynamoDB or AWS Lambda. For more information, see Developing a REST API in API Gateway and Working with REST APIs in the Amazon API Gateway Developer Guide.



Note

Only public API Gateway APIs are supported.

 Your API Gateway REST API must be able to authenticate and authorize calls from the AWS Data Exchange service principal. Every request from AWS Data Exchange to your API uses the Signature Version 4 (SigV4) protocol signed with AWS Data Exchange credentials. AWS Data Exchange works with custom domains and domain key mappings.



Note

AWS Data Exchange doesn't support Amazon Cognito, No-Auth, and AWS Lambda authorizers.

- If your API Gateway REST API uses a custom identity system for authentication and authorization, configure it to use IAM authentication and import an OpenAPI schema describing your API. AWS Data Exchange will invoke your API Gateway REST API with its own service credentials and include subscriber information such as account ID.
- Your API Gateway REST API is responsible for integrating with your backend. To do this, do one
 of the following:
 - Attach a long-lived authentication token to every request that comes through your API Gateway REST API that the backend can verify.
 - Use API Gateway to invoke a Lambda function that can generate credentials and invoke your API.

Your API is invoked per the API integration request specification.

For more information, see the following topics:

Topics

- API data set security
- API integration request specification
- Header forwarding

API data set security

AWS Data Exchange encrypts traffic end to end using Transport Layer Security (TLS) 1.2. All metadata is encrypted at rest. AWS Data Exchange will not store subscriber requests or the responses from your backend. We only extract metering metadata necessary for billing.

API integration request specification

An API on AWS Data Exchange passes through all headers (except for the headers listed in <u>Header forwarding</u>), body, http method, path, and query strings as-is from the customer request and appends the following headers.

```
// These headers help prevent Confused Deputy attacks. They enable the SourceAccount
// and SourceArn variables in IAM policies.
'x-amz-source-account': ACCOUNT_ID,
'x-amz-source-arn': `arn:aws:dataexchange:${REGION}:${OWNER_ACCOUNT_ID}:data-sets/
${DATA_SET_ID}/revisions/${REVISION_ID}/assets/${ASSET_ID}`,
```

```
// These headers identify the API Asset in Data Exchange.
'x-amzn-dataexchange-asset-id': ASSET_ID,
'x-amzn-dataexchange-data-set-id': DATA_SET_ID,
'x-amzn-dataexchange-revision-id': REVISION_ID,

// This header identifies the Data Exchange Product.
'x-amzn-dataexchange-product-id': PRODUCT_ID,

// This header identifies the caller of Data Exchange. It will contain subscriber
// information.
'x-amzn-dataexchange-requester-account-id': REQUESTER_ACCOUNT_ID,

// Providers can attach custom metadata in the form of key/value pairs
// to a particular subscription. We will send these key/value pairs as stringified
// JSON.
'x-amz-dataexchange-subscription-metadata': STRINGIFIED_METADATA,
```

Header forwarding

AWS Data Exchange removes any headers related to authentication or namespaced to Amazon prior to forwarding it to a provider backend. Specifically, AWS Data Exchange removes:

- Authentication header
- Any headers that begin with x-amz

The host header will be overwritten as a consequence of the proxying.

Step 1: Update the API resource policy

If you have an Amazon API Gateway REST API that meets the <u>Prerequisites</u>, you must update your API resource policy to grant AWS Data Exchange the ability to invoke your API when a subscriber makes a request to get your API's schema.

To update your API resource policy

1. Add the following policy to your API's resource policy:

```
{
"Effect": "Allow",
"Principal": {"Service": "dataexchange.amazonaws.com"},
"Action": "execute-api:Invoke",
```

```
"Resource": "*",
"Condition": {"StringEquals": {"aws:SourceAccount": "<account-id>"}}
}
```

2. Replace account-id with the account that will be creating the API data set.

The account with the API Gateway resource does not need to be in the same account that is creating the data set.

This policy restricts these permissions to calls made by the AWS Data Exchange service principal and requires that only your account can authorize AWS Data Exchange to integrate with your API.



If you have a resource policy that explicitly denies AWS Data Exchange from doing this invocation, you must remove or limit this deny.

You're now ready to create an API data set.

Step 2: Create an API data set

Data sets in AWS Data Exchange are dynamic and are versioned using revisions, with each revision containing at least one asset. For more information, see Data in AWS Data Exchange.

You use either the AWS Data Exchange console or the AWS Command Line Interface to create an API data set:

- Creating an API data set (console)
- Creating an API data set (AWS CLI)

Creating an API data set (console)

To create an API data set (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. On the left side navigation pane, under **My data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- In Select data set type, choose Amazon API Gateway API.

- 5. In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see Data set best practices.
- 6. (Optional) Under **Add tags optional**, add tags.
- 7. Choose Create.

You are now ready to create a revision.

Creating an API data set (AWS CLI)

To create an API data set (CLI)

1. Use the create-data-set command to create an API data set:

```
$ AWS dataexchange create-data-set \
--asset-type API_GATEWAY_API \
--description 'Data Set Description' \
--name 'Data Set Name'

{
    "Arn": "arn:aws:dataexchange:us-east-1:123456789012:data-sets/$DATA_SET_ID",
    "AssetType": "API_GATEWAY_API",
    "CreatedAt": "2021-09-11T00:16:46.349000+00:00",
    "Description": "Data Set Description",
    "Id": "$DATA_SET_ID",
    "Name": "Data Set Name",
    "Origin": "OWNED",
    "UpdatedAt": "2021-09-11T00:16:46.349000+00:00"
}
```

2. Note the new Asset Type of API_GATEWAY_API.

You are now ready to create a revision.

Step 3: Create a revision

In the following procedure, you create a revision after you've created a data set. For more information, see Revisions.

You use either the AWS Data Exchange console or the AWS Command Line Interface to create a revision:

- Creating a revision (console)
- Creating a revision (AWS CLI)

Creating a revision (console)

To create a revision (console)

- 1. On the **Data set overview** section of the data set details page:
 - a. (Optional) Choose **Edit name** to edit information about your data set.
 - b. (Optional) Choose **Delete** to delete the data set.
- 2. On the **Revisions** section, choose **Create revision**.
- 3. Under **Define revision**, provide an optional comment for your revision that describes the purpose of the revision.
- 4. (Optional) Under **Add tags optional**, add tags associated with the resource.
- 5. Choose Create revision.
- 6. Review, edit, or delete your changes from the previous step.

You are now ready to add API assets to the revision.

Creating a revision (AWS CLI)

To create a revision (AWS CLI)

Use the create-revision command to create a revision:

```
$ AWS dataexchange create-revision \
--data-set-id $DATA_SET_ID \
--comment 'First Atlas Revision'
{
   "Arn": "arn:aws:dataexchange:us-east-1:123456789012:data-sets/$DATA_SET_ID/
   revisions/$REVISION_ID",
   "Comment": "First Atlas Revision",
   "CreatedAt": "2021-09-11T00:18:49.160000+00:00",
   "DataSetId": "$DATA_SET_ID",
   "Finalized": false,
   "Id": "$REVISION_ID",
   "UpdatedAt": "2021-09-11T00:18:49.160000+00:00"
```

}

Add the API assets to the revision. 2.



Note

You will need to know the ID of the API Gateway REST API you want to import as well as the stage.

Step 4: Add API assets to a revision

API assets contain the information subscribers need to make calls to your API. For more information, see Assets.

In the following procedure, you import data assets, and then finalize the revision.

You use either the AWS Data Exchange console or the AWS CLI to add assets to a revision:

- Adding API assets to a revision (console)
- Adding API assets to a revision (AWS CLI)

Adding API assets to a revision (console)

To add assets to the revision (console)

- 1. Under the API assets section of the data set details page, choose Add API stage.
- Under Select API stage, for Amazon API Gateway API, enter an API in the input box or choose 2. one of the following from the drop-down list:
 - API in another AWS account this is a cross account API that you have been given permission to access.
 - In this AWS account this is an API in your AWS account.
 - If you chose API in another AWS account, enter the API ID and the API Stage name in the a. input boxes.
 - If you chose In this AWS account, choose the API Stage name from the drop-down list



Note

You can create a new API stage by choosing **Create new** and following the steps in the Create new API on Amazon API Gateway modal. Once the new stage has been created, repeat Step 2.

- Under Advanced configuration optional, you can choose to Connect existing Amazon API Gateway usage plan to use the throttling and quota limits as defined in the existing usage plan, and enter the API key.
- Under **Document API for subscribers**, provide details about the API that the subscribers will 4. see after they subscribe to your product.
 - For **API** name, enter a name that subscribers can use to identify the API asset.



Note

If an In this AWS account was selected, the API name is automatically populated, which you can modify if necessary.

If a API in another AWS account was selected, the API name is populated with a default name, which you should modify to so the subscriber can easily understand what it is.

- For **OpenAPI 3.0 specification**, either:
 - i. Enter or copy and paste the OpenAPI 3.0 specification file.
 - ii. Choose Import from .JSON file, and then select the .json file from your local computer to import.

The imported specification appears in the box.

iii. Choose Import from Amazon API Gateway, and then choose a specification to import.

The imported specification appears in the box.

For **Additional documentation - optional**, enter any additional information that is useful C. for the subscriber to know about your API. Markdown is supported.



Note

You can't edit the OpenAPI specification and additional documentation after you add this asset to a revision.

If you want to update this information, and the revision is not finalized, you can replace the asset.

If you want to update this information, and the revision is finalized, you can create a new revision with the updated asset.

5. Choose **Add API stage**.

A job is started to import your asset (in this case, the API) into your data set.



Note

If you do not have an API on Amazon API Gateway, you will be prompted to create one.

- After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**. 6.
- If you have more APIs to add, repeat Step 2. 7.
- 8. Under Revision overview, review your revision and its assets.
- Choose Finalize. 9.

You have successfully finalized a revision for a data set.

You can edit a revision or delete a revision before you add it to a product.

You are now ready to publish a new API data product.

Adding API assets to a revision (AWS CLI)

You can add API assets by running an IMPORT_ASSET_FROM_API_GATEWAY_API job.

To add API assets to a revision (AWS CLI):

Use the create-job command to add API assets to the revision:

```
AWS dataexchange create-job \
--type IMPORT_ASSET_FROM_API_GATEWAY_API \
```

```
--details '{"ImportAssetFromApiGatewayApi":
{"DataSetId":"$DATA_SET_ID","RevisionId":"$REVISION_ID","ApiId":"$API_ID","Stage":"$API_STA
{
    "Arn": "arn:aws:dataexchange:us-east-1:123456789012:jobs/$JOB_ID",
    "CreatedAt": "2021-09-11T00:38:19.875000+00:00",
    "Details": {
        "ImportAssetFromApiGatewayApi": {
            "ApiId": "$API_ID",
            "DataSetId": "$DATA_SET_ID",
            "ProtocolType": "REST",
            "RevisionId": "$REVISION_ID",
            "Stage": "$API_STAGE"
        }
    },
    "Id": "$JOB_ID",
    "State": "WAITING",
    "Type": "IMPORT_ASSET_FROM_API_GATEWAY_API",
    "UpdatedAt": "2021-09-11T00:38:19.875000+00:00"
}
$ AWS dataexchange start-job --job-id $JOB_ID
$ AWS dataexchange get-job --job-id $JOB_ID
{
    "Arn": "arn:aws:dataexchange:us-east-1:0123456789012:jobs/$JOB_ID",
    "CreatedAt": "2021-09-11T00:38:19.875000+00:00",
    "Details": {
        "ImportAssetFromApiGatewayApi": {
            "ApiId": "$API_ID",
            "DataSetId": "$DATA_SET_ID",
            "ProtocolType": "REST",
            "RevisionId": "$REVISION_ID",
            "Stage": "$API_STAGE"
            "ApiEndpoint": "string",
            "ApiKey": "string",
            "ApiName": "string",
            "ApiDescription": "string",
            "ApiSpecificationDownloadUrl": "string",
            "ApiSpecificationDownloadUrlExpiresAt": "string"
        }
    },
    "Id": "$JOB_ID",
    "State": "COMPLETED",
    "Type": "IMPORT_ASSET_FROM_API_GATEWAY_API",
    "UpdatedAt": "2021-09-11T00:38:52.538000+00:00"
```

}

2. Use the list-revision-assets command to confirm that the new asset was created properly:

```
$ AWS dataexchange list-revision-assets \
  --data-set-id $DATA_SET_ID \
  --revision-id $REVISION_ID
{
    "Assets": [
    {
        "Arn": "arn:aws:dataexchange:us-east-1:123456789012:data-sets/$DATA_SET_ID/
revisions/$REVISION_ID/assets/$ASSET_ID",
        "AssetDetails": {
            "ApiGatewayApiAsset": {
                "ApiEndpoint": "https://$API_ID.execute-api.us-
east-1.amazonaws.com/$API_STAGE",
                "ApiId": "$API_ID",
                "ProtocolType": "REST",
                "Stage": "$API_STAGE"
            }
        },
        "AssetType": "API_GATEWAY_API",
        "CreatedAt": "2021-09-11T00:38:52.457000+00:00",
        "DataSetId": "$DATA_SET_ID",
        "Id": "$ASSET_ID",
        "Name": "$ASSET_ID/$API_STAGE",
        "RevisionId": "$REVISION_ID",
        "UpdatedAt": "2021-09-11T00:38:52.457000+00:00"
    }
    ]
}
```

You are now ready to publish the API data product.

Edit a revision

To edit the revision after you've finalized it

On the Revision overview, choose De-finalize.

You see a message that the revision is no longer in the finalized state.

- 2. To edit the revision, from **Revision overview**, choose **Actions**, **Edit**.
- 3. Make your changes, and then choose **Update**.
- 4. Review your changes and then choose **Finalize**.

Delete a revision

To delete the revision after you've finalized it

- 1. On the **Revision overview**, choose **Delete**.
- 2. Type **Delete** in the **Delete revision** dialog box, and then choose **Delete**.

Marning

This deletes the revision and all of its assets. This action cannot be undone.

Step 5: Publish a new product containing APIs

After you've created at least one data set and finalized a revision with assets, you're ready to publish that data set as a part of a product. For more information, see Product best practices in AWS Data Exchange. Make sure that you have all required details about your product and offer.

You use the AWS Data Exchange console or the AWS Marketplace Catalog API to publish a new product containing APIs. For more information about how to publish a new product using the AWS Marketplace Catalog API, see Using AWS Data Exchange with the AWS Marketplace Catalog API.

Publishing a new product containing APIs (console)

Publishing a new product containing APIs (console)

To publish a new product containing APIs

- 1. From the left navigation pane of the <u>AWS Data Exchange console</u>, under **Publish data**, choose **Products**.
- 2. From **Products**, choose **Publish new product** to open the **Publish new product** wizard.
- In Product visibility:

Choose your product's **Product visibility options** as either **Public** or **Private**. a.

All AWS Data Exchange products with visibility set to **Public** require a public offer.

For more information, see Product visibility in AWS Data Exchange.

Choose your product's **Sensitive information** configuration.

For more information, see Sensitive categories of information in AWS Data Exchange.

Choose **Next**. c.

In Add data:

Under Owned data sets, select the check boxes next to the data sets you want to add, and a. then choose Add selected.



Note

The data sets you choose must have a finalized revision. Data sets without finalized revisions aren't added.

b. Go to **Selected data sets** to review your selections.

You can review the **Name** of the data set, the **Type** of data set, and the timestamp of when the data set was Last updated.

Go to Select revision access rules, and choose the revision access rules that you want to set for data sets included in this product.

For more information, see Revision access rules in AWS Data Exchange.

Choose **Next**.

In **Define product**: 5.

Under **Product overview**, enter information about your product, including the **Product** a. name, Product logo, Support contact information, and Product categories.

For more information, see Product best practices in AWS Data Exchange.

- (Optional) Under **Data dictionaries and samples optional**, choose a data set by selecting the option button next to the data set name and then choose **Edit**.
 - i. In the **Edit** dialog box, choose **Upload** to upload a new data dictionary.

You can choose one data dictionary, in .csv format, with a maximum size of 1 MB.

ii. Choose a saved data dictionary from your computer and then choose **Open**.

The data dictionary .csv file appears on the **Edit** dialog box.



Note

Your data dictionary must conform to the AWS Data Exchange data dictionary template. If you don't have a saved data dictionary to upload, you can choose either the blank data dictionary template link or the example data dictionary link in the AWS Data Exchange console.

- iii. Choose **Data dictionary preview** to preview the data dictionary.
- Under Samples optional, choose Upload samples, choose a sample from your iv. computer, and then choose **Open**.

The samples appear on the **Edit** dialog box.



Note

You can upload up to 10 samples with a maximum size of 50 MB. Samples in .csv format can be previewed.

- Enter a description for each sample that will be visible on the product detail page. ٧.
- vi. Choose Save.
- Under **Product definition**, enter a **Short description** and a **Long description** of your product. 6.

If you want to use a template for your long description, select **Apply template**, choose your template type, and then fill out the template with your specific product details.

- 7. Choose Next.
- Configure your offer in either Add public offer (for public offer) or Add custom offer (for 8. private offers):

All AWS Data Exchange products with visibility set to **Public** require a public offer.

For private offers only: a.

- i. Choose one of the listed **Offer types**: **Private offer**, **Renewed private offer**, or **Bring Your Own Subscription (BYOS)**.
- ii. In the **Subscriber account information** section, add at least one subscriber account to which you want to extend the offer.
- b. Choose your **Pricing and access duration** options for the subscription.
- c. For **Metered costs optional**, choose **Add**.
 - i. For **Add metered cost**, select the type of cost for the API call from the **Type** list:
 - Per API request
 - · Per successful API request
 - · Per unit of data transferred in bytes
 - New custom metered cost
 - ii. Enter or update the **Cost display name**, which is visible on the subscriber's invoice.
 - iii. If you're using a **Pre-defined metered cost**, the **Key** is automatically generated, can't be edited, and doesn't need to be sent back in the response header.
 - iv. If you're creating a **New custom metered cost**, enter the **Key**, which is the identifier for the metered cost in the API response header (15 characters maximum).

This **Key** should be sent back as part of the x-amz-dataexchange-metering response header.

Example Custom key

If you have a custom key called **VertexCount** and another custom key called **EdgeCount**, the "x-amz-dataexchange-metering" response header could have a value of VertexCount=3, EdgeCount=10 or you could return two separate header lines:

```
x-amz-dataexchange-metering: VertextCount=3
```

x-amz-dataexchange-metering: EdgeCount=10

- v. Enter the price the subscriber is charged per unit in **Price / unit**.
- vi. (Optional) Enter the number of units to display an example of the cost in the **Metered** cost calculator.
- vii. (Optional) Enter a brief **Description** of the metered cost that appears on the product detail page.

viii. Choose Add.

(Optional) Repeat to add additional metered costs.

The order of the metered costs appears on the product detail page. You can't reorder them.

Note

After the offer is created, you can edit the price and description of a metered cost. For more information, see Updating product and offer details in AWS Data Exchange.

- d. For private offers only, choose the **Offer expiration date** by which the subscriber must accept the offer.
- Choose your **Tax settings**, **Data subscription agreement (DSA)**, and **Refund policy**. e.
- f. (Optional) For public offers only, set **Subscription verification**, which enables you to control who can subscribe to this product. For more information, see Subscription verification for providers in AWS Data Exchange.
- Choose your Offer auto-renewal option. For more information, see Creating an offer for AWS Data Exchange products.
- Choose **Next**.
- In the **Review & publish** section, review your product information. 9.
 - Expand the **Product page preview** to see how the product page will look after a. publication.
 - (Optional) Choose the **Edit** button in any section to edit that section.
- 10. If you're sure that you want to make the product and public offer visible and available to everyone, choose **Publish**.

You've now completed the manual portion of publishing a data product with a public offer. AWS Data Exchange prepares and publishes your product.

On the **Product overview** page, the status of your product is **Awaiting approval** and then changes to **Published** after it's published.

Step 6: (Optional) Copy a product

After you have created your first product, you can copy its details and public offers to create a new product.



Note

You can copy a public, private, published, or unpublished product. Custom offers associated with the product will not be copied, but public offers will be copied.

To copy a product

- Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. From **Products**, choose the button next to the product you want to copy.
- Select the **Actions** dropdown, and then choose **Create copy**. 4.
- 5. Continue through the **Publish a new product** workflow, with details already filled in, based on the product you chose in Step 3. For more information, see Step 5: Publish a new product.

Publishing a product in AWS Data Exchange containing Amazon Redshift data sets

An Amazon Redshift data set contains AWS Data Exchange datashares for Amazon Redshift. When customers subscribe to a product containing datashares, they are granted read-only access to the tables, views, schemas, and user-defined functions that a data provider adds to the datashare.

As a data provider, you create an AWS Data Exchange for Amazon Redshift datashare in your cluster. Then, you add to the datashare the schemas, tables, views, and user-defined functions that you want the subscribers to access. You then import the datashare to AWS Data Exchange, create a data set, add it to a product, and publish the product. Subscribers are granted access to the datashare upon subscription.

After you have set up your Amazon Redshift datashare in Amazon Redshift, you can create a new Amazon Redshift data set in AWS Data Exchange. You can then create a revision, and add Amazon Redshift datashare assets. This allows requests to the AWS Data Exchange endpoint to proxy through to your Amazon Redshift datashare. You can then add this data set to a product and add

pricing. Then, prospective subscribers can view your product and subscribe to it in the AWS Data Exchange catalog.

The following topics describe the process of creating an Amazon Redshift data set and publishing a new product with Amazon Redshift data sets using the AWS Data Exchange console. The process has the following steps:

Steps

- Step 1: Create an Amazon Redshift datashare asset
- Step 2: Create an Amazon Redshift data set
- Step 3: Create a revision
- Step 4: Add Amazon Redshift datashare assets to a revision
- Step 5: Publish a new product containing Amazon Redshift data sets
- Step 6: (Optional) Copy a product

Step 1: Create an Amazon Redshift datashare asset

Assets are the data in AWS Data Exchange. For more information, see Assets.

To create an Amazon Redshift datashare asset

Create a datashare within your Amazon Redshift cluster.

For more information about how to create a datashare, see Working with AWS Data Exchange datashares as a producer in the Amazon Redshift Database Developer Guide.



Note

We recommend setting your datashare as publicly accessible. If you do not, customers with publicly accessible clusters will not be able to consume your data.

Step 2: Create an Amazon Redshift data set. 2.

Step 2: Create an Amazon Redshift data set

An Amazon Redshift data set includes AWS Data Exchange datashares for Amazon Redshift. For more information, see Amazon Redshift data set.

To create an Amazon Redshift data set

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. On the left side navigation pane, under **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- 4. In Select data set type, choose Amazon Redshift datashare.
- 5. In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see Data set best practices.
- 6. Under Add tags optional, add tags.
- 7. Choose **Create**.

Step 3: Create a revision

In the following procedure, you create a revision after you've created a data set in the AWS Data Exchange console. For more information, see Revisions.

To create a revision

- 1. On the **Data set overview** section of the data set details page:
 - a. (Optional) Choose **Edit name** to edit information about your data set.
 - b. (Optional) Choose **Delete** to delete the data set.
- 2. On the **Revisions** section, choose **Create revision**.
- 3. Under **Define revision**, provide an optional comment for your revision that describes the purpose of the revision.
- 4. Under **Add tags optional**, add tags associated with the resource.
- Choose Create.
- 6. Review, edit, or delete your changes from the previous step.

Step 4: Add Amazon Redshift datashare assets to a revision

In the following procedure, you add Amazon Redshift datashare assets to a revision, and then finalize the revision in the AWS Data Exchange console. For more information, see <u>Assets</u>.

To add assets to the revision

- Under the AWS Data Exchange datashares for Amazon Redshift section of the data set details page, choose **Add datashares**.
- Under AWS Data Exchange datashares for Amazon Redshift, select the datashares and then choose Add datashare(s).



Note

You can add up to 20 datashares to a revision.

A job is started to import your assets into your revision.

- After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**. 3.
- 4. If you have more data to add, repeat Step 1.
- 5. Under **Revision overview**, review your revision and its assets.
- Choose Finalize. 6.

You have successfully finalized a revision for a data set.

You can edit or delete a revision before you add it to a product.

Step 5: Publish a new product containing Amazon Redshift data sets

After you've created at least one data set and finalized a revision with assets, you're ready to publish a product with Amazon Redshift data sets. For more information, see Product best practices in AWS Data Exchange. Make sure that you have all required details about your product and offer.

To publish a new product containing Amazon Redshift data sets

- From the left navigation pane of the AWS Data Exchange console, under Publish data, choose Products.
- From **Products**, choose **Publish new product** to open the **Publish new product** wizard.
- In the **Product visibility** section, choose your product's **Product visibility options** and **Sensitive information** configuration, and then choose **Next**. For more information, see

Product visibility in AWS Data Exchange and Sensitive categories of information in AWS Data Exchange.

In the **Add data** section, under **Owned data sets**, select the check boxes next to the data sets that you want to add, and then choose Add selected.



Note

The data sets you choose must have a finalized revision. Data sets without finalized revisions won't be added.

Go to **Selected data sets** to review your selections. a.

You can review the **Name** of the data set, the **Type** of data set, and the timestamp of when the data set was Last updated.

Go to **Select revision access rules**, choose the revision access rules that you want to set for data sets included in this product, and then choose Next.

For more details, see Revision access rules in AWS Data Exchange.

In the **Define product** section, under **Product overview**, enter information about your product, including the Product name, Product logo, Support contact information, and Product categories.

For more information, see Product best practices in AWS Data Exchange.

(Optional) In the **Define product** section, under **Data dictionaries and samples – optional**, choose a data set by selecting the option button next to the data set name and then choose Edit.

For more information, see Data dictionaries in AWS Data Exchange and Sample data in AWS Data Exchange.

In the **Edit** dialog box, under **Upload data dictionary**, choose **Add file** to upload a new data dictionary.

You can choose one data dictionary, in .csv format, with a maximum size of 1 MB.

b. Choose a saved data dictionary from your computer and then choose **Open**.

The data dictionary .csv file appears on the **Edit** dialog box.



Note

Your data dictionary must conform to the AWS Data Exchange data dictionary template. If you don't have a saved data dictionary to upload, you can choose either the blank data dictionary template link or the example data dictionary link in the AWS Data Exchange console.

- c. Choose **Data dictionary preview** to preview it.
- d. Under Samples - optional, choose Upload samples, choose a sample from your computer, and then choose **Open**.

The samples appear on the **Edit** dialog box.



Note

You can upload up to 10 samples with a maximum size of 50 MB. Samples in .csv format can be previewed.

- Enter a description for each sample that will be visible on the product detail page.
- f. Choose Save.
- 7. Under **Product definition**, enter a **Short description** and a **Long description** of your product.

If you want to use a template for your long description, select **Apply template**, choose your template type, and then fill out the template with your specific product details.

- 8. Choose Next.
- Configure your offer.
 - If you are creating a public offer, in the **Add public offer** section, configure your offer. All AWS Data Exchange products with visibility set to **Public** require a public offer.
 - 1. Choose your **Pricing and access duration** options for the subscription.
 - 2. Choose your US sales tax settings, data subscription agreement (DSA), and refund policy.
 - 3. (Optional) Set **Subscription verification**, which enables you to control who can subscribe to this product. For more information, see Subscription verification for providers in AWS Data Exchange.
 - 4. Choose your **Offer auto-renewal** option. For more information, see Creating an offer for AWS Data Exchange products.

5. Choose Next.

- If you are creating a private offer, configure the offer details in the Add custom offer section.
 - 1. In the **Subscriber account information** section, add at least one subscriber account to which you want to extend the offer.
 - 2. Choose your **Pricing and access duration** options for the subscription.
 - 3. Choose the **Offer expiration date** by which the subscriber must accept the offer.
 - 4. Choose your US sales tax settings, data subscription agreement (DSA), and refund policy.
 - 5. Choose your **Offer auto-renewal** option. For more information, see Creating an offer for AWS Data Exchange products.
 - 6. Choose Next.
- 10. In the Review & publish section, review your product information and then expand the **Product page preview** to see how it will look after it's published.
- 11. If you're sure that you want to make the product and public offer visible and available to everyone, choose **Publish**.

You've now completed the manual portion of publishing a data product with a public offer. AWS Data Exchange prepares and publishes your product. On the **Product overview** page, the status of your product is **Awaiting approval** and then changes to **Published** after it's published.

Step 6: (Optional) Copy a product

After you have created your first product, you can copy its details and public offers to create a new product.



Note

You can copy a public, private, published, or unpublished product. Custom offers associated with the product will not be copied, but public offers will be copied.

To copy a product

- Open your web browser and sign in to the AWS Data Exchange console. 1.
- From the left navigation pane, under **Publish data**, choose **Products**. 2.

- 3. From **Products**, choose the button next to the product you want to copy.
- 4. Select the **Actions** dropdown, and then choose **Create copy**.
- 5. Continue through the **Publish a product** workflow, with details already filled in, based on the product you chose in Step 3. For more information, see Step 5: Publish a new product.

Publishing a product in AWS Data Exchange containing Amazon S3 data access

With AWS Data Exchange for Amazon S3, providers can share direct access to Amazon S3 buckets or specific prefixes and Amazon S3 objects. Providers also use AWS Data Exchange to automatically manage subscriptions, entitlements, billing, and payments.

As a data provider, you can share direct access to an entire Amazon S3 bucket or specific prefixes and Amazon S3 objects without creating or managing copies. These shared Amazon S3 objects can be server-side encrypted with customer managed keys stored in AWS Key Management Service (AWS KMS) or with AWS managed keys (SSE-S3). For more information about monitoring your KMS keys and understanding encryption contexts, see the section called "Key management for Amazon S3 data access". When a customer subscribes to your data products, AWS Data Exchange automatically provisions an Amazon S3 access point and updates its resource policies on your behalf to grant subscribers read-only access. Subscribers can use the Amazon S3 access point aliases in places where they use Amazon S3 bucket names to access data in Amazon S3.

When the subscription ends, the subscriber's permissions are revoked. If you choose to end an agreement with a subscriber early, contact <u>AWS Support</u>. You can add terms of subscriptions in the Data Subscription Agreement (DSA).

Before you can publish a product containing Amazon S3 data access, you must meet the following prerequisites:

Prerequisites

- Confirm that the Amazon S3 buckets hosting the data are configured with the Amazon S3
 bucket owner enforced setting turned on ACLs Disabled. For more information, see <u>Controlling</u>
 ownership of objects and disabling ACLs for your bucket in the Amazon Simple Storage Service
 User Guide.
- Your shared objects must be in the Amazon S3 Standard Storage class, or be managed using S3
 Intelligent Tiering, for subscribers to access them successfully. If they're in other storage classes,

or if you have enabled Intelligent Tiering with Deep Archive, your subscribers will receive errors because they won't have permission to RestoreObject.

- Confirm that the Amazon S3 buckets hosting the data has encryption disabled or encrypted with Amazon S3 managed keys (SSE-S3) or customer managed keys stored in AWS Key Management Service (AWS KMS).
- If you're using customer managed keys, you must have the following:
 - IAM permissions to kms: CreateGrant on the KMS keys. You can access these permissions
 through the key policy, IAM credentials, or through an AWS KMS grant on the KMS key. For
 more information about key management and understanding how AWS Data Exchange uses
 AWS KMS grants, see Creating AWS KMS grants.

To provide access, add permissions to your users, groups, or roles:

• Users and groups in AWS IAM Identity Center:

Create a permission set. Follow the instructions in <u>Create a permission set</u> in the *AWS IAM Identity Center User Guide*.

Users managed in IAM through an identity provider:

Create a role for identity federation. Follow the instructions in <u>Create a role for a third-party</u> identity provider (federation) in the *IAM User Guide*.

- IAM users:
 - Create a role that your user can assume. Follow the instructions in <u>Create a role for an IAM</u> user in the *IAM User Guide*.
 - (Not recommended) Attach a policy directly to a user or add a user to a user group. Follow the instructions in Adding permissions to a user (console) in the *IAM User Guide*.

Users need programmatic access if they want to interact with AWS outside of the AWS Management Console. The way to grant programmatic access depends on the type of user that's accessing AWS.

To grant users programmatic access, choose one of the following options.

Which user needs programmatic access?	То	Ву
Workforce identity (Users managed in IAM Identity Center)	Use temporary credentia Is to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Configuring the AWS CLI to use AWS IAM Identity Center in the AWS Command Line Interface User Guide. • For AWS SDKs, tools, and AWS APIs, see IAM Identity Center authentic ation in the AWS SDKs and Tools Reference Guide.
IAM	Use temporary credentia ls to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions in Using temporary credentia ls with AWS resources in the IAM User Guide.

Which user needs programmatic access?	То	Ву
IAM	(Not recommended) Use long-term credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Authenticating using IAM user credentials in the AWS Command Line Interface User Guide. • For AWS SDKs and tools, see Authenticate using long-term credentials in the AWS SDKs and Tools Reference Guide. • For AWS APIs, see Managing access keys for IAM users in the IAM User Guide.

Following is an example JSON policy that shows how you could add to the key policy of the KMS key.

```
{
    "Sid": "AllowCreateGrantPermission",
    "Effect": "Allow",
    "Principal": {
    "AWS": "<IAM identity who will call Dataexchange API>"
      },
      "Action": "kms:CreateGrant",
      "Resource": "*"
}
```

The following policy shows an example policy addition for the IAM identity that is used.

JSON

Note

Cross account KMS keys are also permitted if the kms: CreateGrant permission on the KMS keys are obtained through the earlier step. If another account owns the key, you must have permissions on the key policy and your IAM credentials as detailed in the above examples.

- 2. Make sure to use KMS keys to encrypt existing and new objects in the Amazon S3 bucket using the Amazon S3 bucket key feature. For more details, see Configuring S3 Bucket Keys in the Amazon Simple Storage Service User Guide.
 - For new objects added to your Amazon S3 bucket, you can set up Amazon S3 bucket key
 encryption by default. If existing objects have been encrypted without using the Amazon
 S3bucket key feature, these objects must be migrated to use the Amazon S3 bucket key for
 encryption.

To enable the Amazon S3 bucket key for existing objects, use the copy operation. For more information, see Configuring an Amazon S3 bucket key at the object level using batch operations.

- AWS managed KMS keys or AWS owned keys aren't supported. You can migrate from an unsupported encryption scheme to the ones currently supported. For more information, see Changing your Amazon S3 encryption at the AWS Storage Blog.
- 3. Set the Amazon S3 buckets hosting the data to trust AWS Data Exchange owned access points. You must update these Amazon S3 bucket policies to give AWS Data Exchange permissions to create Amazon S3 access points and grant or remove subscribers' access on your behalf. If the policy statement is missing, you must edit the bucket policy to add the Amazon S3 locations to your data set.

An example policy is shown below. Replace <Bucket ARN> with the appropriate value.

JSON

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "AWS": "*"
            },
            "Action": [
                "s3:GetObject",
                "s3:ListBucket"
            ],
            "Resource": [
                "arn:aws:s3:::BucketName",
                "arn:aws:s3:::/*"
            ],
            "Condition": {
                "StringEquals": {
                     "s3:DataAccessPointAccount": [
                         "337040091392",
                         "504002150500",
                         "366362662752",
                         "330489627928",
                         "291973504423",
                         "461002523379",
                         "036905324694",
                         "540564263739",
                         "675969394711",
                         "108584782536",
```

```
"844053218156"

]
}
}
}
}
```

You can delegate data sharing through AWS Data Exchange to an entire Amazon S3 bucket. However, you can scope delegation to the specific prefixes and objects of the bucket that you want to share in the data set. Following is an example of a scoped policy. Replace <Bucket ARN> and "mybucket/folder1/*" with your own information.

JSON

```
{
 "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DelegateToAdxGetObjectsInFolder1",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::mybucket/folder1/*"
      ],
      "Condition": {
        "StringEquals": {
          "s3:DataAccessPointAccount": [
            "337040091392",
            "504002150500",
            "366362662752",
            "330489627928",
            "291973504423",
            "461002523379",
            "036905324694",
```

```
"540564263739",
            "675969394711",
            "108584782536",
            "844053218156"
          ]
        }
      }
    },
      "Sid": "DelegateToAdxListObjectsInFolder1",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::mybucket",
      "Condition": {
        "StringLike": {
          "s3:prefix": [
            "folder1/*"
          ]
        },
        "StringEquals": {
          "s3:DataAccessPointAccount": [
            "337040091392",
            "504002150500",
            "366362662752",
            "330489627928",
            "291973504423",
            "461002523379",
            "036905324694",
            "540564263739",
            "675969394711",
            "108584782536",
            "844053218156"
          ]
        }
      }
    }
  ]
}
```

Similarly, to scope access to only a single file, a provider can use the following policy.

JSON

```
"Version": "2012-10-17",
  "Statement": [
      "Sid": "DelegateToAdxGetMyFile",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::mybucket/folder1/myfile"
      ],
      "Condition": {
        "StringEquals": {
          "s3:DataAccessPointAccount": [
            "337040091392",
            "504002150500",
            "366362662752",
            "330489627928",
            "291973504423",
            "461002523379",
            "036905324694",
            "540564263739",
            "675969394711",
            "108584782536",
            "844053218156"
          ]
        }
      }
    }
 ]
}
```

The following topics describe the process of creating an Amazon S3 data set and publishing a new product with Amazon S3 data sets using the AWS Data Exchange console. The process has the following steps:

Steps

- Step 1: Create an Amazon S3 data set
- Step 2: Configure Amazon S3 data access
- Step 3: Review and finalize the data set
- Step 4: Add an Amazon S3 data set to an AWS Data Exchange product
- Step 5: Publish a new product containing access to Amazon S3
- Step 6: (Optional) Copy a product

Step 1: Create an Amazon S3 data set

To create an Amazon S3 data set

- 1. On the left side navigation pane, under **Publish data**, choose **Owned data sets**.
- 2. On the left side navigation pane, under **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.
- 4. In Select data set type, choose Amazon S3 data access.
- 5. In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see Data set best practices.
- (Optional) Under Add tags optional, add tags.
- 7. Choose **Create data set** and continue.

Step 2: Configure Amazon S3 data access

Choose the Amazon S3 buckets or Amazon S3 bucket locations that you want to make available to subscribers. You can select an entire Amazon S3 bucket, or specify up to five prefixes or objects within an Amazon S3 bucket. To add more Amazon S3 buckets, you must create another Amazon S3 data share.

To configure shared Amazon S3 data access

- 1. On the Configure Amazon S3 data access page, select Choose Amazon S3 locations.
- In Choose Amazon S3 locations, enter your Amazon S3 bucket name in the search bar or select your Amazon S3 bucket, prefixes, or Amazon S3 files and choose Add selected. Then, choose Add locations.



Note

We recommend choosing a top-level folder where a majority of objects and prefixes are stored so providers don't need to reconfigure which prefixes or objects to share.

- 3. In **Configuration details**, choose your **Requester Pays** configuration. There are two options:
 - Enable Requester Pays (recommended) Requesters will pay for all requests and transfers in the Amazon S3 bucket. We recommend this option because it helps protect against unintended costs from subscriber requests and transfers.
 - Disable Requester Pays You pay for subscriber requests and transfers in the Amazon S3 bucket.

For more information about Requester Pays, see Objects in Requester Pays Buckets in the Amazon Simple Storage Service User Guide.

- Select the **Bucket Policy** that best suits your needs. Choose **General** to use one bucket policy for your entire Amazon S3 bucket. This is a one-time configuration and additional configuration isn't needed to share prefixes or objects in the future. Choose **Specific** to use a bucket policy that is specific to the selected Amazon S3 locations. Your shared Amazon S3 bucket needs a bucket policy in place to create an Amazon S3 data access data set successfully and can't have ACLs enabled.
 - To disable ACLs, navigate to your bucket permissions and set **Object Ownership** to **Bucket** a. owner enforced.
 - To add a bucket policy, copy the bucket statement to your clipboard. In the Amazon S3 console, from the Amazon S3 permissions tab, choose Edit in the bucket policy section, paste the bucket policy into the statement, and **Save changes**.
- If the Amazon S3 bucket contains objects encrypted using AWS KMS customer managed keys, you must share all such KMS keys with AWS Data Exchange. For information about required prerequisites when using KMS keys to encrypt objects in your Amazon S3 bucket, see the section called "Containing Amazon S3 data access". To share these KMS keys with AWS Data Exchange, do the following:
 - From the Configure Amazon S3 data access page, in Customer managed KMS keys, select Choose from your AWS KMS keys or Enter AWS KMS key ARN and select all AWS KMS keys currently being used to encrypt the Amazon S3 shared locations. AWS Data

Exchange uses these KMS keys to create grants for subscribers to access your shared locations. For more information, see Grants in AWS KMS.



Note

AWS KMS has a limit of 50,000 grants per KMS key including pre-existing grants.

Review your Amazon S3 locations, selected KMS keys, and configuration details, and choose Save and continue.

Step 3: Review and finalize the data set

Review and finalize your newly created data set. If you wish to create and add another Amazon S3 data access to share access to additional Amazon S3 buckets, prefixes, objects, choose Add another Amazon S3 data access.



Note

We recommend this when needing to share access to data hosted in a different Amazon S3 bucket than the one previously picked in the initial Amazon S3 data access.

If you would like to make changes prior to publishing, you can save the data set as a draft by choosing **Save draft**. Then, choose **Finalize data set** to add it to your product.

Step 4: Add an Amazon S3 data set to an AWS Data Exchange product

In the following procedure, you add your data set to a new or existing AWS Data Exchange product.

To add a data set to a new or existing AWS Data Exchange product

- On the **Owned data sets** page, under **Data set overview**, you can **Edit name**, **Delete**, or **Create** product from data set.
- Complete the product creation specifying product description, use cases, metadata, pricing, and terms and conditions.
- **Review and publish** the product when finished.



Note

When a customer subscribes to your product, the customer receives access permission to read and use your data using the Amazon S3 access point created on your behalf.

Step 5: Publish a new product containing access to Amazon S3

After you create at least one data set and finalize a revision with assets, you can publish a product with Amazon S3 data access. For more information, see Product best practices in AWS Data Exchange. Make sure that you have all required details about your product and offer.



Note

You don't need to create a new revision when updating the shared Amazon S3 objects unless the Amazon S3 locations have been altered and these objects aren't accessible to subscribers.

To publish a new product containing access to Amazon S3

- From the left navigation pane of the AWS Data Exchange console, under **Publish data**, choose Products.
- From **Products**, choose **Publish new product** to open the **Publish new product** wizard.
- In the **Product visibility** section, choose your product's **Product visibility options** and Sensitive information configuration, and then choose Next. For more information, see Product visibility in AWS Data Exchange and Sensitive categories of information in AWS Data Exchange.
- 4. In the Add data section, under Owned data sets, select the check boxes next to the data sets that you want to add, and then choose **Add selected**.



Note

The data sets you choose must have a finalized revision. Data sets without finalized revisions aren't added.

- Go to **Selected data sets** to review your selections. a.
 - You can review the **Name** of the data set, the **Type** of data set, and the timestamp of when the data set was **Last updated**.
- Go to **Select revision access rules**, choose the revision access rules that you want to set for data sets included in this product, and then choose **Next**.

For more details, see Revision access rules in AWS Data Exchange.

In the **Define product** section, under **Product overview**, enter information about your product, including the Product name, Product logo, Support contact information, and Product categories.

For more information, see Product best practices in AWS Data Exchange.

(Optional) In the **Define product** section, under **Data dictionaries and samples – optional**, choose a data set by selecting the option button next to the data set name and then choose Edit.

For more information, see Data dictionaries in AWS Data Exchange and Sample data in AWS Data Exchange.

In the **Edit** dialog box, under **Upload data dictionary**, choose **Add file** to upload a new a. data dictionary.

You can choose one data dictionary, in .csv format, with a maximum size of 1 MB.

b. Choose a saved data dictionary from your computer and then choose **Open**.

The data dictionary .csv file appears on the **Edit** dialog box.



Note

Your data dictionary must conform to the AWS Data Exchange data dictionary template. If you don't have a saved data dictionary to upload, you can choose either the blank data dictionary template link or the example data dictionary link in the AWS Data Exchange console.

c. Choose **Data dictionary preview** to preview the data dictionary. Under Samples - optional, choose Upload samples, choose a sample from your computer, and then choose **Open**.

The samples appear on the **Edit** dialog box.



Note

You can upload up to 10 samples with a maximum size of 50 MB. Samples in .csv format can be previewed.

- Enter a description for each sample that will be visible on the product detail page.
- f. Choose Save.
- Under **Product definition**, enter a **Short description** and a **Long description** of your product. 7.

If you want to use a template for your long description, select **Apply template**, choose your template type, and then provide your specific product details in the template.

- Choose Next. 8.
- Configure your offer.
 - If you're creating a public offer, in the Add public offer section, configure your offer. All AWS Data Exchange products with visibility set to **Public** require a public offer.
 - 1. Choose your **Pricing and access duration** options for the subscription.
 - 2. Choose your US sales tax settings, data subscription agreement (DSA), and refund policy.
 - 3. (Optional) Set **Subscription verification** to control who can subscribe to this product. For more information, see Subscription verification for providers in AWS Data Exchange.
 - 4. Choose your Offer auto-renewal option. For more information, see Creating an offer for AWS Data Exchange products.
 - 5. Choose **Next**.
 - If you're creating a private offer, configure the offer details in the **Add custom offer** section.
 - 1. In the **Subscriber account information** section, add at least one subscriber account to which you want to extend the offer.
 - 2. Choose your **Pricing and access duration** options for the subscription.
 - 3. Choose the **Offer expiration date** by which the subscriber must accept the offer.
 - 4. Choose your US sales tax settings, data subscription agreement (DSA), and refund policy.

- 5. Choose your Offer auto-renewal option. For more information, see Creating an offer for AWS Data Exchange products.
- 6. Choose **Next**.
- 10. In the Review & publish section, review your product information and then expand the **Product page preview** to see how it will look after it's published.
- 11. If you're sure that you want to make the product and public offer visible and available to everyone, choose Publish.

You've now completed the manual portion of publishing a data product with a public offer. AWS Data Exchange prepares and publishes your product. On the **Product overview** page, the status of your product is **Awaiting approval**. The status changes to **Published** after the product is published.

Step 6: (Optional) Copy a product

After you have created your first product, you can copy its details and public offers to create a new product.



You can copy a public, private, published, or unpublished product. Custom offers associated with the product can't be copied, but public offers can be copied.

To copy a product

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. From **Products**, choose the option next to the product that you want to copy.
- Select the **Actions** dropdown list, and then choose **Create copy**. 4.
- 5. Continue through the **Publish a product** workflow, with details already filled in, based on the product you chose in Step 3. For more information, see Step 5: Publish a new product.

Publishing a product in AWS Data Exchange containing AWS Lake Formation data permission data sets (Preview)

If you're interested in publishing products containing AWS Lake Formation data permission data sets during this Preview, contact AWS Support.

An AWS Lake Formation data permission data set contains a set of LF-tags and permissions for data managed by AWS Lake Formation. When customers subscribe to a product containing Lake Formation data permissions, they are granted read-only access to the databases, tables, and columns associated with the LF-tags added to the data set.

As a data provider, you start by creating LF-tags in AWS Lake Formation and associating those tags with the data you want to make available to subscribers. For more information about tagging your resources in Lake Formation, see <u>Lake Formation Tag-based access control</u> in the *AWS Lake Formation Developer Guide*. Then you import those LF-tags and a set of data permissions into AWS Data Exchange as an asset. Subscribers are granted access to the data associated with those LF-tags upon subscription.

The following topics describe the process of publishing a product containing AWS Lake Formation data permissions. The process has the following steps:

Steps

- Step 1: Create an AWS Lake Formation data set (Preview)
- Step 2: Create an AWS Lake Formation data permission (Preview)
- Step 3: Review and finalize
- Step 5: (Optional) Create a revision
- Step 6: Publish a new product containing AWS Lake Formation data sets (Preview)
- Considerations when publishing an AWS Lake Formation data permission data set (Preview)

Step 1: Create an AWS Lake Formation data set (Preview)

To create an AWS Lake Formation data set

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. In **Owned data sets**, choose **Create data set** to open the **Data set creation steps** wizard.

- In Select data set type, choose AWS Lake Formation data permission. 4.
- In **Define data set**, enter a **Name** and **Description** for your data set. For more information, see 5. the section called "Data set best practices".
- Under Add tags optional, choose Add new tag. 6.
- 7. Choose Create data set and continue.

Step 2: Create an AWS Lake Formation data permission (Preview)

AWS Data Exchange uses LF-Tags to grant data permissions. Choose the LF-Tags that are associated with the data you want to share to grant subscriber permissions to the data.

To create AWS Lake Formation data permission

- On the Create Lake Formation data permission page, choose Add LF-Tag.
- 2. Enter the **Key** and choose your LF-Tag **Values**.
- Choose **Preview resource(s)** to view how your LF-Tags are interpreted. 3.
 - From Preview resource(s), select your Associated data catalog resource(s).



Note

Make sure to revoke IAMAllowedPrincipals group on the following resources. For more information, see Revoking IAM role temporary security credentials in the IAM User Guide.

- Review the interpretation of the LF-Tag expression in the dialog box below and **Permissions** associated with the data set.
- For **Service access**, select your existing service role that allows AWS Data Exchange to assume the role and access, grant, and revoke entitlements to Lake Formation data permissions on your behalf. Then choose **Create Lake Formation data permission**. For more information about creating a role for an AWS service, see Creating a role to delegate permissions to an AWS service.
- In the **Define product** section, under **Product overview**, enter information about your product, including the Product name, Product logo, Support contact information, and Product categories.

For more information, see Product best practices in AWS Data Exchange.

(Optional) In the **Define product** section, under **Data dictionaries and samples – optional**, 7. choose a data set by selecting the option button next to the data set name and then choose Edit.

For more information, see Data dictionaries in AWS Data Exchange and Sample data in AWS Data Exchange.

In the **Edit** dialog box, under **Upload data dictionary**, choose **Add file** to upload a new a. data dictionary.

You can choose one data dictionary, in .csv format, with a maximum size of 1 MB.

Choose a saved data dictionary from your computer and then choose **Open**. b.

The data dictionary .csv file appears on the **Edit** dialog box.



Note

Your data dictionary must conform to the AWS Data Exchange data dictionary template. If you don't have a saved data dictionary to upload, you can choose either the blank data dictionary template link or the example data dictionary link in the AWS Data Exchange console.

- Choose **Data dictionary preview** to preview the data dictionary. c.
- Under Samples optional, choose Upload samples, choose a sample from your computer, d. and then choose **Open**.

The samples appear on the **Edit** dialog box.



Note

You can upload up to 10 samples with a maximum size of 50 MB. Samples in .csv format can be previewed.

- Enter a description for each sample that will be visible on the product detail page. e.
- f. Choose Save.
- Under **Product definition**, enter a **Short description** and a **Long description** of your product. 8.

If you want to use a template for your long description, select **Apply template**, choose your template type, and then provide your specific product details in the template.

9. Choose Next.

- Configure your offer.
 - If you're creating a public offer, in the Add public offer section, configure your offer. All AWS
 Data Exchange products with visibility set to Public require a public offer.
 - 1. Choose your **Pricing and access duration** options for the subscription.
 - 2. Choose your US sales tax settings, data subscription agreement (DSA), and refund policy.
 - 3. (Optional) Set **Subscription verification** to control who can subscribe to this product. For more information, see <u>Subscription verification</u> for providers in AWS <u>Data Exchange</u>.
 - 4. Choose your **Offer auto-renewal** option. For more information, see <u>Creating an offer for</u> AWS Data Exchange products.
 - 5. Choose **Next**.
 - If you're creating a private offer, configure the offer details in the **Add custom offer** section.
 - 1. In the **Subscriber account information** section, add at least one subscriber account to which you want to extend the offer.
 - 2. Choose your **Pricing and access duration** options for the subscription.
 - 3. Choose the **Offer expiration date** by which the subscriber must accept the offer.
 - 4. Choose your US sales tax settings, data subscription agreement (DSA), and refund policy.
 - 5. Choose your **Offer auto-renewal** option. For more information, see <u>Creating an offer for AWS Data Exchange products</u>.
 - 6. Choose **Next**.
- 11. In the **Review & publish** section, review your product information and then expand the **Product page preview** to see how it will look after it's published.
- 12. If you're sure that you want to make the product and public offer visible and available to everyone, choose **Publish**.

You've now completed the manual portion of publishing a data product with a public offer. AWS Data Exchange prepares and publishes your product. On the **Product overview** page, the status of your product is **Awaiting approval**. The status changes to **Published** after the product is published.

Step 3: Review and finalize

After creating your AWS Lake Formation data permission (Preview), you can Review and finalize your data set.

To review and finalize

- Review your **Data set details** and **Tags** in **Step 1** for accuracy.
- Review your LF-Tag expression(s), Add another Lake Formation data permission (optional), 2. **Associated data catalog resource(s)**, and job details.



Note

Job are deleted 90 days after they're created.

Choose Finalize.

Step 5: (Optional) Create a revision

To create a revision

- From the **Owned data sets** section, choose the data set for which you want to add a revision. 1.
- Choose the **Revisions** tab. 2.
- 3. In the **Revisions** section, choose **Create revision**.
- On the Revise Lake Formation data permission page, choose Add LF-Tag. 4.
- Review the **Permissions** for **Database** and **Table**. 5.
- From Service access, select an existing service role and then choose Create Lake Formation data permission.

Step 6: Publish a new product containing AWS Lake Formation data sets (Preview)

After you've created at least one data set and finalized a revision with assets, you're ready to publish a product with AWS Lake Formation data sets. For more information, see the section called "Product best practices". Make sure that you have all required details about your product.

To publish a new product containing AWS Lake Formation data sets (Preview)

Open your web browser and sign in to the AWS Data Exchange console.

- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. From **Products**, choose **Publish new product** to open the **Publish new product** wizard.
- In the **Product visibility** section, choose your product's **Product visibility options** and 4. Sensitive information configuration, and then choose Next. For more information, see the section called "Product visibility" and the section called "Sensitive categories of information".
- In the Add data section, under Owned data sets, select the check boxes next to the data sets that you want to add, and then choose Add selected.

Note

The data sets you choose must have a finalized revision. Data sets without finalized revisions aren't added.

- Go to **Selected data sets** to review your selections. a.
 - You can review the **Name** of the data set, the **Type** of data set, and the timestamp of when the data set was **Last updated**.
- Go to **Select revision access rules**, choose the revision access rules that you want to set for data sets included in this product, and then choose Next.

For more details, see Revision access rules in AWS Data Exchange.

Considerations when publishing an AWS Lake Formation data permission data set (Preview)

To ensure an optimal subscriber experience, we strongly advise against making any of the following modifications to any permissions where your product contains AWS Data Exchange for Lake Formation data sets (Preview) with active subscribers to that product.

- We recommend not deleting or modifying IAM roles passed to AWS Data Exchange in published products containing AWS Lake Formation data sets. If you delete or modify such IAM roles, the following issues occur:
 - AWS accounts that have access to the Lake Formation data permissions might retain access indefinitely.

• AWS accounts that subscribe to your product but have not yet received access to the Lake Formation data permissions will fail to receive access.

AWS Data Exchange will not be liable for any IAM roles that you delete or modify.

- We recommend that you don't revoke granted AWS Lake Formation data permissions from IAM roles passed to AWS Data Exchange in published product containing AWS Lake Formation data sets. If you revoke granted data permissions from such IAM roles, the following issues occur:
 - AWS accounts that have access to the Lake Formation data permissions might retain access indefinitely.
 - AWS accounts that subscribe to your product but have not yet received access to the Lake Formation data permissions will fail to receive access.
- We recommend not revoking granted AWS Lake Formation data permissions from AWS accounts
 with active subscriptions to published products containing AWS Lake Formation data sets. If you
 revoke granted data permissions from AWS accounts subscribed to your product, those accounts
 will lose access, creating a poor customer experience.
- We recommend setting the cross account version in your AWS Glue Data Catalog to version 3
 when publishing products containing AWS Lake Formation data sets. If you downgrade the cross
 account version of your Data Lake Catalog while having published products containing AWS Lake
 Formation data sets, the AWS accounts that subscribe to your product but have not yet received
 access to the Lake Formation data permissions may fail to get access to the data.

Product best practices in AWS Data Exchange

When you publish a product on the AWS Data Exchange console, you must provide the product's details. This section covers some best practices to consider when you're preparing product details.

Topics

- Product visibility in AWS Data Exchange
- Sensitive categories of information in AWS Data Exchange
- AWS Data Exchange product details
- Revision access rules in AWS Data Exchange
- Data dictionaries in AWS Data Exchange
- Sample data in AWS Data Exchange

Product best practices 199

Product visibility in AWS Data Exchange

When you create a product in AWS Data Exchange, you choose its visibility. **Product visibility** can be either **Public** or **Private**:

- Public The product is visible in the public catalog in the AWS Data Exchange console and AWS Marketplace. Public products must have a public offer associated with them, and they might also have custom offers.
- Private The product is not publicly visible in the public catalogs of either AWS Data Exchange or AWS Marketplace, and can only have custom offers created for it. Only the specific accounts for whom you have created a custom offer can see the product and subscribe to it. Subscribers can view custom offers created for them on their My product offers tab of AWS Data Exchange.



Note

You can't modify the visibility of a product after it has been created.

For more information about creating a product (with either public or private visibility), see Step 5: Publish a new product.

Sensitive categories of information in AWS Data Exchange

When you create a product in AWS Data Exchange, you must specify whether your product contains any personal data or sensitive categories of information.

Sensitive categories of information include: biometric or genetic data; health data; racial or ethnic origin; political opinions; religious or philosophical beliefs; sex or sexual orientation; trade union membership; personal payment or financial information (for example, credit history); or other similar categories of information.

Personal data is data that identifies or can be used to identify a natural person.

Before accepting a private offer, prospective subscribers will be alerted on the product detail page that your product contains sensitive categories of personal information and/or personal information that is not otherwise publicly available.

As part of the process described in Step 5: Publish a new product, you choose the options for your product's **Sensitive information** configuration. Choose one of the following options:

Product visibility 200 Option 1 – No personal data that is not otherwise publicly available, and no sensitive categories of information

Choose this option if your product does not contain any personal data that is not otherwise publicly available, and no sensitive categories of information.

Examples include financial market data, weather patterns, or public company filings.

Option 2 – No personal data but contains sensitive categories of information

Choose this option if your product contains non-personal sensitive information.

Examples include aggregated diversity data or anonymized financial data.

 Option 3 – Personal data (i) with sensitive categories of information and/or (ii) not otherwise publicly available and does not include Protected Health Information (PHI) under the Health **Insurance Portability and Accountability Act of 1996 (HIPAA)**

Choose this option if your product contains personal data that is not otherwise publicly available. The product must not include protected health information (PHI) subject to HIPAA.

Examples include PII such as email addresses, Social Security numbers, biometrics, or mobile IDs.



Note

This option is only available to eligible providers enrolled in the Extended Provider Program who have agreed to the Extended Provider Program Addendum to the Terms and Conditions for AWS Marketplace Providers. For more information, see Extended Provider Program (EPP).

 Option 4 – Protected Health Information (PHI) subject to the Health Insurance Portability and Accountability Act of 1996 (HIPAA)

Choose this option if your product contains protected health information (PHI) subject to HIPAA.

Examples include PHI such as patient information disclosed by a covered entity.



Option 4 is only available for private products. Public products may not contain such data.

Note

Option 4 is only available to the following eligible providers:

- Eligible providers enrolled in the Extended Provider Program who have agreed to the Extended Provider Program Addendum to the Terms and Conditions for AWS Marketplace Providers. For more information, see Extended Provider Program (EPP).
- Eligible providers who have agreed to the AWS Business Associate Addendum, as well as the AWS Data Exchange Addendum to the AWS Business Associate Addendum.

Marning

If you are not enrolled in the Extended Provider Program, listing a product with data or information described in Option 3 and Option 4 is a violation of our Publishing guidelines for AWS Data Exchange. AWS removes any product that breaches these guidelines and can suspend the provider from future use of the service.

For more information about creating a product and setting the sensitivity status of the data, see Step 5: Publish a new product.

AWS Data Exchange product details

The following topics provide best practices for the details of a product in AWS Data Exchange.

Product name

Subscribers will search for the names of products in AWS Data Exchange, so make your product name something meaningful.

Product logo

The product logo appears in the AWS Data Exchange product catalog on the console and on AWS Marketplace. The supported formats for the logo are .png, .jpg, and .jpeg.

Product details 202

Support contact

As a provider, you must include valid contact information in AWS Data Exchange. This can be a managed email alias or case management system link for customers to use to get help when they have guestions about your product. We strongly recommend that you don't use a personal email address because the address is publicly visible.

Product categories

All products fit into one or more categories in AWS Data Exchange. By specifying up to two categories for your product, you help subscribers filter and find your products in AWS Data Exchange and AWS Marketplace.

Short description for products

The product short description text appears on the tiles in the product catalog portion of the AWS Data Exchange console. We recommend that you provide a concise description of your product for this field.

Long description for products

Subscribers see the product long description in the product detail page after the product is published in AWS Data Exchange. We recommend that you list the product's features, benefits, usage, and other information specific to the product.

Product information in the description must accurately represent the data being provided to subscribers. This includes data coverage (for example, 30,000 financial instruments or 10,000 location coordinates) and data set update frequency (for example, daily updates or weekly updates).



Note

You can use Markdown templates as a starting point for the long description of a number of popular product types. For more information, see Product description templates in AWS Data Exchange.

Product description additional information

In order to make your product description compelling to prospective subscribers, we recommend you add the following information to your product description:

Product details 203

- Data due diligence questionnaire (DDQ) Typically includes responses to questions regarding
 the firm selling a data set. Examples of the information in a DDQ includes the process that a
 provider goes through to collect the data, or quality control procedures and questions regarding
 regulatory compliance.
- Data set schemas Provide prospective users with detailed descriptions of the structure
 and format of your data sets. Examples of the information in a data set schema include the
 identification of a primary key, field names, field definitions, expected output types for each
 field (for example, string, integer), and acceptable enumerations for each field (for example, 0%–
 100%).
- Trial product listings Many prospective subscribers request trials of data sets before paying for a subscription. Trial products can be published on AWS Data Exchange for subscribers to subscribe to like regular paid products.
- Sample files Sample files are typically smaller versions, or older, out-of-date versions of full production data sets. These sample files give prospective users insights into the outputs they can expect before purchasing a subscription.
- *Product fact sheets* These can be documents, web links, or both to provide subscribers with more granular statistics on the coverage of your data sets, typical use cases for your data sets, and any other factors that differentiate your data sets.

For information about adding links in the description, see Include links in your product description.

Include links in your product description

The long description for an AWS Data Exchange product supports Markdown, which allows you to include links in your product's details page. The following procedure shows you how to add links to websites in your AWS Data Exchange product description.

To include embedded links in your product listing

- 1. Log into the AWS console and navigate to an <u>Amazon S3 bucket</u> that your AWS Data Exchange user has access to. The contents of this bucket are publicly readable.
- 2. Upload the files (for example, documents such as PDF files or Microsoft Excel files) that you want to include in your product listing into the Amazon Simple Storage Service (Amazon S3) bucket. After the upload is complete, make sure you set the file or files to have public read access permissions.
- 3. Choose one of the uploaded files. In the **Overview** tab, you will see a URL for the file. Copy the URL to your clipboard.

Product details 204

- 4. Open the AWS Data Exchange console.
- 5. Choose the product you want to update, and then choose **Edit**.
- 6. From **Product Description**, use the following Markdown formats to link to relevant files (using the URL link you copied previously) or to another URL, like your website.
 - To link to a file stored in an S3 bucket:

```
**_[File name](Object URL from Amazon S3)_**

Description of the object.
```

To link to a trial product listing on AWS Data Exchange:

```
**_[Website Title](URL)_**

Description of the website.
```

7. Choose **Save Changes**. After a few minutes your AWS Data Exchange product listing page should be updated with the new links.

Revision access rules in AWS Data Exchange

Revision access rules specify which revisions subscribers can access when they subscribe to your product in AWS Data Exchange. You choose options for subscribers to get historical and future revisions.

- *Historical revision options* Historical revisions are revisions that you published prior to the subscription start date. You have three options for historical revisions:
 - All pre-existing revisions published prior to subscription Give your subscribers access to all historical revisions.
 - A fixed number of trailing revisions published prior to subscription You choose how many historical revisions your subscribers have access to (from 1 to 100).
 - **No historical revisions** Your subscribers get no access to historical revisions. With this option, your subscribers will initially have no data available, until you publish your next revision after their subscription starts.
- Future revision options Future revisions are revisions that you publish after subscription start. You have two options for future revisions:

Revision access rules 205

- All future revisions published during subscription duration Give your subscribers access to all revisions that you publish until their subscription expires.
- **No future revisions** Your subscribers get no access to future revisions.



Note

You can't choose both No historical revisions and No future revisions. That would create a product with no revisions and no data.

Data dictionaries in AWS Data Exchange

A data dictionary is a visual representation of the contents of your data set in AWS Data Exchange.

Subscribers can view and download a data dictionary before they subscribe to your product to evaluate if your product meets their needs.

You can add one data dictionary to each data set, with a maximum size of 1 MB. The accepted file type for a data dictionary is .csv.

When you create a data dictionary, you include details about what columns are included in the data set and their meaning. Your data dictionary must conform to the AWS Data Exchange data dictionary template. You can download the blank data dictionary template from the AWS Data Exchange console. AWS Data Exchange also provides an example data dictionary for you to view as an example.



Note

A data dictionary is attached to a product and associated with a data set. If you want to have more than one data dictionary for potential subscribers to evaluate, you can create two or more versions of the same product with the same data sets. Then, add a different data dictionary to each product.

For more information about how to add a data dictionary to a product, see Publishing a new product in AWS Data Exchange.

Data dictionaries 206

Sample data in AWS Data Exchange

A sample is a small part of the data in your product on AWS Data Exchange that is intended to show what the entire data set is like.

Subscribers can view and download samples before they subscribe to your product to evaluate if your product meets their needs.

You can upload up to 10 samples to each data set with a maximum size of 50 MB. The accepted file formats for samples are any file type accepted by Amazon S3. Samples in .csv format can be previewed.



Note

Samples are attached to a product and associated with a data set. If you want to have more than 10 samples for potential subscribers to evaluate, you can create two or more versions of the same product with the same data sets. Then, add up to 10 samples to each product.

For more information about how to add a sample to a product, see Publishing a new product in AWS Data Exchange.

Product description templates in AWS Data Exchange

When listing a product on AWS Data Exchange, you should include a long description that contains all the information necessary for subscribers to understand what your product offers. For more information about the product long description, see Long description for products.

This section contains Markdown templates that you can use as a starting point for the long description of a number of popular product types.

You can copy and paste the content below in your long description and use the sections that apply to your data product.

Topics

- Generic long description template for AWS Data Exchange products
- Financial services long description template for AWS Data Exchange products
- Healthcare and life sciences long description template for AWS Data Exchange products
- Marketing and advertising long description template for AWS Data Exchange

Samples 207

- Media and entertainment long description template for AWS Data Exchange products
- Public sector long description template for AWS Data Exchange products
- Retail and location long description template for AWS Data Exchange products

Generic long description template for AWS Data Exchange products

When listing a product on AWS Data Exchange, you should include a long description that contains all the information necessary for subscribers to understand what your product offers. The following is a general, all-purpose template for a long description.

```
## PRODUCT TITLE Data Product Overview
Instructions: Provide a description of the data product and what it contains in this
 section.
## Use Cases
Instructions: Provide a handful of use-cases or guidance of best ways to utilize the
 data product.
## Metadata
Instructions: Provide metadata of your data using a table. Examples include but are not
 limited to:
Description | Value
----|----
Update Frequency | ADD INFO HERE
Data Source(s) | ADD INFO HERE
Original Publisher of data | ADD INFO HERE
Data Creation Date | ADD INFO HERE
Data Modification Date | ADD INFO HERE
Geographic coverage | ADD INFO HERE
Time period coverage | ADD INFO HERE
Is historical data "point-in-time" | YES OR NO
Data Set(s) Format(s) | ADD INFO HERE
Raw or scraped data | ADD INFO HERE
Key Fields | ADD INFO HERE
Key Words | ADD INFO HERE
Number of companies/brands covered | ADD INFO HERE
```

Generic template 208

```
## Key Data Points
Key data points include:
* Key Data Point:
* Key Data Point:
## Additional Information
* [Data Source] (ADD LINK HERE)
* [Data Due Diligence Questionnaire] (ADD LINK HERE)
* [Link to Corresponding ADX Trial Product/ Link to Corresponding ADX Paid Product]
 (ADD LINK HERE)
## Pricing Information
If you would like to tell your subscribers that you would like them to inquire for
custom pricing
(ie you price based on other variables), you can explain here.
## Regulatory and Compliance Information
If this section is applicable, provide an overview of the regulatory guidance and
 compliance
for use of this product. Are there exemptions that need to be linked in order for the
data product to be published?
## Subscription Verification Request Information
If you are enabling subscription verification for your products, you may elect to
indicate the information
that you will require from the prospective subscriber i.e., EIN number, # of
 applications, # of users, # of Regions, etc.
## Need Help?
* If you have questions about our products, contact us using the support information
below.
## About Your Company
Provide a description and/or link about your company
```

Generic template 209

* [Company Fact Sheet] (ADD LINK HERE)

Financial services long description template for AWS Data Exchange products

When listing a product on AWS Data Exchange, you should include a long description that contains all the information necessary for subscribers to understand what your product offers. The following is template for a long description for financial services products.

```
## PRODUCT TITLE Data Product Overview
Instructions: Provide a description of the data product and what it contains in this
 section.
## Use Cases
Instructions: Provide a handful of use-cases or guidance of best ways to utilize the
 data product.
## Metadata
Instructions: Provide metadata of your data using a table. Examples include but are not
 limited to:
Description | Value
----|----
Update Frequency | YOUR INFO HERE
Data Source(s) | YOUR INFO HERE
Original Publisher of data | YOUR INFO HERE
Data Creation Date | YOUR INFO HERE
Data Modification Date | YOUR INFO HERE
Geographic coverage | YOUR INFO HERE
Time period coverage | YOUR INFO HERE
Is historical data "point-in-time" | YES OR NO
Data Set(s) Format(s) | YOUR INFO HERE
Raw or scraped data | YOUR INFO HERE
Key Fields | YOUR INFO HERE
Key Words | YOUR INFO HERE
Number of companies/brands covered | YOUR INFO HERE
Standard entity identifiers | YOUR INFO HERE, EXAMPLE BELOW
examples include(include your identifier above then delete this section)
```

Financial services template 210

```
* CUSIP Number: A unique identification number assigned to all stocks and registered
 bonds in the US & Canada
* ISIN: An International Securities Identification Number that uniquely identifies
 a specific securities issue (a series of stocks/bonds offered to raise funds from
investors)
* RIC: The Reuters Instrument Code is used to identify financial instruments/indices
 used in Refinitiv financial information networks
* Bloomberg ID: 12-digit alpha-numeric ID used to identify securities
* D-U-N-S Number: 9-digit identifier assigned to businesses by Dun & Bradstreet
## Tables
If this section is applicable, you can make a table and include information such as:
Description | Identifier | Format | Frequency
----
FX FWD | FIGI | .CSV | Intraday
USD Deposits | CUSIP | .txt | End of Day
Interest Rate Swaps | ISIN | .json | Daily
Basis Swaps | CUSIP | .xml | Intraday
## Key Data Points
Examples of key data points include:
* Symbol: Ticker symbol for the security
* Exchange: Exchange MIC identifier
* Currency: Trading currency code
* Open: Opening price for the day
* High: High price for the day
* Low: Low price for the day
* Last: Last price for the day
* Volume: Trading volume for the day
* Split Ratio: Ratio of new number of shares to old on the effective date
* Cash Dividend: Cash dividend amount on the ex-dividend date
* Dividend amount:
* Extra dividends:
* Total dividends paid this year:
* Effective dates:
* Textual descriptions of special dividends:
* Dividend Currency: Currency for the cash dividend
## Additional Information
```

Financial services template 211

```
* [Data Source] (ADD LINK HERE)
* [Data Due Diligence Questionnaire] (ADD LINK HERE)
* [Link to Corresponding ADX Trial Product/ Link to Corresponding ADX Paid Product]
 (ADD LINK HERE)
## Pricing Information
If you would like to tell your subscribers that you would like them to inquire for
custom pricing
(ie you price based on other variables), you can explain here.
## Regulatory and Compliance Information
If this section is applicable, provide an overview of the regulatory guidance and
    compliance for use of this product. Are there exemptions that need to be linked in
 order for
    the data product to be published?
## Subscription Verification Request Information
If you are enabling subscription verification for your products, you may elect to
indicate
the information that you will require from the prospective subscriber i.e., EIN number,
# of applications,
# of users, # of Regions, etc.
---
## Need Help?
* If you have questions about our products, contact us using the support information
below.
## About Your Company
Provide a description and/or link about your company
* [Company Fact Sheet] (ADD LINK HERE)
```

Financial services template 212

Healthcare and life sciences long description template for AWS Data Exchange products

When listing a product on AWS Data Exchange, you should include a long description that contains all the information necessary for subscribers to understand what your product offers. The following is template for a long description for healthcare and life sciences products.

```
## PRODUCT TITLE Data Product Overview
Instructions: Provide a description of the data product and what it contains in this
 section.
## Use Cases
Instructions: Provide a handful of use-cases or guidance of best ways to utilize the
 data product.
## Metadata
Instructions: Provide metadata of your data using a table. Examples include but are not
 limited to:
Description | Value
----|----
Update Frequency | YOUR INFO HERE
Data Source(s) | YOUR INFO HERE
Original Publisher of data | YOUR INFO HERE
Data Creation Date | YOUR INFO HERE
Data Modification Date | YOUR INFO HERE
Geographic coverage | YOUR INFO HERE
Time period coverage | YOUR INFO HERE
Is historical data "point-in-time" | YES OR NO
Data Set(s) Format(s) | YOUR INFO HERE
Raw or scraped data | YOUR INFO HERE
Key Fields | YOUR INFO HERE
Key Words | YOUR INFO HERE
Number of companies/brands covered | YOUR INFO HERE
## Key Data Points
Key data points include:
```

```
* Key Data Point:
* Key Data Point:
## Use Cases for the Data Set
Provide a handful of use-cases or quidance of best ways to utilize the data product.
## Target Therapeutic Area / Disease Focus
Provide an overview of which therapeutic areas, diagnoses, procedures, medications,
and more can be analyzed in the data listing, and can other data for different
therapeutic areas be sourced.
## Data Engineering Overview
Provide an overview of how the raw data was engineered. Questions to answer:
* What data models were applied?
* What standards / terminologies applied?
* Was NLP post-processing used in the curation of the data?
## Additional Information
* [Data Source] (ADD LINK HERE)
* [Data Due Diligence Questionnaire] (ADD LINK HERE)
* [Link to Corresponding ADX Trial Product/ Link to Corresponding ADX Paid Product]
 (ADD LINK HERE)
## Pricing Information
If you would like to tell your subscribers that you would like them to inquire for
custom pricing (ie you price based on other variables), you can explain here.
## Regulatory and Compliance Information
If this section is applicable, provide an overview of the regulatory guidance and
compliance for use of this product. Are there exemptions that need to be linked in
order for the data product to be published?
## Subscription Verification Request Information
If you are enabling subscription verification for your products, you may elect to
```

```
indicate the information that you will require from the prospective subscriber i.e.,
EIN number, # of applications, # of users, # of Regions, etc.

---
## Need Help?
* If you have questions about our products, contact us using the support information below.

---
## About Your Company
Provide a description and/or link about your company
* [Company Fact Sheet] (ADD LINK HERE)
```

Marketing and advertising long description template for AWS Data Exchange

When listing a product on AWS Data Exchange, you should include a long description that contains all the information necessary for subscribers to understand what your product offers. The following is template for a long description for marketing and advertising products.

```
## PRODUCT TITLE Data Product Overview
Instructions: Provide a description of the data product and what it contains in this
 section.
## Use Cases
Instructions: Provide a handful of use-cases or guidance of best ways to utilize the
 data product.
## Metadata
Instructions: Provide metadata of your data using a table. Examples include but are not
 limited to:
Description | Value
----|----
Update Frequency | YOUR INFO HERE
Data Source(s) | YOUR INFO HERE
Original Publisher of data | YOUR INFO HERE
Data Creation Date | YOUR INFO HERE
Data Modification Date | YOUR INFO HERE
```

```
Geographic coverage | YOUR INFO HERE
Time period coverage | YOUR INFO HERE
Is historical data "point-in-time" | YES OR NO
Data Set(s) Format(s) | YOUR INFO HERE
Raw or scraped data | YOUR INFO HERE
Key Fields | YOUR INFO HERE
Key Words | YOUR INFO HERE
Number of companies/brands covered | YOUR INFO HERE
Data Channels | Examples include web devices, mobile devices, CTV devices, offline
 purchases, household data, B2B data
## Data Set Specification
The following are examples of data set specifications that you may include if
 applicable:
The data sets are updated at midnight EST daily.
Custom data cuts are available if desired.
## Additional Information
* [Data Source] (ADD LINK HERE)
* [Data Due Diligence Questionnaire] (ADD LINK HERE)
* [Link to Corresponding ADX Trial Product/ Link to Corresponding ADX Paid Product]
 (ADD LINK HERE)
## Pricing Information
If you would like to tell your subscribers that you would like them to inquire for
 custom pricing
(ie you price based on other variables), you can explain here.
## Regulatory and Compliance Information
If this section is applicable, provide an overview of the regulatory guidance and
 compliance for use of this product.
Are there exemptions that need to be linked in order for the data product to be
 published?
## Subscription Verification Request Information
```

```
If you are enabling subscription verification for your products, you may elect to indicate the information
that you will require from the prospective subscriber i.e., EIN number, # of applications, # of users, # of Regions, etc.

---
## Need Help?
* If you have questions about our products, contact us using the support information below.

---
## About Your Company
Provide a description and/or link about your company
* [Company Fact Sheet] (ADD LINK HERE)
```

Media and entertainment long description template for AWS Data Exchange products

When listing a product on AWS Data Exchange, you should include a long description that contains all the information necessary for subscribers to understand what your product offers. The following is template for a long description for media and entertainment products.

```
## PRODUCT TITLE Data Product Overview
Instructions: Provide a description of the data product and what it contains in this section.

## Use Cases
Instructions: Provide a handful of use-cases or guidance of best ways to utilize the data product.

## Metadata
Instructions: Provide metadata of your data using a table. Examples include but are not limited to:

Description | Value
----|-----
Update Frequency | ADD INFO HERE
Data Source(s) | ADD INFO HERE
Original Publisher of data | ADD INFO HERE
```

```
Data Creation Date | ADD INFO HERE
Data Modification Date | ADD INFO HERE
Geographic coverage | ADD INFO HERE
Time period coverage | ADD INFO HERE
Is historical data "point-in-time" | YES OR NO
Data Set(s) Format(s) | ADD INFO HERE
Raw or scraped data | ADD INFO HERE
Key Fields | ADD INFO HERE
Key Words | ADD INFO HERE
Number of companies/brands covered | ADD INFO HERE
Table format examples
## Data Set(s) Inventory
File Description | Format | Initial Size | Revision Frequency | Revision Type
----
New Text Archives | .CSV | 100 GB | Hourly | Incremental
Image Library | .JSON | 1.5 TB | Weekly | Incremental
Ratings | .JSON | 50 MB | Every 5 Min | Republish
## Key Data Points
Examples of key data points include:
* Publisher or Studio
* Title
* Artist Name
* Producer Name
* Director Name
* Distributor
* Distribution Channel
* Release Date
* Publish Date
* Format
* Operating System
* Sale Price
* Number of Transactions
* Number of Streams
* Average rating
* Designated Market Area (DMA)
```

```
* Zip or Postal Code
## Additional Information
* [Data Source] (ADD LINK HERE)
* [Data Due Diligence Questionnaire] (ADD LINK HERE)
* [Link to Corresponding ADX Trial Product/ Link to Corresponding ADX Paid Product]
 (ADD LINK HERE)
## Pricing Information
If you would like to tell your subscribers that you would like them to inquire for
 custom pricing
(i.e., you price based on other variables), you can explain here.
## Regulatory and Compliance Information
If this section is applicable, provide an overview of the regulatory guidance and
 compliance for use of this product.
Are there exemptions that need to be linked in order for the data product to be
 published?
## Subscription Verification Request Information
If you are enabling subscription verification for your products, you may elect to
 indicate the information
that you will require from the prospective subscriber i.e., EIN number, # of
 applications, # of users, # of Regions, etc.
## Need Help?
* If you have questions about our products, contact us using the support information
 below.
## About Your Company
Provide a description and/or link about your company
* [Company Fact Sheet] (ADD LINK HERE)
```

Public sector long description template for AWS Data Exchange products

When listing a product on AWS Data Exchange, you should include a long description that contains all the information necessary for subscribers to understand what your product offers. The following is template for a long description for public sector products.

```
## PRODUCT TITLE Data Product Overview
Instructions: Provide a description of the data product and what it contains in this
 section.
## Applicable Industries for Data Product Usage
Provide a list of industries that this data product is applicable to.
## Use Cases
Instructions: Provide a handful of use-cases or guidance of best ways to utilize the
 data product.
## Metadata
Instructions: Provide metadata of your data using a table. Examples include but are not
 limited to:
Description | Value
----|----
Update Frequency | YOUR INFO HERE
Data Source(s) | YOUR INFO HERE
Original Publisher of data | YOUR INFO HERE
Data Creation Date | YOUR INFO HERE
Data Modification Date | YOUR INFO HERE
Geographic coverage | YOUR INFO HERE
Time period coverage | YOUR INFO HERE
Is historical data "point-in-time" | YES OR NO
Data Set(s) Format(s) | YOUR INFO HERE
Raw or scraped data | YOUR INFO HERE
Key Fields | YOUR INFO HERE
Key Words | YOUR INFO HERE
Number of companies/brands covered | YOUR INFO HERE
```

Public sector template 220

```
## Additional Information
* [Data Source] (ADD LINK HERE)
* [Data Due Diligence Questionnaire] (ADD LINK HERE)
* [Link to Corresponding ADX Trial Product/ Link to Corresponding ADX Paid Product]
 (ADD LINK HERE)
## Pricing Information
If you would like to tell your subscribers that you would like them to inquire for
custom pricing (ie you price based on other variables), you can explain here.
## Regulatory and Compliance Information
If this section is applicable, provide an overview of the regulatory guidance and
compliance for use of this product. Are there exemptions that need to be linked in
order for the data product to be published?
## Subscription Verification Request Information
If you are enabling subscription verification for your products, you may elect to
indicate the information that you will require from the prospective subscriber i.e.,
EIN number, # of applications, # of users, # of Regions, etc.
## Need Help?
* If you have questions about our products, contact us using the support information
below.
## About Your Company
Provide a description and/or link about your company
* [Company Fact Sheet] ADD LINK HERE
```

Public sector template 221

Retail and location long description template for AWS Data Exchange products

When listing a product on AWS Data Exchange, you should include a long description that contains all the information necessary for subscribers to understand what your product offers. The following is template for a long description for retail and location products.

```
## PRODUCT TITLE Data Product Overview
Instructions: Provide a description of the data product and what it contains in this
 section.
## Use Cases
Instructions: Provide a handful of use-cases or quidance of best ways to utilize the
data product.
## Metadata
Instructions: Provide metadata of your data using a table. Examples include but are not
 limited to:
Description | Value
----|----
Update Frequency | YOUR INFO HERE
Data Source(s) | YOUR INFO HERE
Original Publisher of data | YOUR INFO HERE
Data Creation Date | YOUR INFO HERE
Data Modification Date | YOUR INFO HERE
Geographic coverage | YOUR INFO HERE
Time period coverage | YOUR INFO HERE
Is historical data "point-in-time" | YES OR NO
Data Set(s) Format(s) | YOUR INFO HERE
Raw or scraped data | YOUR INFO HERE
Key Fields | YOUR INFO HERE
Key Words | YOUR INFO HERE
Number of companies/brands covered | YOUR INFO HERE
Data Channels | Examples include web devices, mobile devices, CTV devices, offline
purchases, household data, B2B data
```

Retail and location template 222

```
## Data Set Specification
The following are examples of data set specifications that you can include if
applicable:
The data sets are updated at midnight EST daily.
The data sets are tied to a home address, and attributes correspond to the household
level.
Provider processes opt-outs on a daily basis and remove records from future files.
Custom data cuts are available if desired.
## Additional Information
* [Data Source] (ADD LINK HERE)
* [Data Due Diligence Questionnaire] (ADD LINK HERE)
* [Link to Corresponding ADX Trial Product/ Link to Corresponding ADX Paid Product]
 (ADD LINK HERE)
## Pricing Information
If you would like to tell your subscribers that you would like them to inquire for
 custom pricing
    (i.e., you price based on other variables), you can explain here.
## Regulatory and Compliance Information
If this section is applicable, provide an overview of the regulatory quidance and
 compliance
    for use of this product. Are there exemptions that need to be linked in order for
 the data product
    to be published?
## Subscription Verification Request Information
If you are enabling subscription verification for your products, you may elect to
indicate
    the information that you will require from the prospective subscriber i.e., EIN
number, # of applications, # of users, # of Regions, etc.
## Need Help?
```

Retail and location template 223

```
* If you have questions about our products, contact us using the support information below.

---

## About Your Company

Provide a description and/or link about your company

* [Company Fact Sheet] (ADD LINK HERE)
```

Creating an offer for AWS Data Exchange products

To make a product available, you must create an *offer* in the AWS Data Exchange console. Offers define the terms that subscribers are agreeing to when they subscribe to a product. Products with visibility set to **Public** must have a public offer available to all subscribers. You can also create custom offers for selected subscribers. When you create an offer for your product, you define:

- The data subscription agreement, which defines the terms that a prospective subscriber must agree to before purchasing a subscription for your product.
- Available pricing and duration combinations.
- Whether US sales tax is collected.
- The Terms and Conditions for the refund policy, if any.
- Whether the subscriber must fill out a questionnaire to request a subscription using subscription verification.
- Whether auto-renewal is available for the offer.

You can also create custom offers that you extend to a select AWS account. The custom offer makes it possible for you to set specific terms and pricing for your product. The following topics provide more information about creating all offers.

Topics

- Offer pricing
- US sales and use tax
- Data Subscription Agreement
- Refund policy
- Subscription verification
- Offer auto-renewal

Creating offers 224

- Creating private offers in AWS Data Exchange
- Creating Bring Your Own Subscription offers in AWS Data Exchange
- Viewing AWS Data Exchange subscriptions

Offer pricing

When you define the pricing information, you define the total price and duration of the subscription. Durations are 1–36 months. For public offers, you can specify up to 5 different durations in a single offer.

We recommend that you choose durations that you plan to support for the long run. If you discontinue a duration, AWS cancels the subscription renewal for those affected subscribers who opted into an auto-renewal policy.

The only supported currency for pricing is US dollars (USD). You must specify a price for each duration. For example, you can specify different prices for durations of 1 month, 6 months, 12 months, 24 months, and 36 months in a single offer. All options are available to prospective subscribers. They must choose a single price and duration when they subscribe to your offer, and they must agree to your offer terms and pay upfront for the purchase charges.

US sales and use tax

You can enable US sales tax collection for the offer, based on your tax nexus settings. For more information, see US sales and use tax.

Data Subscription Agreement

The Data Subscription Agreement (DSA) is the standard contract template that AWS Data Exchange offers as the default. The DSA describes the Terms and Conditions for the data product. As a provider, you control the legal terms and usage rights. These terms are part of each offer you create for your product.

You can download the default DSA template on the AWS Data Exchange console and edit it to add your own Terms and Conditions. Or, you can specify your own custom terms by uploading the DSA of your choice. AWS Data Exchange associates the DSA that you specify for the product's offer without any further modifications.

The DSA was developed in collaboration with the subscriber and provider community to address the needs of both parties. The DSA proactively defines common ground across key contractual

Offer pricing 225

clauses like use, warranty, indemnification and governing law. AWS Data Exchange providers can offer the DSA as the EULA for self-service transactions, or private offers. Subscribers can search for, subscribe to, and use data from providers that offer the DSA, and can request a standard DSA for private offers. For private offers, subscribers can request a DSA template from the provider. The DSA terms can be amended to address custom transaction requirements as agreed upon between the parties.

Refund policy

As a provider, you control the refund policy for your product's subscribers. Although AWS Data Exchange doesn't require you to offer refunds, you must clearly specify your refund policy in the offer details. We encourage you to provide these details in a clear and concise manner so that subscribers can contact you in case of any questions or requests. AWS can process refunds that you authorize on your behalf, but as the provider, you must authorize the refunds.

For AWS to process authorized refunds, submit a refund approval form to AWS Support through the AWS Marketplace Management Portal. Your refund request is processed, and the refund is issued to the subscriber. You can view all refunds that AWS processed on your behalf in the monthly billed revenue report.

Subscription verification

As a provider, you have the option to enable subscription verification for your data products on AWS Data Exchange. For more information, see Subscription verification for providers in AWS Data Exchange.

Offer auto-renewal

As a provider, you control the availability of auto-renewal. When you first create an offer, you can choose to enable auto-renewal, which gives subscribers the option to subscribe to the product with automatic renewals. You cannot change this parameter once the offer has been created.



Note

If you set up a flexible payment schedule for a custom private offer, the offer can't be set to auto-renewal.

Refund policy 226

Creating private offers in AWS Data Exchange

AWS Data Exchange gives providers the option to create custom offers, such as private offers.

As a data provider, you can provide your data product to a subscriber at terms that are different from the offer terms available to the general public. For products that are not publicly visible, your private offers are the only terms available to customers, and only customers you create private offers for can see the product. Private offers allow you to create a custom offer for one or more AWS accounts. A private offer can be different from other offers in any dimension, including price, duration, payment schedule, data subscription agreement, or refund policy.

As a provider, after you have created a product, you can then create a private offer and make it available to a group of subscribers of your choosing. For publicly visible products, you must create a public offer before you can create a private offer.

To create a private offer

- 1. Sign in to the AWS Management Console and open the AWS Data Exchange console.
- 2. From the left navigation pane of the console, choose **Products**, and then choose the product for which you want to make a private offer.
- 3. From the **Private offer** tab, choose **Create**.
- 4. On the **Select Offer Type** page, select **Private offer** or **Renewed private offer**, and choose **Next**.



Note

Choose Renewed private offer if this is a renewal of an expired private offer or a preexisting subscription that is being upgraded on AWS Data Exchange. If you choose this option, AWS might audit and verify that your offer is a renewal or upgrade. If AWS is unable to do so, then we may revoke the offer and entitlements to your subscribers.

- 5. Under **Subscriber AWS account ID**, enter the 12-digit account number of the account you are creating a private offer for. Because a single private offer can be extended to multiple accounts, you can add more than one account.
- 6. Under **Description**, provide a short description of the account (for example, the company name of the account).
- 7. Under **Pricing and duration**, provide the offer details, including the duration and pricing information.

Private offers 227 8. Choose the **Specify payment schedule** check box if you want to distribute the **Total price** to the subscriber over multiple payments. You can add an **Upfront payment** that will be invoiced at the time of subscription. You can then choose for the subscriber to make additional monthly or custom payments. If you choose the **Monthly** option, the dates are automatically populated. If you choose the **Custom** option, you must enter the invoice dates (up to 36 payments).

Note

The Offer expiration date is the date by which the subscriber must accept the offer. The private offer is no longer available for subscribing if it is not accepted by this date. The expiration date must be before the second payment.

If you need to expire an offer already created prior to the expiry date, you can return to the offer page, and choose **Expire**. This will expire the offer for all potential subscribers.

- 9. Provide US sales tax and use tax settings, data subscription agreement, auto-renewal settings, and support information.
- 10Choose Next. If you selected Renewed private offer, you must select the check box to indicate that you acknowledge the terms of the renewed private offer.
- 11Make sure that the information is correct, and then choose **Publish**.



Note

After you create the private offer, you can edit all of the fields except for the price and invoice dates.

Creating Bring Your Own Subscription offers in AWS Data Exchange

AWS Data Exchange gives providers the option to create custom offers, such as Bring Your Own Subscription (BYOS) offers.

As a data provider, you might already have subscribers for your data products. BYOS offers allow you to migrate and fulfill existing subscriptions with AWS customers at no additional cost.

With BYOS offers, any billing relationship between you and your subscribers continues. BYOS offers are not subject to fulfillment fees. Subscribers receive an AWS Marketplace invoice for the subscription with no charge. After you create a BYOS offer, we review it and contact you if we have any issues or questions.

BYOS offers 228 Because the lifecycle of the subscription begins outside of AWS Data Exchange, the workflow for migrating an existing subscription to AWS Data Exchange using BYOS requires collaboration between you and the subscriber.

Important

With BYOS offers, you're migrating a subscription that pre-dates the availability of this product on AWS. AWS might verify your BYOS offer with the existing subscription agreement. If AWS cannot verify your BYOS offer, the offer and entitlements might be revoked without notice.

Before creating or accepting a BYOS offer on AWS Data Exchange, the provider and subscriber should perform the following steps together:

Prerequisites

- 1. The provider and the subscriber contact each other about implementing a BYOS AWS Data Exchange solution.
- The subscriber provides the AWS account ID that they want to use to subscribe to data 2. products on AWS Data Exchange.

If you are the provider, follow these steps to create the BYOS offer.

To create a BYOS offer

- Open your web browser and sign in to the AWS Data Exchange console. 1.
- In the navigation pane, under **Publish data**, choose **Products**. 2.
- Choose the product that you want to create the BYOS offer for by selecting the option button next to the product name in the **Products** list.
- From Actions, choose Create custom offer. 4.
- 5. On the **Select offer type** page, for **Offer types**, select the Bring Your Own Subscription (BYOS) option and then choose Next.
- On the Enter pre-existing subscription details page, for Existing agreement, choose Add file to upload your pre-existing subscription and verify that the agreement pre-dates when you created the product on AWS.
- For **Pre-existing subscription start date**, choose the calendar icon and select the start date.

BYOS offers 229

- 8. For **Duration**, enter the number of months applicable.
- 9. On **Auto renew terms**, select **Yes** or **No** to specify if the pre-existing agreement included autorenewal upon expiry of the current subscription.
- 10. In **Refund policy**, enter information regarding the refund policy stated in your pre-existing subscription agreement and then choose **Next**.
- 11. On the Enter subscriber details page, for Subscriber details, enter the subscriber's 12-digit AWS Account ID and a Description and then choose Next.
- 12. On the **Review & publish** page, verify all of the information. Choose **Edit** to make changes to sections if needed.
- 13. In the **Acknowledgement** section, select the check box to acknowledge that you're migrating a pre-existing subscription that pre-dates the availability of this product on AWS.
- 14. Choose Publish.



Auto-renewal settings can't be changed after the BYOS offer is created. Only one AWS account can be added to a BYOS. If multiple accounts are required, create additional BYOS offers.

Viewing AWS Data Exchange subscriptions

You can view all of the subscriptions for any of your products through the **Product overview** page. You can also view subscriptions for each of your offers.

Viewing subscriptions for a product

To view subscriptions for a product

- 1. Open your web browser and sign in to the <u>AWS Data Exchange console</u>.
- 2. From the left navigation pane, expand **Publish data** and choose **Products**.
- 3. From **Products**, choose the product you want to view offers for.
- 4. Choose the Subscriptions tab. From here, you can view all the subscriptions for your product.

Viewing subscriptions 230

You can choose to filter to currently active subscriptions or to archived (expired and ended) subscriptions from the dropdown at the top left of the **Subscriptions** tab.

Viewing subscriptions for an offer

To view subscriptions for a specific offer

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, expand **Publish data** and choose **Products**.
- 3. From **Products**, choose the product you want to view offers for.
- 4. Choose either the **Public offer** or **Custom offers** tab. From here, you can view all the subscriptions for your offer.

You can choose to filter to currently active subscriptions or to archived (expired and ended) subscriptions from the dropdown at the top left of the **Subscriptions** section.

Updating products in AWS Data Exchange

The following sections describe how to update your AWS Data Exchange products. The instructions assume that you're a provider who is familiar with <u>Data in AWS Data Exchange</u>. After you publish a product, you can edit the product's details and its public offer. You can also update the underlying data sets by publishing new revisions to subscribers. For more information, see <u>Revisions</u>.

Topics

- Updating product and offer details in AWS Data Exchange
- Updating a data dictionary in AWS Data Exchange
- Updating a sample in AWS Data Exchange
- Updating custom metadata in AWS Data Exchange
- Publishing a new data set revision in AWS Data Exchange
- Unpublish a product in AWS Data Exchange
- Removing a revision in AWS Data Exchange
- Revoking access to revisions in AWS Data Exchange

Updating products 231

Updating product and offer details in AWS Data Exchange

After you publish a product, you can use the AWS Data Exchange console to edit the product details. You can also edit the product's public or custom offers and change the offer terms. When you update your product's offer terms, subscribers with an active subscription keep their existing offer terms as long as their subscription is active. Subscribers who have chosen auto-renewals use the new offer terms.

Keep the following in mind when you update products:

- You can't remove or edit a subscription duration in your offers. This ensures that existing subscribers retain the ability to renew. If you no longer want to offer a specific subscription duration, you can unpublish your existing product and then publish a new product. For more information, see Unpublish a product in AWS Data Exchange.
- You can't remove data sets from a product after it is published, regardless of how many subscribers have subscribed to your product.
- If you're updating the metered costs for a product that contains APIs:
 - A metered costs price decrease appears immediately on the product detail page for new subscribers.

Marning

If you undo a price decrease for metered costs, you are increasing the price for metered costs. See the following point for more information about metered costs price increases.

• A metered costs price increase will go into effect on the first day of the month, 90 days after the price increase is submitted for existing subscribers OR upon renewal (whichever is sooner). The email is sent to existing subscribers when the price change is submitted. The price increase appears on the product detail page immediately for new subscribers.

Example Example

You submit a metered costs price increase on May 10. Existing subscribers receive an email about the price change. The price increase goes into effect on September 1.



Marning

You can't undo a price increase (because that action decreases the price) before the price increase goes into effect for existing subscribers.

To update a product, data set, or offer details

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. From **Products**, choose the product you want to update. Make sure its status is **Published**.
- From **Product details**: 4.
 - If you're editing a public offer, choose the **Public offer** tab, choose **Edit**, and then follow the instructions to edit the product.
 - If you're editing a private offer, choose the **Custom offers** tab, choose the option button next to the private offer that you want to edit, choose **Edit**, and then follow the instructions to edit the product.
 - For products containing APIs with metered costs, in Metered costs optional, select the option button next to the **Type** of metered costs that you want to edit, and then choose Edit.
 - In the **Edit metered cost** dialog box, update the **Price / unit** or **Description**.
 - C. Choose **Update**.

The updated metered costs appears under **Metered costs – optional**.

- 5. From Data sets, under Sensitive information, choose Edit, and then follow the instructions to edit the information.
- From **Data evaluation**, update the data dictionary or sample by selecting the option button next to the data dictionary or sample **Name** and then choosing **Actions**. For more information, see Updating a data dictionary in AWS Data Exchange and Updating a sample in AWS Data Exchange.
- Configure your offer, depending on the offer type:

- If your product is a public offer, from **Public offer**, choose **Edit**, and then follow the instructions to edit the public offer.
- If your product is a custom offer, from **Custom offers**, choose **Edit**, and then follow the instructions to edit the custom offer.
- If your product is a private offer, from **Private offers**, choose **Edit**, and then follow the instructions to edit the private offer.
- 8. Choose **Update**.

Updating a data dictionary in AWS Data Exchange

You can update a data dictionary in AWS Data Exchange by first removing the existing data dictionary and then uploading a new one.

To update a data dictionary

- 1. Open your web browser and sign in to the <u>AWS Data Exchange console</u>.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. From **Products**, choose the product you want to update and confirm its status is **Published**.
- 4. Choose the **Data evaluation** tab.
- 5. Under **Data dictionary and samples**, expand the data set by choosing the plus icon, and then choose the data dictionary by selecting the option button next to the data dictionary **Name**.
 - a. Choose **Actions**, and then **Remove data dictionary**.
 - The data dictionary is removed.
 - Select the option button next to the data set, choose **Actions**, and then **Upload data** dictionary.
 - c. Choose Add file.
 - d. Select a new data dictionary and then click **Open**.
 - e. Choose Upload.
- 6. (Optional) Choose the data dictionary by selecting the option button next to the data dictionary **Name**, choose **Actions**, and then choose **Download data dictionary (CSV)** to download the data dictionary to your computer.

Updating a data dictionary 234

Updating a sample in AWS Data Exchange

After you publish a product, you can use the AWS Data Exchange console to update the product sample.

To update a sample

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. From **Products**, choose the product you want to update and confirm its status is **Published**.
- Choose the **Data evaluation** tab.
- 5. Under **Data dictionary and samples**, select the option button next to a data set.
- 6. Choose **Actions**, and then choose **Add samples**.
 - a. Choose **Upload samples**.
 - b. Select a new sample from your computer, and then choose **Open**.
 - c. Enter an optional **Description**, and then choose **Add**.
- 7. (Optional) Select the option button next to the sample **Name**, choose **Actions**, and then choose one of the following actions:
 - Download selected sample
 - Preview sample (CSV only)
 - Remove selected sample

Updating custom metadata in AWS Data Exchange

After you publish a product, you can use the AWS Data Exchange console to edit the product's custom metadata.

To update custom metadata

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. From **Products**, choose the product you want to update. Make sure its status is **Published**.
- 4. (Optional) From **Subscriptions**, choose **View custom metadata**, and view the metadata, and then choose **Close**.

Updating a sample 235

- 5. From **Subscriptions**, choose **Edit custom metadata**, and then follow the instructions to edit the metadata or add new metadata.
- 6. Choose Save.

Publishing a new data set revision in AWS Data Exchange

AWS Data Exchange supports dynamically updated products. Subscribers subscribe to the product for a certain duration and access all of the published data sets as long as their subscription is active. For example, a provider might want to provide a product that contains daily closing stock prices for US equities, which would be updated every day with the day's closing prices. You can create and finalize new revisions that will be available in your product's data sets, or add new data sets to your product.

Your product includes some or all historical and future revisions as part of a subscription. For more information, see Revision access rules in AWS Data Exchange.

In the following procedure, you create and finalize a new revision for a data set that has already been published using the AWS Data Exchange console. The data set revision is then automatically published to all products the data set belongs to. For more information, see Revisions.

A provider can revoke subscriber access to a revision and then delete the assets of the revision using the console or the AWS Data Exchange API. For more information, see Revoking access to revisions in AWS Data Exchange.

To publish a new data set revision to a product

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. On the left side navigation pane, under **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose the data set you want to update.
- 4. Navigate to the **Products** tab to make sure that the data set is associated with a published product.
- 5. From the **Revisions** tab, choose **Create revision** to open the **Create revision** page.

- a. (Optional) Under **Revision settings**, provide an optional comment for your revision that describes the purpose of the revision.
- b. (Optional) Under Add tags optional, add tags associated with the resource.
- c. Choose Create revision.

Your new revision is created.

- 6. Under the **Jobs** section, choose either **Import from Amazon S3** or **Upload** (to upload from your computer), depending on if the assets you want to include are stored in an Amazon S3 bucket you own or on your local computer.
 - a. Follow the prompts, depending on your selection. A job is started to import your asset into your data set.
 - b. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.
- 7. Under **Revision overview**, review your revision and its assets, and then choose **Finalize**.

The revision has been published to the product and is now available to subscribers.

Suggested approach for historical data published with the Files delivery type

Some dynamic products contain historical content that subscribers can access. For example, if your product includes a 30-year history of daily closing stock price for US equities, subscribers would get access to that data in addition to the dynamic updates every day.

For these kinds of products that contain a historical record of data, a best practice is to publish all historical data in a single revision of the data set. You can use the optional comment for the revision to indicate that this revision is a single upload of all data history from a specific date.

If the single historical revision contains a time series of multiple objects, you might consider labeling your object names to describe the underlying data periodicity. For example, if your single revision of history contains 200 files each with a week of historical data, you can name each file with a date for the week the data history begins.

Suggested approaches for updates

You can dynamically update your data sets in a number of ways. Here are three example approaches, all of which create a new revision for each update, but the content of the new revision is different.

- Use a new revision for each update that contains only the items that have changed since the last revision Your revision size would be smaller because only those items that have changed are updated. This approach is suitable for data sets for which the updates affect only a small subset of the data and subscribers are focused only on the items that have changed.
- Use a new revision for each update that contains the updated data The new revision contains a full updated file. All items are included in the new revision, including those that have not changed since the last revision. This approach is convenient for subscribers who want to maintain a single up-to-date file for your data. Subscribers export the latest revision's asset or assets to the same destination and override the previous file or files.
- Use a new revision for each update that contains the full history and updated data The new revision contains the full history of the data, including the latest state of the data and the history of the previous revisions. This approach is more storage-heavy. It's suitable for data sets for which subscribers are interested in the latest comprehensive view of the data's history, including any potential past corrections or adjustments. In this approach, each revision is self-sufficient and provides a full view of the data set history with no dependency on previous revisions.

Unpublish a product in AWS Data Exchange

After your product is published in AWS Data Exchange, it's available for all to find and subscribe to, based on the product's visibility settings. You can unpublish a product if you want to achieve any of the following results:

- Remove a product you created for the <u>Publishing a new product in AWS Data Exchange</u> exercise.
- Clean up your resources.
- Remove a product from the publicly listed products on AWS Data Exchange.
- Stop subscribers from auto-renewing your product.

Keep the following in mind when you unpublish a product:

- You can unpublish a product whenever you want.
- If you unpublish a product, it is no longer visible in the AWS Data Exchange catalog or on AWS Marketplace.
- Subscribers with an active subscription maintain access to the data product until the term of their subscription expires.

Unpublish a product 238

- Active subscriptions that expire after you have unpublished your product are not renewed, even if the subscriber has enabled auto-renewal.
- Existing subscribers can still view the product details until their subscription expires.

To unpublish a product

- Open your web browser and sign in to the AWS Data Exchange console. 1.
- 2. From the left navigation pane, under **Publish data**, choose **Products**.
- 3. From **Products**, choose the product you want to remove. Make sure its status is **Published**.
- From **Product overview**, choose **Unpublish**, and then follow the instructions to unpublish the product.



Important

This action can't be undone.

After you complete these steps, your product's status is **Unpublished**. An unpublished product can't be published again, but you can create a new product (with a new product ID) that has the same data sets, product details, and offer details.

Removing a revision in AWS Data Exchange

A provider can revoke subscriber access to a revision and then delete the assets of the revision using the console or the AWS Data Exchange API. For more information, see Revoking access to revisions in AWS Data Exchange.

You can edit or delete a revision after it's finalized, but before you add it to a product. For more information, see the following topics:

- Edit a revision
- Delete a revision

Revoking access to revisions in AWS Data Exchange

As a provider of data products in AWS Data Exchange, you can revoke subscriber access to a specific revision at any time. This action is typically done by providers for compliance reasons. Revoking a

Removing a revision 239 revision doesn't delete the underlying assets. After you have revoked the revision, all subscribers receive an Amazon EventBridge (formerly known as CloudWatch Events) notification that the revision has been revoked. Subscribers can then view the reason for the revoked revision on the AWS Data Exchange console. Subscribers can't export or query the data within a revoked revision.

To be able to revoke revisions, providers who manage their own IAM policies must add dataexchange: RevokeRevision as a new action. Providers who use the <u>managed policies for AWS Data Exchange</u> don't need to make any changes.

After a revision is revoked, you can delete the assets of the revision by using the console or the AWS Data Exchange DeleteAsset API operation.

Topics

- Revoking access to an AWS Data Exchange asset revision (AWS CLI)
- Revoking access to a single AWS Data Exchange asset revision as a provider (console)
- Revoking multiple AWS Data Exchange asset revisions as a provider (console)
- Editing an AWS Data Exchange asset revocation reason as a provider (console)
- Viewing revoked revisions as a subscriber (console)

Revoking access to an AWS Data Exchange asset revision (AWS CLI)

As a provider of AWS Data Exchange data products, you can use the AWS CLI to revoke subscriber access to a revision using the following instructions.

To revoke a revision (AWS CLI)

Use the revoke-revision command to revoke a revision.

```
$ AWS dataexchange revoke-revision \
--data-set-id $DATA_SET_ID \
--revision-id $REVISION_ID \
--comment 'Revoking Revision Example'

{
  "Id": "ab7859881EXAMPLEdd3e8a4b88fc6a8d",
  "Arn": "arn:aws:dataexchange:us-east-1:427362365172:data-sets/$DATA_SET_ID/
  revisions/$REVISION_ID",
  "Comment": "Revoking Revision Example",
  "CreatedAt": "2022-03-08T18:54:20.746Z",
```

Revoking revisions 240

```
"UpdatedAt": "2022-03-09T20:28:53.105Z",
"DataSetId": "24d30f8446a878237c35d011e7b22d0b",
"Finalized": true,
"Revoked": true,
"RevokedAt": "2022-03-09T20:28:53.105Z",
"RevocationComment": "revoking revision example"
}
```

2. After a revision is revoked, you can delete the assets of the revision using the AWS Data Exchange DeleteAsset API operation.

Revoking access to a single AWS Data Exchange asset revision as a provider (console)

As a provider of AWS Data Exchange data products, you can use the AWS Data Exchange console to revoke subscriber access to a single revision using the following instructions.

To revoke revision as a provider (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose the data set that has the revision you want to revoke.
- On the **Revisions** tab, under **Revisions**, choose the revision.
- On the revision page, under **Revision overview**, for **Actions**, choose **Revoke**. 5.
- In the **Revoke revision** dialog box, enter a short description of your reason for revoking the revision. Subscribers will see this description.
- 7. Choose **Revoke**.

The **Status** of the revision is set to **Revoked**.

Marning

This revokes the revision and all of its assets. Subscribers can view the reason for revocation but can't access or export the assets. This action can't be undone.

After a revision is revoked, you can delete the assets of the revision by navigating to the revision page, selecting the assets you want to delete in the **Imported assets** table, and then choosing **Delete**.

Revoking revisions 241 To edit the reason for a revoked revision, see Editing an AWS Data Exchange asset revocation reason as a provider (console).

Revoking multiple AWS Data Exchange asset revisions as a provider (console)

As a provider of AWS Data Exchange data products, you can use the AWS Data Exchange console to revoke subscriber access to multiple revisions using the following instructions.

To revoke multiple revisions as a provider (console)

- Open your web browser and sign in to the AWS Data Exchange console. 1.
- 2. In the left side navigation pane, for **Publish data**, choose **Owned data sets**.
- In **Owned data sets**, choose the data set that has the revisions you want to revoke. 3.
- 4. On the **Revisions** tab, choose up to 10 revisions.
- 5. Choose Revoke.
- In the **Revoke {x} revisions** dialog box, enter a short description of your reason for revoking 6. the revisions. Subscribers will see this description. Then, choose **Revoke**.

The **Status** of the revisions are set to **Revoked**.

Marning

This revokes the revisions and all of the assets. Subscribers can view the reason for revocation but can't access or export the assets. This action can't be undone.

After a revision is revoked, you can delete the assets of the revision by navigating to the revision page, selecting the assets you want to delete in the **Imported assets** table, and then choosing **Delete**.

To edit the reason for a revoked revision, see Editing an AWS Data Exchange asset revocation reason as a provider (console).

Editing an AWS Data Exchange asset revocation reason as a provider (console)

As a provider of AWS Data Exchange data products, you can use the AWS Data Exchange console to edit the reason for the revocation using the following instructions.

Revoking revisions 242

To edit a revocation revision as a provider (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **Publish data products**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose the data set that has the revision you revoked.
- 4. On the **Revisions** tab, choose the revoked revision.
- 5. On the revision page, choose **Edit revocation reason**.
- 6. In the **Edit revocation revision** dialog box, enter a short description of your reason for revoking the revision.
- 7. Choose **Save**.

The **Status** of the revision is set to **Revoked**.

The updated revocation reason is displayed on the revision page.

Viewing revoked revisions as a subscriber (console)

As a subscriber to AWS Data Exchange data products, you can use the AWS Data Exchange console to view the reason for revocation of access to a revision using the following instructions.

To view a revoked revision as a subscriber (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, under **My subscriptions**, choose **Entitled data**.
- 3. Under **Products**, choose a product, and then expand the data set under the product to see a list of revisions.
- On the data set page, under the Revisions tab, view the Status of the revision (Published or Revoked).
- 5. Choose a revision.
- 6. View the revision reason on the top of the revision detail page.

Revoking revisions 243

Subscription verification for providers in AWS Data Exchange

(1) Important regulatory update

Effective April 8, 2025, you may not use AWS Data Exchange to provide products containing Bulk U.S. Sensitive Personal Data or U.S. Government-related Data to Countries of Concern or Covered Persons, as each is defined in the U.S. Department of Justice Final Rule on Preventing Access to U.S. Sensitive Personal Data and Government-Related Data by Countries of Concern and Covered Persons ("DOJ Rule"), available here. You are responsible for compliance with the DOJ Rule. To support compliance, subscription verification is now enabled for all AWS Data Exchange public offers and requires subscribers to identify if they are in a Country of Concern or a Covered Person.

Subscription verification gives you the ability to review a potential subscriber's identity and approve that subscriber for your product. Approving subscription requests to your product is useful when you have restricted or regulated products, or you have products that you want to limit access to. Subscription verification is on the **Add public offer** section of the **Publish new product** page.

Potential subscribers must complete and submit a form for your review. The form requires the following information:

- Prospective subscriber's contact details, including contact name, company name, email address, and whether the subscriber is in a Country of Concern or a Covered Person. For more information on this requirement, see the Publishing guidelines for AWS Data Exchange.
- Prospective subscriber's intended use case.
- Prospective subscriber's AWS account ID.

Important

AWS Data Exchange doesn't review or validate the information provided by a prospective subscriber on the request form. You are solely responsible for reviewing and verifying the information that the subscriber provides.

To view, approve, or decline all subscription verification requests for all of your products, in the AWS Data Exchange console, under **Published to AWS Marketplace**, choose **Verify subscriptions**. For more information, see Approve or decline requests for subscription verification in AWS Data Exchange.



(i) Note

Each subscription request is uniquely identified using its ID. The ID is visible to both the provider and the subscriber. You can use the subscription request ID in your communications with the subscriber.

If you change the product offer terms after a subscriber makes the request, the terms for that subscriber reflect the terms as they were at the time of the request, not the updated terms. Examples of changes to terms include the price, refund policy, or data subscription agreement. If you changed the product offer terms after the request was submitted, a message is displayed in the approval pane of the AWS Data Exchange console to indicate there is a difference between current terms and the terms in place when the request was made.

The AWS Data Exchange console maintains a history of requests. You control when you delete the subscriber's contact details and personally identifiable information (PII). For more information about how to view the request history, see Viewing subscription verification requests.

The following topics provide more information about subscription verification for providers.

Topics

- Email notifications for subscription verification in AWS Data Exchange
- Viewing subscription verification requests
- Approve or decline requests for subscription verification in AWS Data Exchange

Email notifications for subscription verification in AWS Data Exchange

You will receive an email message to your AWS Marketplace registered Seller AWS account email address to notify you when an AWS Data Exchange subscription request is received, or when its status has changed to cancelled or expired. Although most subscription request status changes result in an email notification, the delivery of these email messages is on a best-effort basis.

Email notifications 245



Note

You will not receive email notifications for subscription request status changes that you have initiated yourself (for example, when you approve a subscription).

You can create rules in your mail client to forward subscriber verification emails. These notification emails have the subject "AWS Data Exchange - Subscription Verification Request".

Viewing subscription verification requests

After you publish a public offer and receive subscription verification requests, you can view the requests.

To view subscription verification requests

- Sign in to the AWS Mangement Console and open the AWS Data Exchange console. 1.
- 2. In the left navigation pane, under **Published to AWS Marketplace**, choose **Verify** subscriptions.
- To view pending requests, choose **View pending requests**. Choose **View history** to view all other requests.

Approve or decline requests for subscription verification in AWS Data **Exchange**

Important regulatory update

Effective April 8, 2025, you may not use AWS Data Exchange to provide products containing Bulk U.S. Sensitive Personal Data or U.S. Government-related Data to Countries of Concern or Covered Persons, as each is defined in the U.S. Department of Justice Final Rule on Preventing Access to U.S. Sensitive Personal Data and Government-Related Data by Countries of Concern and Covered Persons ("DOJ Rule"), available here. You are responsible for compliance with the DOJ Rule. To support compliance, subscription verification is now enabled for all AWS Data Exchange public offers and requires subscribers to identify if they are in a Country of Concern or a Covered Person.

The subscriber information you collect through subscription verification must be used in accordance with AWS Marketplace Terms and Conditions.

After you receive the subscription request for AWS Data Exchange, you have 45 days to approve or reject it. If you don't approve the request in that period of time, the request expires. Potential subscribers can resubmit a rejected request at any time, any number of times.

Approving requests

To approve a subscription request

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, under **Published to AWS Marketplace**, choose **Verify subscriptions**.
- 3. On the **Subscription verification** page, choose **View pending requests**.
- 4. Choose **Approve**.

Approving requests for products containing APIs

You can approve a subscription request for a product containing APIs. You can also add custom metadata to product containing APIs that is sent in the header of each AWS Data Exchange request for the specific subscription. The custom metadata isn't visible to subscribers.

To approve a subscription request for a product containing APIs

- 1. Open your web browser, and sign in to the <u>AWS Data Exchange console</u>.
- 2. In the left navigation pane, under **Published to AWS Marketplace**, choose **Verify subscriptions**.
- 3. On the **Subscription verification** page, choose **View pending requests**.
- 4. Choose Approve and add custom API metadata.
- On the modal, enter the key-value pair and then choose Approve and add custom API metadata.



Note

You can add additional key-value pairs if necessary by choosing **Add** and then entering an additional key-value pair.

- You are returned to the **Subscription verification** page. A message informs you that you have successfully accepted the subscription request.
- To view the custom metadata, go to **Products**, select your product with APIs and then select the **Subscriptions** tab.
- Under **Public and custom subscriptions**, you can:
 - Select the subscription, and choose View custom metadata to see the key-value pairs you a. added.
 - Select the subscription, and choose **Edit custom metadata** to edit, add, or remove the key-value pairs for this subscription.



Note

If you add three or more key-value pairs, the **Custom metadata for APIs** column in the Public and custom subscriptions table displays the first key-value pair, and then displays the number of key-value pairs underneath the first key-value pair. For example: keyExample-valueExample +2 more

Declining requests

To decline a subscription request

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left navigation pane, under **Published to AWS Marketplace**, choose **Verify** subscriptions.
- On the **Subscription verification** page, choose **View pending requests**. 3.
- 4. Choose **Decline**.

Provider-generated notifications in AWS Data Exchange

As a provider in AWS Data Exchange, you can send provider-generated notifications to inform your subscribers about important events related to your data sets. You can contact your subscribers in a structured manner and help them to process their entitled data related events in a consistent manner across providers.

Using provider-generated notifications, you do the following to help your subscribers:

- Send notifications for data updates, delays, schema changes, and deprecations using the AWS Data Exchange Console or the AWS SDK.
- Include comments and expected actions for subscribers to follow.

To send provider-generated notifications to subscribers, follow these steps:

- 1. Open and sign in to the AWS Data Exchange console.
- 2. From the left navigation pane, choose **Send notification**.
- 3. Select your **Notification type** from the dropdown menu. Notification types include:
 - Data update the data source has been updated.
 - Data delay the data source hasn't updated as expected.
 - **Schema change** the data source will include a structural change.
 - **Deprecation** the data source will no longer be updated.
- 4. Select the impacted data set from the dropdown menu and view your **Notification details** for the **date**, **time**, and **list** of subscriber actions. You can also provide location metadata for specifying what is affected by this event.
- 5. Choose **Preview notification** and publish your notification.

AWS Data Exchange provider financials on AWS Marketplace

The following topics cover financial information about providing data through AWS Data Exchange.

AWS Data Exchange is integrated with AWS Marketplace. If you want to register as an AWS Data Exchange provider, you must first register as an AWS Marketplace seller. For more information, see Step 2: Register to be a provider.

As an AWS Data Exchange provider, you benefit from AWS Marketplace features, such as Seller Reports and the AWS Marketplace Commerce Analytics Service. For more information, see <u>Seller Reports and Data Feeds</u>.

Payments

AWS disburses payments monthly directly to the bank account associated with the AWS account registered as a seller, minus AWS Marketplace service fees. Payment is disbursed on a rolling monthly basis based on when the account was created, not the beginning of each month. Funds are disbursed to you only after they are collected from the subscriber. For more information, see Disbursement in the AWS Marketplace Seller Guide.

US sales and use tax

AWS Marketplace Tax Calculation Service makes it possible to calculate and collect US sales and use tax for existing and new products. Some states are not eligible for Tax Calculation Service because AWS Marketplace is required by law to collect and remit applicable sales tax attributable to taxable sales of your products to subscribers based in these states. To use the service, configure your tax nexus settings for your provider profile, and then assign product tax codes to your products.

To configure your tax nexus settings

Open the <u>AWS Marketplace Management Portal</u>. On the **Settings** tab, configure the applicable tax nexus settings.

For more information, see Seller registration process in the AWS Marketplace Seller Guide.

AWS Marketplace seller reports

As an AWS Data Exchange provider, you receive reports detailing the subscription activity of your products. There are several reports available to track daily and monthly data. The reports include information about the subscription activity for your offers, payment received from subscribers, and money being disbursed to you. Disbursement doesn't occur until payment is received from the AWS customer. For more information, see Seller reports in the AWS Marketplace Seller Guide.

AWS Data Exchange providers who use the payment scheduler for their private offers can see this data in a monthly report. For more information, see <u>Monthly billed revenue report</u> in the AWS Marketplace Seller Guide.

Payments 250

Subscriber refund requests

As a provider, you control the refund policy for your products, which you must specify when you create your product. AWS Data Exchange doesn't require you to offer refunds. You must approve all requests for refunds before AWS processes them on your behalf.

Submit a <u>refund approval form</u> to AWS Support. They process your request and issue the refund to the subscriber. You can view all refunds that AWS processed on your behalf in the monthly billed revenue report.

Subscriber refund requests 251

Jobs in AWS Data Exchange

AWS Data Exchange jobs are asynchronous import or export operations.

As a provider of data products in AWS Data Exchange, you can create and manage your data sets that you want to publish to a product. You can download (export) or copy your assets or revisions to Amazon Simple Storage Service (Amazon S3) or a signed URL. In addition, providers can import assets from an Amazon API Gateway API or import assets from an Amazon Redshift data set.

As a subscriber, you can view and access the data sets that you have an entitlement to through a subscription. You can use the API operations to download (export) or copy your entitled data sets to Amazon S3 for use with a variety of AWS analytics and machine learning services.

To create or copy assets or copy revisions through jobs, you can use the AWS Management Console, AWS Command Line Interface (AWS CLI), your own REST application, or one of the AWS SDKs.

Jobs are deleted 90 days after they are created.

Topics

- Job properties
- AWS Regions and jobs
- Importing assets to AWS Data Exchange
- Exporting assets from AWS Data Exchange
- Exporting revisions from AWS Data Exchange

Job properties

Jobs have the following properties:

- Job ID An ID generated when the job is created that uniquely identifies the job.
- **Job type** The following job types are supported:
 - Import from Amazon S3
 - Import an AWS Lake Formation data permission (Preview)
 - · Import from signed URL
 - Import from Amazon API Gateway API
 - Import from an AWS Data Exchange datashare for Amazon Redshift

Job properties 252

- Import an Amazon S3 data access
- Export to Amazon S3
- Export to signed URL
- Amazon Resource Name (ARN) A unique identifier for AWS resources.
- **Job state** The job states are WAITING, IN_PROGRESS, COMPLETED, CANCELLED, ERROR, or TIMED_OUT. When a job is created, it's in the WAITING state until the job is started.
- **Job details** Details of the operation to be performed by the job, such as export destination details or import source details.

Example job resource

```
{
    "Arn": "arn:aws:dataexchange:us-
east-1:123456789012:jobs/6cEXAMPLE818f7c7a23b3d0EXAMPLE1c",
    "Id": "6cEXAMPLE818f7c7a23b3d0EXAMPLE1c",
    "State": "COMPLETED",
    "Type": "IMPORT_ASSETS_FROM_S3",
    "CreatedAt": "2019-10-11T14:12:24.640Z",
    "UpdatedAt": "2019-10-11T14:13:00.804Z",
    "Details": {
        "ImportAssetsFromS3": {
            "AssetSources": [
                {
                    "Bucket": "amzn-s3-demo-bucket",
                    "Key": "MyKey"
                }
            ],
            "DataSetId": "14EXAMPLE4460dc9b005a0dEXAMPLE2f",
            "RevisionId": "e5EXAMPLE224f879066f999EXAMPLE42"
        }
    }
}
```

AWS Regions and jobs

If you import or export an asset to or from an Amazon S3 bucket that is in an AWS Region that is different than the data set's Region, your AWS account is charged for the data transfer costs, according to Amazon S3 data transfer pricing policies.

AWS Regions and jobs 253

If you export assets to a signed URL, your AWS account is charged for data transfer costs from Amazon S3 to the internet according to Amazon S3 pricing policies.

When your subscription to an AWS Data Exchange for Files data set ends, you retain access to any files that you already exported. Review your Data Subscription Agreement to verify if your agreement requires that you delete exported data when ending a subscription.

Importing assets to AWS Data Exchange

You can create an AWS Data Exchange job to import data sets that you want to publish to a product. The following sections describe how to import these assets from a variety of locations.

Topics

- Importing AWS Data Exchange assets from an S3 bucket
- Importing AWS Data Exchangeassets from a signed URL
- Importing AWS Data Exchange assets from an Amazon API Gateway API
- Importing AWS Data Exchange assets from an AWS Data Exchange datashare for Amazon Redshift
- Importing AWS Data Exchange assets from AWS Lake Formation (Preview)

Importing AWS Data Exchange assets from an S3 bucket

When you import assets from Amazon S3 to AWS Data Exchange, the AWS Identity and Access Management (IAM) permissions you use must include the ability to write to the AWS Data Exchange service S3 buckets and to read from the S3 bucket where your assets are stored. You can import from any S3 bucket that you have permission to access, regardless of ownership. For more information, see Amazon S3 permissions.

You can import up to 100 assets in a single job.

Topics

- Importing assets from an S3 bucket (AWS SDKs)
- Importing assets from an S3 bucket (console)

Importing assets 254

Importing assets from an S3 bucket (AWS SDKs)

To import assets from an Amazon S3 bucket (AWS SDKs)

- 1. Create a CreateJob request of type IMPORT_ASSETS_FROM_S3.
- 2. Include the following in the request:
 - AssetSources
 - Bucket
 - Key
 - DataSetID
 - RevisionID
- Start the CreateJob request with a StartJob operation that requires the JobId returned in step 1.
- 4. (Optional) Update the assets' name property after they are created.

Importing assets from an S3 bucket (console)

To import an asset from an S3 bucket (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose the data set that has the revision you want to update.
- 4. On the **Revisions** tab, choose **Create revision** to open the **Create revision** page.
 - a. For **Revision settings**, provide an optional comment for your revision that describes the purpose of the revision.
 - b. For **Add tags optional**, add tags associated with the resource.
 - c. Choose Create.

Your new revision is created.

- 5. For the **Jobs** section, choose **Import from Amazon S3**.
- 6. Follow the prompts in the **Import from Amazon S3** window, and then choose **Import assets**.

A job is started to import your asset into your data set. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

From an S3 bucket 255

Importing AWS Data Exchangeassets from a signed URL

You can use signed URLs to import assets that are not stored in Amazon S3.

Topics

- Importing assets from a signed URL (AWS SDKs)
- Importing assets from a signed URL (console)

Importing assets from a signed URL (AWS SDKs)

To import assets from a signed URL (AWS SDKs)

- Create a CreateJob request of type IMPORT_ASSET_FROM_SIGNED_URL.
- 2. Include the following in the request:
 - AssetName
 - DataSetID
 - Md5Hash
 - RevisionID
- 3. Start the CreateJob request with a StartJob operation that requires the JobId returned in step 1.
- 4. (Optional) Update the assets' name property after they are created.
- 5. The response details include the SignedUrl that you can use to import your file.

Note

The signed URL expires one minute after it's created.

Importing assets from a signed URL (console)

To import an asset from a signed URL (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **Publish data**, choose **Owned data sets**.

From a signed URL 256

- 3. In **Owned data sets**, choose the data set that has the asset you want to update.
- 4. On the **Revisions** tab, choose **Create revision** to open the **Create revision** page.
 - a. For **Revision settings**, provide an optional comment for your revision that describes the purpose of the revision.
 - b. For **Add tags optional**, add tags associated with the resource.
 - c. Choose Create.

Your new revision is created.

- 5. For the **Jobs** section, choose **Upload**.
- 6. Follow the prompts in the upload window, and then choose **Open**.

A job is started to import your asset into your data set. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Importing AWS Data Exchange assets from an Amazon API Gateway API

AWS Data Exchange subscribers can use their IAM credentials and AWS SDKs to call APIs from data providers. AWS Data Exchange manages access to APIs by handling authentication and subscription entitlements.

Importing API assets from an Amazon API Gateway API (AWS SDKs)



Currently, the SendApiAsset operation is not supported for the following SDKs:

- · SDK for .NET
- AWS SDK for C++
- AWS SDK for Java 2.x

To import assets from an Amazon API Gateway API (AWS SDKs)

- Create a CreateJob request of type IMPORT_ASSET_FROM_API_GATEWAY_API.
- Include the following in the request:
 - ApiID

- DataSetID
- ProtocolType
- RevisionID
- Stage
- Start the CreateJob request with a StartJob operation that requires the JobId returned in step 1.
- 4. (Optional) Poll the GetJob operation to wait for the job to complete.
- 5. (Optional) Update the assets' name property after they are created.

Importing API assets from an Amazon API Gateway API (console)

To import an asset from an Amazon API Gateway API (console)

- 1. Open your web browser and sign in to the <u>AWS Data Exchange console</u>.
- 2. In the left side navigation pane, for **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose the data set that has the asset you want to update.
- 4. On the **Revisions** tab, choose **Create revision** to open the **Create revision** page.
 - a. For **Revision settings**, provide an optional comment for your revision that describes the purpose of the revision.
 - b. For **Add tags optional**, add tags associated with the resource.
 - c. Choose **Create**.

Your new revision is created.

- 5. For the API assets section, choose Add API stage.
- 6. On the **Add API stage** page, select the **Amazon API Gateway API** and the **Stage name** from your AWS account or another account.
- 7. For **Document API for subscribers**:
 - a. Update the **API name** to a clear and concise name that subscribers can understand.
 - b. Document the OpenAPI 3.0 specification by entering the specification in the field, importing the specification by choosing Import from .JSON file, or importing the specification by choosing Import from Amazon API Gateway.
- 8. Choose **Add API stage**.

A job is started to import your API assets into your data set. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Importing AWS Data Exchange assets from an AWS Data Exchange datashare for Amazon Redshift

When you import assets using AWS Data Exchange datashare for Amazon Redshift, you can begin querying, analyzing, and operationalizing third-party Amazon Redshift tables after subscribing.

Importing assets from an AWS Data Exchange datashare for Amazon Redshift (AWS SDKs)

To import assets from an AWS Data Exchange datashare for Amazon Redshift (AWS SDKs)

- 1. Create a CreateJob request of type IMPORT_ASSETS_FROM_REDSHIFT_DATA_SHARES.
- 2. Include the following in the request:
 - AssetSources
 - DataShareArn
 - DataSetID
 - RevisionID
- Start the CreateJob request with a StartJob operation that requires the JobId returned in step 1.
- 4. (Optional) Poll the GetJob operation to wait for the job to complete.
- 5. (Optional) Update the assets' name property after they are created.

Importing assets from an AWS Data Exchange datashare for Amazon Redshift (console)

To import an asset from an ADE datashare (for Amazon Redshift console)

- 1. Open your web browser and sign in to the <u>AWS Data Exchange console</u>.
- 2. In the left side navigation pane, for Publish data, choose Owned data sets.
- 3. In **Owned data sets**, choose the data set that has the asset you want to update.

- 4. On the **Revisions** tab, choose **Create revision** to open the **Create revision** page.
 - a. For **Revision settings**, provide an optional comment for your revision that describes the purpose of the revision.
 - b. For **Add tags optional**, add tags associated with the resource.
 - c. Choose Create.

Your new revision is created.

- For the AWS Data Exchange datashares for Amazon Redshift section, choose Add datashares.
- 6. On the **Add AWS Data Exchange datashare to revision** page, select the datashare or datashares that you want to add.
- 7. Choose Add datashare(s).

A job is started to import your assets into your data set. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Importing AWS Data Exchange assets from AWS Lake Formation (Preview)

When you import assets from AWS Lake Formation to AWS Data Exchange, the IAM permissions that you use must include the following abilities:

- · Write to, grant, and revoke Lake Formation permissions
- Create resource shares for tables, databases, and catalogs
- Update, delete, associate, and disassociate resource shares for any resource share beginning with Lake Formation

For more information about required and recommended IAM permissions, see <u>the section called</u> "Identity and access management".

Importing assets from AWS Lake Formation (Preview) (AWS SDKs)

To import assets from AWS Lake Formation (Preview) (AWS SDKs)

- Create a CreateJob request of type
 Import_Assets_From_Lake_Formation_Tag_Policy.
- 2. Include the following in the request:
 - AssetSources
 - CatalogId
 - Database
 - Expression
 - TagKey
 - TagValues
 - Permissions
 - Table
 - Expression
 - TagKey
 - TagValues
 - Permissions
 - RoleArn
 - DataSetId
 - RevisionId
- 3. Start the CreateJob request with a StartJob operation that requires the JobId.
- 4. (Optional) Poll the GetJob operation to wait for the job to complete.
- 5. (Optional) Update the assets' name property after they are created.

Importing assets from AWS Lake Formation (Preview) (console)

To import an asset from AWS Lake Formation (Preview) (console)

1. Open your web browser and sign in to the AWS Data Exchange console.

- 3. In **Owned data sets**, choose the data set that has the revision you want to update.
- 4. On the **Revisions** tab, choose **Create revision** to open the **Create revision** page.
 - a. For **Revision settings**, provide an optional comment for your revision that describes the purpose of the revision.
 - b. For **Add tags optional**, add tags associated with the resource.
 - c. Choose Create.

Your new revision is created.

- 5. For the Lake Formation data permission section, choose Add LF-Tag.
- Choose the Key and Values that you want to add and choose Add LF-Tag.
 - (Optional) Choose **Preview Resource(s)** to view the associated data catalog resources that you are granting permission.
- 7. In **Service access**, select the **Role** to import the AWS Lake Formation resources into AWS Data Exchange.
- 8. Choose Create Lake Formation data permission.

A job is started to import your assets into your data set. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Exporting assets from AWS Data Exchange

Both providers and subscribers can export data sets from a published revision of a product in AWS Data Exchange. You can export these assets to an S3 bucket, or to a signed URL. The following sections describe how to do this.

Topics

- Exporting AWS Data Exchange assets to an S3 bucket
- Exporting AWS Data Exchange assets to a signed URL

Exporting AWS Data Exchange assets to an S3 bucket

When you export assets to Amazon S3, the IAM permissions you use must include the ability to read from the AWS Data Exchange service S3 buckets and to write to the S3 bucket where your

Exporting assets 262

assets are stored. You can export to any S3 bucket you have permission to access, regardless of ownership. For more information, see Amazon S3 permissions.

AWS Data Exchange supports configurable encryption parameters when exporting data sets to Amazon S3. In your export job details, you can specify the Amazon S3 server-side encryption configuration that you want to apply to the exported objects. You can choose to use server-side encryption with Amazon S3-Managed Keys (SSE-S3) or server-side encryption with AWS KMS keys stored in AWS Key Management Service (SSE-KMS). For more information, see Protecting data using server-side encryption in the Amazon Simple Storage Service User Guide.

Important

We recommend that you consider Amazon S3 security features when exporting data to Amazon S3. For information about general guidelines and best practices, see Security best practices for Amazon S3 in the Amazon Simple Storage Service User Guide.

Important

If the provider has marked a product as containing protected health information (PHI) subject to the Health Insurance Portability and Accountability Act of 1996 (HIPAA), you may not export the product's data sets into your AWS account unless such AWS account is designated as a HIPAA account (as defined in the AWS Business Associate Addendum found in AWS Artifact).

You can export up to 100 assets in a single job.

Topics

- Exporting AWS Data Exchange assets to an S3 bucket (AWS SDKs)
- Exporting AWS Data Exchange assets to an S3 bucket as a subscriber (console)
- Exporting AWS Data Exchange assets to an S3 bucket as a provider (console)

The following video explains more about how to export assets from AWS Data Exchange.

To an S3 bucket 263

Exporting AWS Data Exchange assets to an S3 bucket (AWS SDKs)

You can use the AWS SDKs to export AWS Data Exchange assets to an S3 bucket using the following instructions.

To export assets to an S3 bucket (AWS SDKs)

- Create a CreateJob request of type EXPORT_ASSETS_TO_S3.
- 2. Include the following in the request:
 - AssetDestinations
 - AssetID
 - Bucket
 - Key
 - DataSetID
 - Encryption
 - KmsKeyArn
 - Type
 - RevisionID
- Start the CreateJob request with a StartJob operation that requires the JobId returned in step 1.
- 4. (Optional) Update the assets' name property after they are created.

Note

For information about exporting an entire revision as a single job, see <u>Exporting revisions</u> from AWS Data Exchange.

Exporting AWS Data Exchange assets to an S3 bucket as a subscriber (console)

As a subscriber to AWS Data Exchange data products, you can use the AWS Data Exchange console to export AWS Data Exchange assets to an S3 bucket using the following instructions.

To an S3 bucket 264

To export an asset to an S3 bucket as a subscriber (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for My subscriptions, choose Entitled data.
- 3. In **Entitled data**, choose the product that has the revision you want to export.
- 4. In **Entitled data sets**, choose the data set.
- 5. On the **Revisions** tab, choose the revision.
- 6. From the **Assets** tab, select the check box next to the assets that you want to export.
- 7. Select **Export actions** and then choose **Export selected assets to Amazon S3**.
- 8. Follow the prompts in the **Export to Amazon S3** window and then choose **Export**.

A job is started to export your asset. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Exporting AWS Data Exchange assets to an S3 bucket as a provider (console)

As a provider of AWS Data Exchange data products, you can use the AWS Data Exchange console to export AWS Data Exchange assets to an S3 bucket using the following instructions.

To export an asset to an S3 bucket as a provider (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose the data set that has the asset you want to export.
- 4. Navigate to the **Products** tab to make sure that the data set is associated with a published product.
- 5. From the **Revisions** tab, select the revision.
- 6. For the **Imported assets** section, select the check box next to the asset name.
- 7. Select **Export actions** and then choose **Export selected assets to Amazon S3**.
- 8. Follow the prompts in the **Export to Amazon S3** window and then choose **Export**.

A job is started to export your asset. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

To an S3 bucket 265

Exporting AWS Data Exchange assets to a signed URL

You can use signed URLs to export AWS Data Exchange assets that are not stored in Amazon S3.

Topics

- Exporting AWS Data Exchange assets to a signed URL (AWS SDKs)
- Exporting assets to a signed URL as a subscriber (console)
- Exporting assets to a signed URL as a provider (console)

Exporting AWS Data Exchange assets to a signed URL (AWS SDKs)

You can use the AWS SDKs to export AWS Data Exchange assets to destinations other than S3 buckets.

To export assets to a signed URL (AWS SDKs)

- 1. Create a CreateJob request of type EXPORT_ASSET_TO_SIGNED_URL.
- 2. Include the following in the request:
 - AssetID
 - DataSetID
 - RevisionID
- Start the CreateJob request with a StartJob operation that requires the JobId returned in step 1.
- 4. (Optional) Update the assets' name property after they are created.
- 5. The response details include the SignedUrl that you can use to import your file.



The signed URL expires one minute after it's created.

To a signed URL 266

Exporting assets to a signed URL as a subscriber (console)

As a subscriber to AWS Data Exchange data products, you can use the AWS Data Exchange console to export AWS Data Exchange assets to destinations other than S3 buckets using the following instructions.

To export an asset to a signed URL as a subscriber (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for My subscriptions, choose Entitled data.
- 3. In **Entitled data**, choose the product that has the revision you want to export.
- 4. In **Entitled data sets**, choose the data set.
- 5. On the **Revisions** tab, choose the revision.
- 6. From the **Assets** tab, select the check box next to the assets that you want to export.
- 7. Select **Export actions** and then choose **Download selected assets**.

A job is started to export your asset. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Exporting assets to a signed URL as a provider (console)

As a provider of AWS Data Exchange data products, you can use the AWS Data Exchange console to export AWS Data Exchange assets to destinations other than S3 buckets using the following instructions.

To export an asset to a signed URL as a provider (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose the product that has the revision you want to export.
- 4. Navigate to the **Products** tab to make sure that the data set is associated with a published product.
- 5. On the **Revisions** tab, choose the revision.
- 6. For the **Imported assets** section, select the check box next to the asset name.
- 7. Select **Export actions** and then choose **Download selected assets**.

To a signed URL 267

A job is started to export your asset. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Exporting revisions from AWS Data Exchange

Both providers and subscribers can export revisions of a data set to an S3 bucket that they have permissions to access.

AWS Data Exchange supports configurable encryption parameters when exporting revisions to Amazon S3. In your export job details, you can specify the Amazon S3 server-side encryption configuration that you want to apply to the exported objects. You can choose to use server-side encryption with Amazon S3-Managed Keys (SSE-S3) or server-side encryption with KMS keys stored in AWS Key Management Service (SSE-KMS). For more information, see Protecting data using server-side encryption in the Amazon Simple Storage Service Developer Guide.

If the provider has marked a product as containing protected health information (PHI) subject to the Health Insurance Portability and Accountability Act of 1996 (HIPAA), you may not export the product's data sets into your AWS account unless such AWS account is designated as a HIPAA account (as defined in the AWS Business Associate Addendum found in AWS Artifact).

Topics

- Key patterns when exporting asset revisions from AWS Data Exchange
- Exporting AWS Data Exchange asset revisions to an S3 bucket (AWS SDKs)
- Exporting AWS Data Exchange asset revisions to an S3 bucket as a subscriber (console)
- Exporting AWS Data Exchange asset revisions to an S3 bucket as a provider (console)
- Automatically exporting AWS Data Exchange asset revisions to an S3 bucket as a subscriber

The following video explains more about how to export assets from AWS Data Exchange (starting at 2:18).

Exporting revisions 268

Key patterns when exporting asset revisions from AWS Data Exchange

When you export an asset revision from AWS Data Exchange, each asset becomes an object in the S3 bucket. The names of the objects are based on a key pattern that you provide. You can use dynamic references that represent asset attributes to create a pattern for the names that are automatically generated during the export. Use the dynamic references shown in the following table.

Dynamic references	Description	
\${Asset.Id}	The Id of the asset.	
\${Asset.Name}	The name of the asset.	
\${DataSet.Id}	The Id of the data set being exported.	
<pre>\${DataSet.Name}</pre>	The name of the data set being exported.	
<pre>\${Revision.CreatedAt}</pre>	The UTC date and time the revision was created, in the following format: YYYY-MM-DDTHH:MM:SSZ. For example: 2021-10-08T16:33:19.787Z	
<pre>\${Revision.Created At.Day}</pre>	The day of the month the revision was created.	
<pre>\${Revision.Created At.Month}</pre>	The month the revision was created.	
<pre>\${Revision.Created At.Year}</pre>	The year the revision was created.	
<pre>\${Revision.Id}</pre>	The Id of the revision being exported.	

You can use these dynamic references to create the key patterns for your asset names. You must include at least one of the two Asset dynamic references, which are \${Asset.Name} and \${Asset.Id}.

For example, using **\${Revision.Id}/\${Asset.Name}** as a key pattern results in Amazon S3 objects that use the revision Id and asset name (separated by a slash) as the object name.

If you export a revision with the Id testRevisionId that has two assets named asset1 and asset2, the assets are exported to the following locations in Amazon S3:

- <bucket>/testRevisionId/asset1
- <bucket>/testRevisionId/asset2

Note

Your resulting objects must have unique names. If they have the same names as existing objects in the S3 bucket, your export will overwrite existing objects. If the revision you are exporting has non-unique names (for example, two assets with the same name), the export will fail. The only dynamic reference that is unique is \${Asset.Id}.

Exporting AWS Data Exchange asset revisions to an S3 bucket (AWS SDKs)

You can use the AWS SDKs to export AWS Data Exchange asset revisions to an S3 bucket using the following instructions.

To export a revision to an S3 bucket (AWS SDKs)

- 1. Create a CreateJob request of type EXPORT_REVISIONS_TO_S3.
- Include the following in the request:
 - DataSetId
 - Encryption
 - KmsKeyArn
 - Type
 - RevisionDestinations
 - Bucket
 - KeyPattern
 - RevisionId
- Start the CreateJob request with a StartJob operation that requires the JobId returned in step 1.

Using AWS SDKs 270 4. The newly created assets have a name property equal to the original S3 object's key. The Amazon S3 object key defaults to the key pattern \${Asset.Name}.

You can update the assets' name property after they are created.

For more information about key patterns, see <u>Key patterns when exporting asset revisions</u> from AWS Data Exchange.

Note

If you are using DataSet.Name as the dynamic reference, you must have the IAM permission dataexchange: GetDataSet. For more information, see <u>AWS Data Exchange</u> API permissions: actions and resources reference.

Exporting AWS Data Exchange asset revisions to an S3 bucket as a subscriber (console)

As a subscriber to AWS Data Exchange data products, you can use the AWS Data Exchange console to export AWS Data Exchange assets to an S3 bucket using the following instructions.

To export a revision to an S3 bucket as a subscriber (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **My subscriptions**, choose **Entitled data**.
- 3. In **Entitled data**, choose the product that has the revision you want to export.
- 4. In Entitled data sets, choose the data set.
- 5. On the **Revisions** tab, select the revision, and then choose **Export to Amazon S3**.
- 6. In **Export revision to Amazon S3**, select a destination option, Amazon S3 bucket folder destination, configure encryption options, and then choose **Export**.

A job is started to export your revision. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Exporting AWS Data Exchange asset revisions to an S3 bucket as a provider (console)

As a provider of AWS Data Exchange data products, you can use the AWS Data Exchange console to export AWS Data Exchange assets to an S3 bucket using the following instructions.

To export a revision to an S3 bucket as a provider (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **Publish data**, choose **Owned data sets**.
- 3. In **Owned data sets**, choose the product that has the revision you want to export.
- 4. Navigate to the **Products** tab to make sure that the data set is associated with a published product.
- 5. On the **Revisions** tab, choose the revision.
- 6. For the **Imported assets** section, select the check box next to the asset name.
- 7. Select **Export actions** and then choose **Export selected assets to Amazon S3**.
- 8. Follow the prompts in the **Export to Amazon S3** window and then choose **Export**.

A job is started to export your asset. After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

Automatically exporting AWS Data Exchange asset revisions to an S3 bucket as a subscriber

When the provider publishes new revisions, you can select to automatically export new revisions to your Amazon S3 bucket. You can export new revisions to up to five S3 buckets. New revisions will automatically appear in the S3 buckets you have selected.

Topics

- Prerequisites for S3 bucket policy permissions
- Automatically exporting revisions to an S3 bucket as a subscriber (console)
- Automatically exporting revisions to an S3 bucket as a subscriber (AWS SDKs)

Using the console (Provider) 272



Note

To automatically export revisions to an S3 bucket of your choice, your S3 bucket must have a bucket policy with permissions set to allow AWS Data Exchange to export data into it. For more information, see Prerequisites for S3 bucket policy permissions.

Prerequisites for S3 bucket policy permissions

Before you can automatically export revisions to an Amazon S3 bucket, you must disable requester pays and your Amazon S3 bucket must have a bucket policy with permissions set to allow AWS Data Exchange to export data into it. The following procedures provide information about how to either edit your existing S3 bucket policy or create an S3 bucket policy with these permissions.

If your S3 bucket is configured for SSE-KMS encryption, the user configuring the auto-export job must have CreateGrant permission on the KMS key for AWS Data Exchange to copy the objects into your S3 bucket.



Important

To verify that the prerequisites for S3 bucket policy permissions are met, an object with the naming format _ADX-TEST-ACCOUNTID# is added to the S3 bucket during the automatic export process.

Topics

- Editing an existing S3 bucket policy
- Creating an S3 bucket policy

Editing an existing S3 bucket policy

If your S3 bucket has a bucket policy, complete the following procedure to allow AWS Data Exchange to export data to it.

To edit an existing S3 bucket policy

1. Navigate to the bucket to which you want to export revisions.

- 2. Select the **Permissions** tab, and choose **Edit** in the bucket policy section.
- 3. Copy the following statement and paste it at the end of the statement list.

```
{
  "Effect": "Allow",
  "Principal": {
  "Service": "dataexchange.amazonaws.com"
  },
  "Action": [
    "s3:PutObject",
    "s3:PutObjectAcl"
 ],
  "Resource": "arn:aws:s3:::<BUCKET-NAME>/*",
  "Condition": {
    "StringEquals": {
      "aws:SourceAccount": "<AWS ID>"
   }
 }
}
```

- 4. Replace <BUCKET-NAME> with the name of your S3 bucket and replace <AWS ID> with your AWS ID.
- 5. Choose **Save changes**.
- 6. If you want to add more buckets as a destination for your auto-export jobs, repeat the procedure, starting from Step 1.

Creating an S3 bucket policy

If your S3 bucket does not have a bucket policy, complete the following procedure to create an S3 bucket policy to allow AWS Data Exchange to export data to it.

To create an S3 bucket policy

- 1. Navigate to the bucket to which you want to export revisions.
- 2. Select the **Permissions** tab, and choose **Edit** in the bucket policy section.
- 3. Copy the following full bucket policy and paste it into the bucket policy editor.

JSON

```
"Version": "2012-10-17",
  "Statement": [
      "Effect": "Allow",
      "Principal": {
      "Service": "dataexchange.amazonaws.com"
      },
      "Action": [
        "s3:PutObject",
        "s3:PutObjectAcl"
      ],
      "Resource": "arn:aws:s3:::<BUCKET-NAME>/*",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "<AWS ID>"
      }
    }
 ]
}
```

- Replace <BUCKET-NAME> with the name of your S3 bucket and replace <AWS ID> with your AWS ID.
- 5. Choose Save changes.
- 6. If you want to add more buckets as a destination for your auto-export jobs, repeat the procedure, starting from Step 1.

Automatically exporting revisions to an S3 bucket as a subscriber (console)



To automatically export revisions to an S3 bucket of your choice, your S3 bucket must have a bucket policy with permissions set to allow AWS Data Exchange to export data into it. For more information, see Prerequisites for S3 bucket policy permissions.

To automatically export a revision to an S3 bucket as a subscriber (console)

- 1. Open your web browser and sign in to the AWS Data Exchange console.
- 2. In the left side navigation pane, for **My subscriptions**, choose **Entitled data**.
- 3. In **Entitled data**, choose the product that has the revision you want to export.
- 4. In **Entitled data sets**, choose the data set.
- 5. On the **Revisions** tab, under **Auto-export job destinations**, choose **Actions** and then choose **Add auto-export job destination**.
- 6. In **Add auto-export job destination**, choose either the **Simple** or **Advanced** destination option.
 - a. If you choose the **Simple** option, select the Amazon S3 bucket folder destination from the dropdown list and the encryption options, and then choose **Add bucket destination**.
 - b. If you choose the **Advanced** option, select the Amazon S3 bucket folder destination from the dropdown list, select the Key naming pattern and append it to the path.
- 7. Review the **Output**.
- 8. Set the **Encryption options**, review the **Amazon S3 pricing**, and then choose **Add bucket destination**.

The Amazon S3 bucket destination appears on the **Revisions** tab under **Auto-export job** destinations.

A job is started to automatically export your revision.

To verify that the prerequisites for S3 bucket policy permissions are met, an object with the naming format _ADX-TEST-ACCOUNTID# is added to the S3 bucket.

After the job is finished, the **State** field in the **Jobs** section is updated to **Completed**.

To add another destination, choose **Actions**, and then **Add auto-export job destination**.

To edit, select the destination you want to edit, choose **Actions**, and then **Edit destination configuration**.

To delete, choose **Actions**, and then choose **Remove auto-export job destination**.

Automatically exporting revisions to an S3 bucket as a subscriber (AWS SDKs)



Note

To automatically export revisions to an S3 bucket of your choice, your S3 bucket must have a bucket policy with permissions set to allow AWS Data Exchange to export data into it. For more information, see Prerequisites for S3 bucket policy permissions.

To automatically export a revision to an S3 bucket (AWS SDKs)

- 1. Create a Create_Event_Action request.
- Include the following in the request: 2.
 - Action
 - ExportRevisionToS3
 - Encryption
 - KmsKeyArn
 - Type
 - RevisionDestination
 - Bucket
 - KeyPattern
 - Event
 - RevisionPublished
 - DataSetId
 - Tags
- Modify the key pattern if necessary. The Amazon S3 object key defaults to the key pattern {Revision.CreatedAt}/{Asset.Name}.

For more information about key patterns, see Key patterns when exporting asset revisions from AWS Data Exchange.

To verify that the prerequisites for S3 bucket policy permissions are met, an object with the naming format _ADX-TEST-ACCOUNTID# is added to the S3 bucket.

AWS Data Exchange quotas

The following sections provide information about the service quotas, endpoints, guidelines for export and import jobs across AWS Regions, and constraints related to resource fields for AWS Data Exchange for an AWS account.

Service quotas

For information about service quotas, see <u>AWS Data Exchange endpoints and quotas</u> in the *AWS General Reference*.

Service endpoints

For information about service endpoints, see <u>AWS Data Exchange endpoints and quotas</u> in the *AWS General Reference*.

Export and import job guidelines

The following table provides guidelines for export and import jobs. For more information, see <u>AWS</u> Regions and data sets.

Resource, descriptor, or operation	Maximum value	Description
File size for assets imported from a signed URL	5 GB	The maximum size, in GB, of an asset that can be imported using IMPORT_ASSET_FROM_ SIGNED_URL .
File size of a cross-Region revision export to Amazon Simple Storage Service (Amazon S3)	1,000 GB	The maximum size, in GB, of a revision that can be exported to a different Region from the provider data set using an ExportRevision job.
Number of assets that can be imported from a signed URL in a single job	1	The number of assets that can be imported using a single IMPORT_ASSET_FROM_SIGNED_URL job.

Service quotas 278

Resource, descriptor, or operation	Maximum value	Description
Number of assets that can be exported to Amazon S3 in a single cross-Region ExportRevision job	10,000	The number of assets that can be exported from one Region to another from the provider data set using an ExportRevision job.
Number of assets that can be exported to Amazon S3 in a single ExportRevision job	10,000	The number of assets that can be exported to Amazon S3 using an ExportRevision job.
Number of revisions that can be exported to Amazon S3 in a single ExportRevision job	1	The number of revisions that can be exported to Amazon S3 using an ExportRevision job.
Event actions per resource	5	The maximum number of event actions per resource.
Event actions per account	50	The maximum number of event actions per account.
Payload size for APIs imported from API Gateway	10 MB	The maximum payload size for APIs that have been imported from Amazon API Gateway. For more information about quotas for Amazon API Gateway APIs, see Amazon API Gateway API Developer Guide .
SendApiAsset	10 transacti ons per second (TPS)	The default requests per second to SendApiAs set for a customer per Region.

Constraints for resource fields

The following table provides constraints related to resource fields that providers encounter in the AWS Data Exchange console when creating data sets, revisions, products, and product offers. The table also provides constraints related to resource fields that subscribers encounter when making subscription requests.

Resource	Field	Maximum length or size
Dataset	Name	256 characters
Dataset	Description	16,384 characters
Revision	Comment	128 characters
Product details	Name	72 characters
Product details	Short description	500 characters
Product details	Long description	30,000 characters
Product details	Logo	100 KB
Product offer	DSA	10 MB
Product offer	Refund policy	200 characters
Subscription request	company name	40 characters
Subscription request	name	40 characters
Subscription request	email address	100 characters
Subscription request	intended use-case	500 characters

Constraints for resource fields 280

Logging and monitoring in AWS Data Exchange

Monitoring is an important part of the well-architected nature of AWS Data Exchange. You should collect monitoring data from each part of your AWS solution so that you can more easily debug a multi-point failure, if one occurs. AWS provides several tools for monitoring your resources and activity in AWS Data Exchange so you can plan for and respond to potential incidents.

The logging of actions and events in AWS Data Exchange is accomplished through its integration with Amazon CloudWatch.

The following sections describe monitoring and logging in AWS Data Exchange:

Topics

- Monitoring AWS Data Exchange
- Amazon EventBridge events for AWS Data Exchange
- AWS User Notifications for AWS Data Exchange events
- Logging AWS Data Exchange API calls with AWS CloudTrail
- Upcoming changes in AWS Data Exchange CloudTrail logging

Monitoring AWS Data Exchange

Monitoring is an important part of maintaining the reliability, availability, and performance of AWS Data Exchange and your other AWS solutions. AWS provides the following monitoring tools to watch AWS Data Exchange, report when something is wrong, and take automatic actions when appropriate:

- Amazon CloudWatch Events delivers a near-real-time stream of system events that describe changes in AWS resources. CloudWatch Events enables automated event-driven computing. You can write rules that watch for certain events and respond with automated actions in other AWS services when these events occur. For more information, see the <u>Amazon CloudWatch Events</u> User Guide.
- Amazon CloudWatch Logs makes it possible for you to monitor, store, and access your log
 files from Amazon Elastic Compute Cloud (Amazon EC2) instances, AWS CloudTrail, and other
 sources. CloudWatch Logs can monitor information in the log files and notify you when certain
 thresholds are met. You can also archive your log data in highly durable storage. For more
 information, see the Amazon CloudWatch Logs User Guide.

Monitoring 281

CloudTrail captures API calls and related events made by or on behalf of your AWS account and
delivers the log files to an Amazon Simple Storage Service (Amazon S3) bucket that you specify.
You can identify which users and accounts called AWS, the source IP address from which the calls
were made, and when the calls occurred.

Amazon EventBridge events for AWS Data Exchange

AWS Data Exchange is integrated with Amazon EventBridge, formerly called Amazon CloudWatch Events. EventBridge is an event bus service that you can use to connect your applications with data from a variety of sources. For more information, see the Amazon EventBridge User Guide.

As a subscriber with an active subscription to a product, you receive an *event* from AWS Data Exchange every time the provider publishes new revisions or adds new data sets to an existing product. The event contains the DataSetId and the list of RevisionIds that have been published.

Providers can send notifications corresponding to data updates, data delays, schema changes, and deprecations. Providers have the option to include comments and expected actions for subscribers to follow. Subscribers receive these notifications as events in Amazon EventBridge, which they can use to build automated workflows or deliver human-readable notifications to emails and chat programs using <u>AWS User Notifications</u>.

Data product related events are emitted in the AWS Region where the provider published the data set. You must set up EventBridge rules that use these events in the same AWS Region or see Sending and receiving Amazon EventBridge events between AWS Regions for more options.

This topic provides detailed information about each event listed in the following table. The table includes events received by a subscriber when a provider adds a data set to a product, adds a revision to a product, revokes a revision to a product, or removes access to a product.

Actions	Event received	Related topic
Adds a file-based data set to a product and publishes it	Data Sets Published To Product	the section called "Events for adding file-based data sets"
Adds an Amazon S3 data access data set to a product and publishes it	Amazon S3 Data Access Data Sets Published To Product	the section called "Events for adding Amazon S3 data access data sets"

Actions	Event received	Related topic
Adds an AWS Lake Formation data permission data set and publishes it	AWS Lake Formation Data Permission Data Set Published To Product	the section called "Events for adding AWS Lake Formation data permission data sets"
Adds an Amazon Redshift data set to a product and publishes it	Redshift Data Shares Data Sets Published To Product	Events for adding Amazon Redshift datashare data sets
Adds an Amazon API Gateway data set to a product and publishes it	API Gateway API Data Sets Published To Product	the section called "Events for adding Amazon API Gateway API data sets "
Adds a file-based data set revision to a product and publishes it	Revision Published To Data Set	Events for adding revisions
Adds an Amazon S3 data access data set revision to a product and publishes it	Revision Published to Amazon S3 Data Access Data Set	the section called "Events for adding Amazon S3 data access data set revisions"
Adds an AWS Lake Formation data permission data set revision to a product and publishes it	Revision Published To Lake Formation Data Permission Data Set	the section called "Events for adding AWS Lake Formation data permission data set revisions (Preview)"
Adds an Amazon Redshift datashare data set revision to a product and publishes it	Revision Published To Redshift Data Shares Data Set	Events for adding Amazon Redshift datashare data set revisions
Adds an Amazon API Gateway data set revision to a product and publishes it	Revision Published To API Gateway API Data Set	the section called "Events for adding Amazon API Gateway API data set revisions"
Revokes revision to a product	Revision Revoked	Events for revoking revisions

Actions	Event received	Related topic
Takes an action on their Amazon Redshift resources that <i>might</i> remove access from a subscriber	Action Performed On Redshift Data Share By Provider	Events for an action performed on an Amazon Redshift resource
Takes an action on their Amazon Redshift resources that removes access from a subscriber	Redshift Data Share Access Lost	Events for losing access to an Amazon Redshift datashare
Sends a notification for a data update	Data Updated in Data Set	Events for a provider- generated notification of a data update
Sends a notification for a schema change	Schema Change Planned for Data Set	Events for a provider- generated notification of a schema change
Sends a notification for a data delay	Data Set Update Delayed	Events for a provider- generated notification of a data delay
Sends a notification for a data deprecation	Deprecation Planned for Data Set	Events for a provider- generated notification of a data deprecation
Sends an event when a data consumer accepts a data grant	Data Grant Accepted	Events for accepting a data grant
Sends an event when a data producer extends a data grant	Data Grant Extended	Events for extending data grants
Sends an event when a data producer revokes a data grant	Data Grant Revoked	Events for revoking a data grant

Actions	Event received	Related topic
Auto-export job completed	Auto-export Job Completed	Events for an auto-export job completed
Auto-export job failed	Auto-export Job Failed	Events for an auto-export job failed

Note

AWS Data Exchange emits events on a best effort basis. For more information about event delivery, see Events from AWS services.

Events for adding file-based data sets

When a provider adds file-based data sets to a product and publishes it, the subscriber receives an event with the Data Sets Published To Product detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Data Sets Published To Product",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2020-07-29T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "prod-uEXAMPLEabc1d"
    ],
    "detail": {
        "DataSetIds": [
            "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
            "5bqd734EXAMPLE100f7qdd9EXAMPLEe9"
        ],
        "DataSets": [
               "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
```

```
"Name": "Data_Set_Hello_World_One"

},
{
    "Id": "5bgd734EXAMPLE100f7gdd9EXAMPLEe9",
    "Name": "Data_Set_Hello_World_Two"
}
],
"Product":
{
    "Id": "prod-uEXAMPLEabc1d",
    "Name": "Product_Hello_World"
}
}
```

Events for adding Amazon S3 data access data sets

When a provider adds an Amazon S3 data access data set to a product and publishes it, the subscriber receives an event with the following detail type: Amazon S3 Data Access Data Set(s) Published To Product.

```
{
 "version": "0",
 "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
 "detail-type": "S3 Data Access Data Set(s) Published to Product",
"source": "aws.dataexchange",
"account": "123456789012",
 "time": "2020-07-29T18:24:04Z",
"region": "us-east-1",
 "resources": [
  "prod-uEXAMPLEabc1d"
],
 "detail": {
 "DataSetIds": [
  "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
  "5bqd734EXAMPLE100f7qdd9EXAMPLEe9"
 ],
  "DataSets": [{
    "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
    "Name": "Data_Set_Hello_World_One"
  },
```

```
{
    "Id": "5bgd734EXAMPLE100f7gdd9EXAMPLEe9",
    "Name": "Data_Set_Hello_World_Two"
}
],
    "Product": {
    "Id": "prod-uEXAMPLEabc1d",
    "Name": "Product_Hello_World"
}
}
```

Events for adding AWS Lake Formation data permission data sets

When a provider adds an AWS Lake Formation data permission data set to a product and publishes it, the subscriber receives an event with the Lake Formation Data Permission Data Sets Published To Product detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Lake Formation Data Permission Data Sets Published To Product",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2021-12-15T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "prod-uEXAMPLEabc1d"
    ],
    "detail": {
        "DataSetIds": [
            "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
            "5bqd734EXAMPLE100f7qdd9EXAMPLEe9"
        ],
        "DataSets": [
            {
                "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                "Name": "Data_Set_Hello_World_One"
            },
                "Id": "5bgd734EXAMPLE100f7gdd9EXAMPLEe9",
```

```
"Name": "Data_Set_Hello_World_Two"
}
],

"Product": {
    "Id": "prod-uEXAMPLEabc1d",
    "Name": "Product_Hello_World"
}
}
```

Events for adding Amazon Redshift datashare data sets

When a provider adds an Amazon Redshift datashare data set to a product and publishes it, the subscriber receives an event with the Redshift Data Shares Data Sets Published To Product detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Redshift Data Shares Data Sets Published To Product",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2021-12-15T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "prod-uEXAMPLEabc1d"
    ],
    "detail": {
        "DataSetIds": [
            "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
            "5bqd734EXAMPLE100f7qdd9EXAMPLEe9"
        ],
        "DataSets": [
            {
               "Id" : "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
               "Name": "Data_Set_Hello_World_One"
            },
               "Id" : "5bqd734EXAMPLE100f7qdd9EXAMPLEe9",
               "Name": "Data_Set_Hello_World_Two"
            }
```

```
],
    "Product":
{
        "Id" : "prod-uEXAMPLEabc1d",
        "Name": "Product_Hello_World"
}
```

Events for adding Amazon API Gateway API data sets

When a provider adds an Amazon API Gateway API data set to a product and publishes it, the subscriber receives an event with the Amazon API Gateway Data Sets Published To Product detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "API Gateway API Data Sets Published To Product",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2021-12-15T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "prod-uEXAMPLEabc1d"
    ],
    "detail": {
        "DataSetIds": [
            "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
            "5bqd734EXAMPLE100f7qdd9EXAMPLEe9"
        ],
        "DataSets": [
            {
                "Id" : "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                "Name": "Data_Set_Hello_World_One"
            },
                "Id" : "5bqd734EXAMPLE100f7qdd9EXAMPLEe9",
                "Name": "Data_Set_Hello_World_Two"
            }
```

```
],
"Product": {
    "Id" : "prod-uEXAMPLEabc1d",
    "Name": "Product_Hello_World"
    }
}
```

Events for adding revisions

When a provider adds a data set to a product and publishes it, the subscriber receives an event with the Revision Published To Data Set detail type.

The following example shows the event body for the detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Revision Published To Data Set",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2020-07-29T04:16:28Z",
    "region": "us-east-1",
    "resources": [
        "aae4c2cdEXAMPLE54f9369dEXAMPLE66"
    ],
    "detail": {
        "RevisionIds": [
            "3afc623EXAMPLE099e6fcc8EXAMPLEe7"
        ],
        "Revisions": [
            {
                "Id" : "3afc623EXAMPLE099e6fcc8EXAMPLEe7",
                "Comment": "Revision_Comment_One"
            }
         ],
        "DataSets": [
            {
                "Id" : "aae4c2cdEXAMPLE54f9369dEXAMPLE66",
                "Name": "Data_Set_Hello_World_One"
            }
         ],
        "Product": {
```

Events for adding revisions 290

```
"Id" : "prod-uEXAMPLEabc1d",

"Name": "Product_Hello_World"

}
}
```

Events for adding Amazon S3 data access data set revisions

When a provider adds an Amazon S3 data access data set revision to a product and publishes it, the subscriber receives an event with the Revision Published To Amazon S3 Data Access Data Set detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Revision Published to S3 Data Access Data Set(s)",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2020-07-29T04:16:28Z",
    "region": "us-east-1",
    "resources": [
        "aae4c2cdEXAMPLE54f9369dEXAMPLE66"
    ],
    "detail": {
        "RevisionIds": [
            "3afc623EXAMPLE099e6fcc8EXAMPLEe7"
        ],
        "Revisions": [
            {
                "Id": "3afc623EXAMPLE099e6fcc8EXAMPLEe7",
                "Comment": "Revision_Comment_One"
            }
         ],
        "DataSets": [
            {
                "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                "Name": "Data_Set_Hello_World_One"
            }
         ],
        "Product": {
            "Id" : "prod-uEXAMPLEabc1d",
```

```
"Name": "Product_Hello_World"
}
}
```

Events for adding AWS Lake Formation data permission data set revisions (Preview)

When a provider adds an AWS Lake Formation data permission data set revision to a product and publishes it, the subscriber receives an event with the Revision Published to Lake Formation Data Permission Data Set detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Revision Published to Lake Formation Data Permission Data Set",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2021-12-15T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "prod-uEXAMPLEabc1d"
    ],
    "detail": {
        "DataSetIds": [
            "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
            "5bqd734EXAMPLE100f7qdd9EXAMPLEe9"
        ],
        "DataSets": [
            {
                "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                "Name": "Data_Set_Hello_World_One"
            },
            {
                "Id": "5bqd734EXAMPLE100f7qdd9EXAMPLEe9",
                "Name": "Data_Set_Hello_World_Two"
            }
        ],
        "Product": {
            "Id": "prod-uEXAMPLEabc1d",
```

```
"Name": "Product_Hello_World"
}
}
```

Events for adding Amazon Redshift datashare data set revisions

When a provider adds an Amazon Redshift datashare data set revision to a product and publishes it, the subscriber receives an event with the Revision Published To Redshift Data Shares Data Set detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Revision Published To Redshift Data Shares Data Set",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2021-12-15T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "aae4c2cdEXAMPLE54f9369dEXAMPLE66"
    ٦,
    "detail": {
        "RevisionIds": [
            "3afc623EXAMPLE099e6fcc8EXAMPLEe7"
        ],
        "Revisions": [
            {
                "Id" : "3afc623EXAMPLE099e6fcc8EXAMPLEe7",
                "Comment": "Revision_Comment_One,"
            }
         ],
        "DataSets": [
            {
                "Id" : "aae4c2cdEXAMPLE54f9369dEXAMPLE66",
                "Name": "Data_Set_Hello_World_One"
            }
         ],
        "Product": {
            "Id" : "prod-uEXAMPLEabc1d",
            "Name": "Product_Hello_World"
```

```
}
}
```

Events for adding Amazon API Gateway API data set revisions

When a provider adds an Amazon API Gateway API data set revision to a product and publishes it, the subscriber receives an event with the Revision Published To API Gateway Data Set detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Revision Published To API Gateway API Data Set",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2021-12-15T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "aae4c2cdEXAMPLE54f9369dEXAMPLE66"
    ],
    "detail": {
        "RevisionIds": [
            "3afc623EXAMPLE099e6fcc8EXAMPLEe7"
        ],
        "Revisions": [
            {
                "Id": "3afc623EXAMPLE099e6fcc8EXAMPLEe7",
                "Comment": "Revision_Comment_One"
            }
         ],
        "DataSets": [
            {
                "Id" : "aae4c2cdEXAMPLE54f9369dEXAMPLE66",
                "Name": "Data_Set_Hello_World_One"
            }
         ],
        "Product": {
            "Id" : "prod-uEXAMPLEabc1d",
            "Name": "Product_Hello_World"
```

}

The following table describes the API Gateway API data set revision error codes.

Error code	Message	Description
CLUSTER_DELETED	The datashare is unavailable because the provider deleted their cluster. Please contact the provider for more information.	This message is sent when the datashare is no longer available because the provider deleted the cluster containing the datashare.
CLUSTER_ENCRYPTION _DISABLED	The datashare is unavailable because the provider disabled encryption on their cluster. Please contact the provider for more information.	This message is sent when the datashare is no longer available because the provider disabled encryption on their cluster. To use a datashare, both the provider and the subscriber must have encryption enabled.
DATASHARE_DELETED	The datashare is unavailable because the provider deleted the datashare. Please contact the provider for more information.	This message is sent when the datashare is no longer available because the provider deleted it. The provider must create a new datashare so that you can regain access to the data.
DATASHARE_DEAUTHOR IZED	The datashare is unavailable because the provider de-author ized the datashare . Please contact the provider for more information.	This message is sent when the datashare is no longer available because the provider reauthorized the datashare. The provider must create a new datashare so

Error code	Message	Description
		that you can regain access to the data.
DATASHARE_PUBLIC_C ONSUMER_BLOCKED	You cannot access a non-publicly accessible datashare from a publicly accessible cluster. You must turn off public accessibility on your cluster to access this datashare . Please contact your provider for more information.	This message is sent when a provider sets the Publicly accessible option to Disable on the cluster that contains their datashare. If the subscriber's cluster has the Publicly accessible option set to Disable, it will not affect their ability to access the datashare. For the subscriber to access the datashare, either the subscriber must set the Publicly accessible option to Disable on their cluster, or the provider must set the Publicly accessible option to Enable on their cluster. Disable on the cluster that contains their datashare. If the subscriber's cluster has the Publicly accessible option set to Disable, it will not affect their ability to access the datashare. For the subscriber to access the datashare, either the subscriber must set the Publicly accessible option to Disable on their cluster, or the provider must set the Publicly accessible option to Enable on their cluster, or the provider must set the Publicly accessible option to Enable on their cluster.

Events for revoking revisions

When a provider revokes a revision to a product and publishes it, the subscriber receives an event with the Revision Revoked detail type.

The following example shows the event body for the detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Revision Revoked",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2022-02-17T21:25:06Z",
    "region": "us-east-1",
    "resources": [
        "aae4c2cdEXAMPLE54f9369dEXAMPLE66"
    ],
    "detail": {
        "RevisionIds": [
            "3afc623EXAMPLE099e6fcc8EXAMPLEe7"
        "RevocationComment": "example revocation comment",
        "Revisions": [
            {
                "Id": "3afc623EXAMPLE099e6fcc8EXAMPLEe7",
                "Comment": "Revision_Comment_One"
            }
         ],
        "DataSets": [
            {
                "Id" : "aae4c2cdEXAMPLE54f9369dEXAMPLE66",
                "Name": "Data_Set_Hello_World_One"
            }
         ],
        "Product": {
            "Id" : "prod-uEXAMPLEabc1d",
            "Name": "Product_Hello_World"
        }
    }
}
```

Events for revoking revisions 297

Events for an action performed on an Amazon Redshift resource

When a provider takes an action on their Amazon Redshift resources that *might* remove access from a subscriber, the subscriber receives an event with the Action Performed On Redshift Data Share By Provider detail type.

For example, if a provider changes the data share's public accessibility setting from true to false, the subscriber receives an event.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Action Performed On Redshift Data Share By Provider",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2021-12-15T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:redshift:us-east-1:098765432123:datashare:01234567-2590-7654-1234-
f57ea0081234/test_data_share"
    ],
    "detail": {
        "Message": "This is an example message which explains why you may have lost
 access.",
        "AssociatedProducts": [
            {
                "ProductId": "aae4c2cdEXAMPLE54f9369dEXAMPLE66",
                "DataSetIds": [
                    "4afc623EXAMPLE099e6fcc8EXAMPLEe8"
                ],
                "DataSets": [
                    {
                        "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                        "Name": "Data_Set_Hello_World_One"
                    }
                ],
                "Product": {
                    "Id": "aae4c2cdEXAMPLE54f9369dEXAMPLE66",
                    "Name": "Product_Hello_World"
```

```
}
}
```

Events for losing access to an Amazon Redshift datashare

When a provider takes an action on their Amazon Redshift resources that removes access from a subscriber, the subscriber receives an event with the Redshift Data Share Access Lost detail type.

For example, if a provider deletes an Amazon Redshift datashare or deletes a cluster, the subscriber receives an event.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Redshift Data Share Access Lost",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2021-12-15T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:redshift:us-east-1:098765432123:datashare:01234567-2590-7654-1234-
f57ea0081234/test_data_share"
    ],
    "detail": {
        "Message": "This is an example message which explains why you may have lost
 access.",
        "AssociatedProducts": [
                "ProductId": "aae4c2cdEXAMPLE54f9369dEXAMPLE66",
                "DataSetIds": [
                    "4afc623EXAMPLE099e6fcc8EXAMPLEe8"
                ],
                "DataSets": [
                    {
                        "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                        "Name": "Data_Set_Hello_World_One"
                    }
                ],
                "Product": {
```

```
"Id" : "prod-uEXAMPLEabc1d",

"Name": "Product_Hello_World"

}

}

}
```

Events for an auto-export job completed

After an auto-export job moves all the data in a newly published File data set revision to the subscriber's chosen Amazon S3 bucket, the subscriber receives an event with the Auto-export Job Completed detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Auto-export Job Completed",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2020-07-29T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "prod-uEXAMPLEabc1d"
    ],
    "detail": {
        "RevisionIds": [
            "3afc623EXAMPLE099e6fcc8EXAMPLEe7"
        ],
        "Revisions": [{
            "Id": "3afc623EXAMPLE099e6fcc8EXAMPLEe7",
            "Comment": "Revision_Comment_One"
        }],
        "DataSetIds": [
            "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
        ],
        "DataSets": [{
            "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
            "Name": "Data_Set_Hello_World_One"
        }, ],
        "Product": {
```

```
"Id": "prod-uEXAMPLEabc1d",
}
}
```

Events for an auto-export job failed

When an auto-export job fails, the subscriber receives an event with the Auto-export Job Failed detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Auto-Export job failed",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2020-07-29T18:24:04Z",
    "region": "us-east-1",
    "resources": [
        "prod-uEXAMPLEabc1d"
    ],
    "detail": {
        "RevisionIds": [
            "3afc623EXAMPLE099e6fcc8EXAMPLEe7"
        ],
        "Revisions": [{
            "Id": "3afc623EXAMPLE099e6fcc8EXAMPLEe7",
            "Comment": "Revision_Comment_One"
        }],
        "DataSetIds": [
            "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
            "5bqd734EXAMPLE100f7qdd9EXAMPLEe9"
        ],
        "DataSets": [{
                "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                "Name": "Data_Set_Hello_World_One"
            },
                "Id": "5bqd734EXAMPLE100f7qdd9EXAMPLEe9",
                "Name": "Data_Set_Hello_World_Two"
            }
```

```
],
    "Product": {
        "Id": "prod-uEXAMPLEabc1d",
    }
}
```

Events for a provider-generated notification of a data update

When a provider sends a notification for a data update, the subscriber receives an event with the Data Updated in Data Set detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Data Updated in Data Set",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2023-08-21T10:29:48Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:dataexchange:us-east-1::data-sets/b5538f9f45e4613d448eb9eEXAMPLEc6"
    ],
    "detail": {
        "DataSet": {
            "Id": "b5538f9f45e4613d448eb9eEXAMPLEc6",
            "Name": "Example Data Set",
            "AssetType": "S3_DATA_ACCESS"
        },
        "Product": {
            "Id": "prod-7ip6EXAMPLEhs",
            "Name": "Example Data Product",
            "ProviderContact": "no-reply@marketplace.aws"
        },
        "Notification": {
            "Comment": "This is a test DATA_UPDATE notification.",
            "Type": "DATA_UPDATE",
            "Details": {
                "DataUpdate": {
                    "DataUpdatedAt": "2023-07-12T00:00:00Z"
```

Events for a provider-generated notification of a schema change

When a provider sends a notification for a schema change, the subscriber receives an event with the Schema Change Planned for Data Set detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Schema Change Planned for Data Set",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2023-08-21T10:29:48Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:dataexchange:us-east-1::data-sets/b5538f9f45e4613d448eb9eEXAMPLEc6"
    ],
    "detail": {
        "DataSet": {
            "Id": "b5538f9f45e4613d448eb9eEXAMPLEc6",
            "Name": "Example Data Set",
            "AssetType": "S3_DATA_ACCESS"
        },
        "Product": {
            "Id": "prod-7ip6EXAMPLEhs",
            "Name": "Example Data Product",
```

```
"ProviderContact": "no-reply@marketplace.aws"
        },
        "Notification": {
            "Comment": "This is a test SCHEMA_CHANGE notification.",
            "Type": "SCHEMA_CHANGE",
            "Details": {
                "SchemaChange": {
                    "Changes": [{
                             "Type": "ADD",
                             "Description": "This object is being added to the bucket,
 or a field is being added to the object.",
                             "Name": "KeyA"
                        },
                        {
                             "Type": "REMOVE",
                             "Description": "This object is being removed from the
 bucket or a field is being removed from the object.",
                             "Name": "KeyB"
                        },
                        {
                             "Type": "MODIFY",
                             "Description": "The usage or meaning of this key prefix is
 changing, or something is changing about every file under this key prefix.",
                             "Name": "KeyPrefix"
                        }
                    ],
                    "SchemaChangeAt": "2023-09-08T13:46:01Z"
                }
            },
            "Scope": {
                "S3DataAccesses": [{
                    "KeyPrefixes": [
                         "KeyPrefix"
                    ],
                    "Keys": [
                        "KeyA",
                         "KeyB"
                    ]
                }]
            }
        }
    }
}
```

Events for a provider-generated notification of a data delay

When a provider sends a notification for a data delay, the subscriber receives an event with the following detail type: **Data Set Update Delayed**.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Data Set Update Delayed",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2023-08-21T10:29:48Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:dataexchange:us-east-1::data-sets/b5538f9f45e4613d448eb9eEXAMPLEc6"
    ],
    "detail": {
        "DataSet": {
            "Id": "b5538f9f45e4613d448eb9eEXAMPLEc6",
            "Name": "Example Data Set",
            "AssetType": "S3_DATA_ACCESS"
        },
        "Product": {
            "Id": "prod-7ip6EXAMPLEhs",
            "Name": "Example Data Product",
            "ProviderContact": "no-reply@marketplace.aws"
        },
        "Notification": {
            "Comment": "This is a test DATA_DELAY notification.",
            "Type": "DATA_DELAY",
            "Scope": {
                "S3DataAccesses": [{
                    "KeyPrefixes": [
                         "KeyPrefix"
                    ],
                    "Keys": [
                         "KeyA",
                         "KeyB"
                    ]
                }]
            }
```

```
}
}
}
```

Events for a provider-generated notification of a data deprecation

When a provider sends a notification for a data deprecation, the subscriber receives an event with the Deprecation Planned for Data Set detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Deprecation Planned for Data Set",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2023-08-21T10:29:48Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:dataexchange:us-east-1::data-sets/b5538f9f45e4613d448eb9eEXAMPLEc6"
    ],
    "detail": {
        "DataSet": {
            "Id": "b5538f9f45e4613d448eb9eEXAMPLEc6",
            "Name": "Example Data Set",
            "AssetType": "S3_DATA_ACCESS"
        },
        "Product": {
            "Id": "prod-7ip6EXAMPLEhs",
            "Name": "Example Data Product",
            "ProviderContact": "no-reply@marketplace.aws"
        },
        "Notification": {
            "Comment": "This is a test DEPRECATION notification.",
            "Type": "DEPRECATION",
            "Details": {
                "Deprecation": {
                    "DeprecationAt": "2023-09-08T13:46:01Z"
            },
            "Scope": {
                "S3DataAccesses": [{
```

Events for accepting a data grant

When a data consumer accepts a data grant, the data owner receives an event with the Data Grant Accepted detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Data Grant Accepted",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2022-02-17T21:25:06Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:dataexchange:us-east-1::data-sets/4afc623EXAMPLE099e6fcc8EXAMPLEe8"
    ],
    "detail": {
        "DataSets": [
            {
                "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                "Name": "Data_Set_Hello_World_One"
            }
         ],
        "DataGrant": {
            "Arn" : "arn:aws:dataexchange:us-east-1:123456789012:data-
grants/4afc623EXAMPLE099e6fcc8EXAMPLEe9",
            "Name": "DataGrant_Hello_World"
```

```
}
}
```

Events for extending data grants

When a data owner extends a data grant, the data consumer receives an event with the Data Grant Extended detail type.

The following example shows the event body for the detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Data Grant Extended",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2022-02-17T21:25:06Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:dataexchange:us-east-1::data-sets/4afc623EXAMPLE099e6fcc8EXAMPLEe8"
    ],
    "detail": {
        "DataSets": [
            {
                "Id": "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                "Name": "Data_Set_Hello_World_One"
            }
         ],
         "DataGrant": {
            "Arn" : "arn:aws:dataexchange:us-east-1:123456789012:data-
grants/4afc623EXAMPLE099e6fcc8EXAMPLEe9",
            "Name": "DataGrant_Hello_World"
        }
    }
}
```

Events for revoking a data grant

When a data owner revokes a data grant, the data consumer receives an event with the Data Grant Revoked detail type.

```
{
    "version": "0",
    "id": "dc529cb6-2e23-4c5f-d020-EXAMPLE92231",
    "detail-type": "Data Grant Revoked",
    "source": "aws.dataexchange",
    "account": "123456789012",
    "time": "2022-02-17T21:25:06Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:dataexchange:us-east-1::data-sets/4afc623EXAMPLE099e6fcc8EXAMPLEe8"
    ],
    "detail": {
        "DataSets": [
            {
                "Id" : "4afc623EXAMPLE099e6fcc8EXAMPLEe8",
                "Name": "Data_Set_Hello_World_One"
            }
         ],
         "DataGrant": {
            "Arn" : "arn:aws:dataexchange:us-east-1:123456789012:data-
grants/4afc623EXAMPLE099e6fcc8EXAMPLEe9",
            "Name": "DataGrant_Hello_World"
        }
    }
}
```

AWS User Notifications for AWS Data Exchange events

You can use <u>AWS User Notifications</u> to set up delivery channels that notify you about AWS Data Exchange events. You receive a notification when an event matches a specified rule. You can receive notifications for events through multiple channels, including email, Amazon Q Developer in chat applications chat notifications, or AWS Console Mobile Application push notifications. You can also see notifications using the Console Notifications Center in the AWS User Notifications console. AWS User Notifications supports aggregation, which can reduce the number of notifications you receive during specific events. For more information, see the AWS User Notifications User Guide.

To use AWS User Notifications, you must have the correct AWS Identity and Access Management (IAM) permissions. For more information about configuring your IAM permissions, see <u>Configuring</u> AWS User Notifications in the AWS User Notifications User Guide.

The following table provides more information about the notifications that you can configure for AWS Data Exchange events using AWS User Notifications.

Actions	Notification received by subscriber
Adds a file-based data set to a product and publishes it	Data Sets Published To Product
Adds an Amazon Redshift data set to a product and publishes it	Redshift Data Shares Data Sets Published To Product
Adds a file-based data set revision to a product and publishes it	Revision Published To Data Set
Revokes revision to a product	Revision Revoked
Adds an Amazon Redshift data set revision to a product and publishes it	Revision Published To Redshift Data Shares Data Set
Takes an action on Amazon Redshift resources that might remove access from a subscriber	Action Performed On Redshift Data Share By Provider
Takes an action on Amazon Redshift resources that removes access from a subscriber	Redshift Data Share Access Lost
Adds an Amazon API Gateway data set to a product and publishes it	API Gateway API Data Sets Published To Product
Adds an Amazon API Gateway data set revision to a product and publishes it	Revision Published To API Gateway API Data Set
Adds an AWS Lake Formation data set to a product and publishes it (Preview)	Lake Formation Data Permission Data Sets Published To Product (Preview)
Adds an AWS Lake Formation data set revision to a product and publishes it (Preview)	Revision Published To Lake Formation Data Permission Data Set (Preview)

Actions	Notification received by subscriber
Auto-export job completed	Auto-export Job Completed
Auto-export job failed	Auto-export Job Failed
Sends notification for a data update	Data Updated in Data Set
Sends notification for a schema change	Schema Change Planned for Data Set
Sends notification for a data delay	Data Set Update Delayed
Sends notification for a data deprecation	Deprecation Planned for Data Set

Logging AWS Data Exchange API calls with AWS CloudTrail

AWS Data Exchange is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in AWS Data Exchange. AWS CloudTrail captures all calls to AWS Data Exchange API operations as events, including calls from the AWS Data Exchange console and from code calls to the AWS Data Exchange API operations.

If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon Simple Storage Service (Amazon S3) bucket, including events for AWS Data Exchange. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in Event **history**. Using the information collected by CloudTrail, you can determine the request that was made to AWS Data Exchange, the IP address from which the request was made, who made the request, when it was made, and other details.



Important

Some actions you can take are console-only actions. There is no corresponding API in the AWS SDK or AWS Command Line Interface (AWS CLI). These are actions that rely on AWS Marketplace functionality, such as publishing or subscribing to a product. AWS Data Exchange provides CloudTrail logs for a subset of these console-only actions. See the following list of console-only actions for which CloudTrail logs are provided. For more information, see What Is AWS CloudTrail?

Logging AWS Data Exchange API calls with AWS CloudTrail

In addition to CloudTrail events for all the <u>AWS Data Exchange APIs</u> and corresponding console actions, AWS Data Exchange also provides CloudTrail trails for a subset of the AWS Marketplace-backed console-only actions. AWS Data Exchange provides a CloudTrail log for the following console-only actions:

Subscriber actions

- Subscribe to a product
- Send subscription verification request
- Enable subscription auto-renewal
- Disable subscription auto-renewal
- Cancel subscription verification request
- List active subscriptions
- Check subscription status
- List targeted private offers
- · View details of a specific product and offer
- View details of a specific subscription
- · View details of a specific subscription verification request

Provider actions

- Publish a product
- Unpublish a product
- Edit a product
- Create custom offer
- Edit custom offer
- Approve subscription verification request
- Decline subscription verification request
- Delete subscriber contact information
- List subscription verification requests
- View details of a specific subscription verification request
- · Send a notification for a data set

AWS Data Exchange information in CloudTrail

CloudTrail is enabled when you create your AWS account. When activity occurs in AWS Data Exchange, the activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see <u>Viewing Events</u> with <u>CloudTrail Event History</u> in the <u>AWS CloudTrail User Guide</u>.

For an ongoing record of events in your AWS account, including events for AWS Data Exchange, create a trail. CloudTrail uses this trail to deliver log files to an S3 bucket. By default, when you use the console to create a trail, it applies to all AWS Regions. The trail logs events from all Regions and delivers the log files to the S3 bucket that you specify. You can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see:

- Overview for Creating a Trail
- CloudTrail Supported Services and Integrations
- Configuring Amazon SNS Notifications for CloudTrail
- Receiving CloudTrail Log Files from Multiple Regions
- Receiving CloudTrail Log Files from Multiple Accounts

All AWS Data Exchange actions are documented in the AWS Data Exchange API Reference. Every AWS Data Exchange action, except for SendAPIAsset, is logged by CloudTrail. For example, calls to the CreateDataSet, StartImportAssetsFromS3Workflow, and ListRevisionAssets API operations generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM)
 credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see CloudTrail userIdentity Element.

Understanding AWS Data Exchange log file entries

A trail is a configuration that makes it possible to deliver events as log files to an S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files are not an ordered stack trace of the public API calls, so they do not appear in any order.

Note

These examples have been formatted to improve readability. In a CloudTrail log file, all entries and events are concatenated into a single line. This example has been limited to a single AWS Data Exchange entry. In a real CloudTrail log file, you see entries and events from multiple AWS services.

The following example shows a CloudTrail log entry that demonstrates the CreateDataSet operation.

```
{
    "eventVersion": "1.05",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE:account_name",
        "arn": "arn:aws:sts::123456789012:user/Mary_Major",
        "accountId": "123456789012",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2018-06-20T18:32:25Z"
            },
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AIDACKCEVSQ6C2EXAMPLE",
                "arn": "arn:aws:iam::123456789012:role/Admin",
                "accountId": "123456789012",
                "userName": "username"
            }
        }
    },
```

```
"eventTime": "2018-06-20T19:04:36Z",
    "eventSource": "dataexchange.amazonaws.com",
    "eventName": "CreateDataSet",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "203.0.113.12",
    "userAgent": "console.amazonaws.com",
    "requestParameters": {
        "Name": "MyDataSet",
        "AssetType": "S3_SNAPSHOT",
        "Description": "This is my data set"
    },
    "responseElements": {
        "Origin": "OWNED",
        "AssetType": "S3_SNAPSHOT",
        "Name": "MyDataSet",
        "CreatedAt": 1726255485679,
        "UpdatedAt": 1726255485679,
        "Arn": "arn:aws:dataexchange:us-east-1:123456789012:data-sets/
DataSetIdentifier",
        "Id": "DataSetIdentifier",
        "Description": "This is my data set"
    },
    "requestID": "cb8c167e-EXAMPLE",
    "eventID": "e3c6f4ce-EXAMPLE",
    "readOnly": false,
    "eventType": "AwsApiCall",
    "recipientAccountId": "123456789012"
}>
```

Upcoming changes in AWS Data Exchange CloudTrail logging

This section summarizes the upcoming changes for logging API calls in AWS CloudTrail for AWS Data Exchange. The effective date for the change is on or after September 1, 2023. We recommend reviewing your CloudTrail usage to make sure this change will not impact your monitoring, analysis, or auditing. For questions or concerns, please send an email message to Support.

Customer persona	Event description	Previous eventName	New eventName	Previous eventSource	New eventSource
Subscriber	Subscribe to a product	Subscribe	CreateAgr eementReq uest and AcceptAgr eementReq uest	aws-marke tplace.am azonaws.com	agreement -marketpl ace.amazo naws.com
Subscriber	Send subscription verification request	Subscribe	CreateAgr eementReq uest and AcceptAgr eementReq uest	aws-marke tplace.am azonaws.com	agreement -marketpl ace.amazo naws.com
Subscriber	Enable subscription auto-renewal	Subscribe	CreateAgr eementReq uest and AcceptAgr eementReq uest	aws-marke tplace.am azonaws.com	agreement -marketpl ace.amazo naws.com
Subscriber	Disable subscription auto-renewal	Unsubscri be	CreateAgr eementReq uest and AcceptAgr eementReq uest	aws-marke tplace.am azonaws.com	agreement -marketpl ace.amazo naws.com
Subscriber	Cancel subscription verification request	CancelAgr eementReq uest	CancelAgr eementReq uest	aws-marke tplace.am azonaws.com	agreement -marketpl ace.amazo naws.com

Customer persona	Event description	Previous eventName	New eventName	Previous eventSource	New eventSource
Provider	Publish a product	StartChan geSet	StartChan geSet	aws-marke tplace.am azonaws.com	marketpla cecatalog .amazonaw s.com
Provider	Edit a product	StartChan geSet	StartChan geSet	aws-marke tplace.am azonaws.com	marketpla cecatalog .amazonaw s.com
Provider	Unpublish a product	StartChan geSet	StartChan geSet	aws-marke tplace.am azonaws.com	marketpla cecatalog .amazonaw s.com
Provider	Create custom offer	StartChan geSet	StartChan geSet	aws-marke tplace.am azonaws.com	marketpla cecatalog .amazonaw s.com
Provider	Edit custom offer	StartChan geSet	StartChan geSet	aws-marke tplace.am azonaws.com	marketpla cecatalog .amazonaw s.com
Provider	Approve subscription verification request	AcceptAgr eementApp rovalRequ est	AcceptAgr eementApp rovalRequ est	aws-marke tplace.am azonaws.com	agreement -marketpl ace.amazo naws.com
Provider	Decline subscription verification request	RejectAgr eementApp rovalRequ est	RejectAgr eementApp rovalRequ est	aws-marke tplace.am azonaws.com	agreement -marketpl ace.amazo naws.com

Customer persona	Event description	Previous eventName	New eventName	Previous eventSource	New eventSource
Provider	Delete subscribe r contact information	UpdateAgr eementApp rovalRequ est	UpdateAgr eementApp rovalRequ est	aws-marke tplace.am azonaws.com	agreement -marketpl ace.amazo naws.com

Security

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from multiple data centers and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The <u>shared responsibility model</u> describes this as security of the cloud and security in the cloud:

- Security of the cloud AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. The effectiveness of our security is regularly tested and verified by third-party auditors as part of <u>AWS compliance programs</u>. To learn about the compliance programs that apply to AWS Data Exchange, see AWS Services in Scope by Compliance Program.
- **Security in the cloud** Your responsibility is determined by the AWS services that you use. You are also responsible for other factors, including the sensitivity of your data, your organization's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when you use AWS Data Exchange. The following topics show you how to configure AWS Data Exchange to meet your security and compliance objectives. You also learn how to use other AWS services that help you monitor and secure your AWS Data Exchange resources.

Data protection in AWS Data Exchange

The AWS <u>shared responsibility model</u> applies to data protection in AWS Data Exchange. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the <u>Data Privacy FAQ</u>. For information about data protection in Europe, see the <u>AWS Shared Responsibility Model and GDPR</u> blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

Data protection 319

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see <u>Working with CloudTrail trails</u> in the AWS CloudTrail User Guide.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-3.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with AWS Data Exchange or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

AWS Data Exchange provides the following options that you can use to help secure the content that exists in your data sets:

Topics

- Encryption at rest
- Encryption in transit
- Restrict access to content

Encryption at rest

AWS Data Exchange always encrypts all data products stored in the service at rest without requiring any additional configuration. This encryption is automatic when you use AWS Data Exchange.

Encryption at rest 320

Encryption in transit

AWS Data Exchange uses Transport Layer Security (TLS) and client-side encryption for encryption in transit. Communication with AWS Data Exchange is always done over HTTPS so your data is always encrypted in transit. This encryption is configured by default when you use AWS Data Exchange.

Restrict access to content

As a best practice, you should restrict access to the appropriate subset of users. With AWS Data Exchange, you can do this by ensuring that users, groups, and roles who use your AWS account have the right permissions. For more information about roles and policies for IAM entities, see <u>IAM</u> *User Guide*.

Key management for Amazon S3 data access

This page is specific to the Amazon S3 data access type where the provider is sharing objects encrypted using SSE-KMS. The subscriber must have a grant on the keys used for access.

If your Amazon S3 bucket contains data encrypted using AWS KMS customer managed keys, you must share these AWS KMS keys with AWS Data Exchange to configure your Amazon S3 data access data set. For more information, see the section called "Step 2: Configure Amazon S3 data access".

Topics

- Creating AWS KMS grants
- Encryption context and grant constraints
- Monitoring your AWS KMS keys in AWS Data Exchange

Creating AWS KMS grants

When you provide AWS KMS keys as part of your Amazon S3 data access data set, AWS Data Exchange creates an AWS KMS grant on each AWS KMS key shared. This grant, known as the *parent grant*, is used to give AWS Data Exchange permission to create additional AWS KMS grants for subscribers. These additional grants are known as *child grants*. Each subscriber is permitted one AWS KMS grant. Subscribers get permission to decrypt the AWS KMS key. Then, they can decrypt and use the encrypted Amazon S3 objects shared with them. For more information, see <u>Grants in AWS KMS</u> in the *AWS Key Management Service Developer Guide*.

Encryption in transit 321

AWS Data Exchange also uses the AWS KMS parent grant to manage the lifecycle of the AWS KMS grant that it creates. When a subscription ends, AWS Data Exchange retires the AWS KMS child grant created for the corresponding subscriber. If the revision is revoked, or the data set is deleted, AWS Data Exchange retires the AWS KMS parent grant. For more information about AWS KMS actions, see the AWS KMS API reference.

Encryption context and grant constraints

AWS Data Exchange uses grant constraints to permit the decrypt operation only when the request includes the specified encryption context. You can use the Amazon S3 Bucket Key feature to encrypt your Amazon S3 objects and share it with AWS Data Exchange. The bucket Amazon Resource Name (ARN) is implicitly used by Amazon S3 as the encryption context. The following example shows that AWS Data Exchange uses the bucket ARN as the grant constraint for all AWS KMS grants that it creates.

```
"Constraints": {
   "EncryptionContextSubset": "aws:s3:arn": "arn:aws:s3:::<Bucket ARN>"
   }
}
```

Monitoring your AWS KMS keys in AWS Data Exchange

When you share AWS KMS customer managed keys with AWS Data Exchange, you can use <u>AWS CloudTrail</u> to track requests that AWS Data Exchange or data subscribers send to AWS KMS. The following are examples of what your CloudTrail logs will look like for the CreateGrant and Decrypt calls to AWS KMS.

CreateGrant for parent

CreateGrant is for parent grants created by AWS Data Exchange for itself.

```
"eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AROAIGDTESTANDEXAMPLE:Provider01",
        "arn": "arn:aws:sts::<your-account-id>:assumed-role/Admin/Provider01",
        "accountId": "<your-account-id>",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
```

```
"type": "Role",
            "principalId": "AROAIGDTESTANDEXAMPLE",
            "arn": "arn:aws:iam::<your-account-id>:role/Admin/Provider01",
            "accountId": "<your-account-id>",
            "userName": "Admin"
        },
        "webIdFederationData": {},
        "attributes": {
            "creationDate": "2023-02-16T17:29:23Z",
            "mfaAuthenticated": "false"
        }
    },
    "invokedBy": "datax.amazonaws.com"
},
"eventTime": "2023-02-16T17:32:47Z",
"eventSource": "kms.amazonaws.com",
"eventName": "CreateGrant",
"awsRegion": "us-east-2",
"sourceIPAddress": "datax.amazonaws.com",
"userAgent": "datax.amazonaws.com",
"requestParameters": {
    "keyId": "<Key ARN of the Key you shared with AWS Data Exchange>",
    "operations": [
        "CreateGrant",
        "Decrypt",
        "RetireGrant"
    ],
    "granteePrincipal": "dataexchange.us-east-2.amazonaws.com",
    "retiringPrincipal": "dataexchange.us-east-2.amazonaws.com",
    "constraints": {
        "encryptionContextSubset": {
            AWS:s3:arn": "arn:aws:s3:::<Your Bucket ARN>"
        }
    }
},
"responseElements": {
    "grantId": "<KMS Grant ID of the created Grant>",
    "keyId": "<Key ARN of the Key you shared with AWS Data Exchange>"
},
"requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
"eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
"readOnly": false,
"resources": [
    {
```

CreateGrant for child

CreateGrant is for child grants created by AWS Data Exchange for subscribers.

```
{
      "eventVersion": "1.08",
      "userIdentity": {
         "type": "AWSService",
         "invokedBy": "datax.amazonaws.com"
     },
     "eventTime": "2023-02-15T23:15:49Z",
     "eventSource": "kms.amazonaws.com",
     "eventName": "CreateGrant",
     "awsRegion": "us-east-2",
     "sourceIPAddress": "datax.amazonaws.com",
     "userAgent": "datax.amazonaws.com",
     "requestParameters": {
         "keyId": "<Key ARN of the Key you shared with AWS Data Exchange>",
         "operations": [
             "Decrypt"
         ],
         "granteePrincipal": "<Subscriber's account Id>",
         "retiringPrincipal": "dataexchange.us-east-2.amazonaws.com",
         "constraints": {
             "encryptionContextSubset": {
                 "aws:s3:arn": "arn:aws:s3:::<Your Bucket ARN>"
             }
         }
     },
     "responseElements": {
         "grantId": "<KMS Grant ID of the created Grant>",
         "keyId": "<Key ARN of the Key you shared with AWS Data Exchange>"
```

```
},
     "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
     "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
     "readOnly": false,
     "resources": [
         {
             "accountId": "<Your Account Id>",
             "type": "AWS::KMS::Key",
             "ARN": "<Key ARN of the Key you shared with AWS Data Exchange>"
         }
     ],
     "eventType": "AwsApiCall",
     "managementEvent": true,
     "recipientAccountId": "<Your Account Id>",
     "sharedEventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE ",
     "eventCategory": "Management"
}
```

Decrypt

Decrypt is called by subscribers when they attempt to read the encrypted data in which they're subscribed.

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AWSAccount",
        "principalId": "AROAIGDTESTANDEXAMPLE:Subscriber01",
        "accountId": "<subscriber-account-id>",
        "invokedBy": "<subscriber's IAM identity>"
    },
    "eventTime": "2023-02-15T23:28:30Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "Decrypt",
    "awsRegion": "us-east-2",
    "sourceIPAddress": "<subscriber's IP address>",
    "userAgent": "<subscriber's user agent>",
    "requestParameters": {
        "encryptionContext": {
            "aws:s3:arn": "arn:aws:s3:::<Your Bucket ARN>"
        },
        "encryptionAlgorithm": "SYMMETRIC_DEFAULT"
    },
```

```
"responseElements": null,
    "requestID": ""ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": ""ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": true,
    "resources": [
        {
            "accountId": "<Your Account Id>",
            "type": "AWS::KMS::Key",
            "ARN": "<Key ARN of the Key you shared with AWS Data Exchange>"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "602466227860",
    "sharedEventID": "bcf4d02a-31ea-4497-9c98-4c3549f20a7b",
    "eventCategory": "Management"
}
```

Identity and access management in AWS Data Exchange

To perform any operation in AWS Data Exchange, such as creating an import job using an AWS SDK, or subscribing to a product in the AWS Data Exchange console, AWS Identity and Access Management (IAM) requires that you authenticate that you're an approved AWS user. For example, if you're using the AWS Data Exchange console, you authenticate your identity by providing your AWS sign-in credentials.

After you authenticate your identity, IAM controls your access to AWS with a defined set of permissions on a set of operations and resources. If you're an account administrator, you can use IAM to control the access of other users to the resources that are associated with your account.

Topics

- Authentication
- Access control
- AWS Data Exchange API permissions: actions and resources reference
- AWS managed policies for AWS Data Exchange

Authentication

You can access AWS with any of the following types of identities:

- AWS account root user When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see Tasks that require root user credentials in the *IAM User Guide*.
- User A <u>user</u> is an identity in your AWS account that has specific custom permissions. You can
 use your IAM credentials to sign in to secure AWS webpages like the AWS Management Console
 or the AWS Support Center.
- IAM role An IAM role is an IAM identity that you can create in your account that has specific permissions. An IAM role is similar to an IAM user in that it is an AWS identity with permissions policies that determine what the identity can and cannot do in AWS. However, instead of being uniquely associated with one person, a role is intended to be assumable by anyone who needs it. Also, a role does not have standard long-term credentials such as a password or access keys associated with it. Instead, when you assume a role, it provides you with temporary security credentials for your role session. Roles with temporary credentials are useful in the following situations:
 - **Federated user access** Instead of creating a user, you can use existing identities from AWS Directory Service, your enterprise user directory, or a web identity provider. These are known as *federated users*. AWS assigns a role to a federated user when access is requested through an identity provider. For more information about federated users, see Federated Users and Roles.
 - AWS service access A service role is an IAM role that a service assumes to perform actions in your account on your behalf. When you set up some AWS service environments, you must define a role for the service to assume. This service role must include all the permissions that are required for the service to access the AWS resources that it needs. Service roles vary from service to service, but many allow you to choose your permissions as long as you meet the documented requirements for that service. Service roles provide access only within your account and cannot be used to grant access to services in other accounts. You can create, modify, and delete a service role from within IAM. For example, you can create a role that allows Amazon Redshift to access an Amazon S3 bucket on your behalf and then load data

Authentication 327

from that bucket into an Amazon Redshift cluster. For more information, see Creating a Role to Delegate Permissions to an AWS Service.

• Applications running on Amazon EC2 – You can use an IAM role to manage temporary credentials for applications that are running on an Amazon EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys in the Amazon EC2 instance. To assign an AWS role to an Amazon EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the Amazon EC2 instance to get temporary credentials. For more information, see Using an IAM Role to Grant Permissions to Applications Running on Amazon EC2 Instances.

Access control

To create, update, delete, or list AWS Data Exchange resources, you need permissions to perform the operation and to access the corresponding resources. To perform the operation programmatically, you also need valid access keys.

Overview of managing access permissions to your AWS Data Exchange resources

Every AWS resource is owned by an AWS account, and permissions to create or access a resource are governed by permissions policies. An account administrator can attach permissions policies to users, groups, and roles. Some services (such as AWS Lambda) also support attaching permissions policies to resources.



Note

An account administrator (or administrator) is a user with administrator privileges. For more information, see IAM Best Practices.

To provide access, add permissions to your users, groups, or roles:

- Users and groups in AWS IAM Identity Center:
 - Create a permission set. Follow the instructions in Create a permission set in the AWS IAM Identity Center User Guide.
- Users managed in IAM through an identity provider:

Create a role for identity federation. Follow the instructions in <u>Create a role for a third-party</u> identity provider (federation) in the *IAM User Guide*.

- IAM users:
 - Create a role that your user can assume. Follow the instructions in <u>Create a role for an IAM user</u> in the *IAM User Guide*.
 - (Not recommended) Attach a policy directly to a user or add a user to a user group. Follow the instructions in Adding permissions to a user (console) in the *IAM User Guide*.

Topics

- AWS Data Exchange resources and operations
- Understanding resource ownership
- Managing access to resources
- Specifying policy elements: actions, effects, and principals
- Specifying conditions in a policy

AWS Data Exchange resources and operations

In AWS Data Exchange, there are two different kinds of primary resources with different control planes:

- The primary resources for AWS Data Exchange are *data sets* and *jobs*. AWS Data Exchange also supports *revisions* and *assets*.
- To facilitate transactions between providers and subscribers, AWS Data Exchange also uses AWS
 Marketplace concepts and resources, including products, offers, and subscriptions. You can use
 the AWS Marketplace Catalog API or the AWS Data Exchange console to manage your products,
 offers, subscription requests, and subscriptions.

Understanding resource ownership

The AWS account owns the resources that are created in the account, regardless of who created the resources. Specifically, the resource owner is the AWS account of the <u>principal entity</u> (that is, the AWS account root user, a user, or a role) that authenticates the resource creation request. The following examples illustrate how this works.

Resource ownership

Any IAM entity in an AWS account with the correct permissions can create AWS Data Exchange data sets. When an IAM entity creates a data set, their AWS account owns the data set. Published data products can contain data sets that are owned only by the AWS account that created them.

To subscribe to an AWS Data Exchange product, the IAM entity needs permissions to use AWS Data Exchange, in addition to the aws-marketplace:subscribe, aws-marketplace:aws-marketplace:CreateAgreementRequest, and aws-marketplace:AcceptAgreementRequest IAM permissions for AWS Marketplace (assuming they pass any related subscription verifications). As a subscriber, your account has read access to entitled data sets; however, it does not own the entitled data sets. Any entitled data sets that are exported to Amazon S3 are owned by the subscriber's AWS account.

Managing access to resources

This section discusses using IAM in the context of AWS Data Exchange. It doesn't provide detailed information about the IAM service. For complete IAM documentation, see What Is IAM? in the IAM User Guide. For information about IAM policy syntax and descriptions, see AWS Identity and Access Management Policy Reference in the IAM User Guide.

A *permissions policy* describes who has access to what. The following section explains the options for creating permissions policies.

Policies attached to an IAM identity are referred to as *identity-based* policies (IAM policies). Policies attached to a resource are referred to as *resource-based* policies. AWS Data Exchange supports only identity-based policies (IAM policies).

Topics

- Identity-based policies and permissions
- Resource-based policies

Identity-based policies and permissions

AWS Data Exchange provides a set of managed policies. For more information about them and their permissions, see AWS managed policies for AWS Data Exchange.

Amazon S3 permissions

When importing assets from Amazon S3 to AWS Data Exchange, you need permissions to write to the AWS Data Exchange service S3 buckets. Similarly, when exporting assets from AWS Data Exchange to Amazon S3, you need permissions to read from the AWS Data Exchange service S3 buckets. These permissions are included in the policies mentioned previously, but you can also create your own policy to allow just what you want your users to be able to do. You can scope these permissions to buckets that contain aws-data-exchange in their name and use the CalledVia permission to restrict the usage of the permission to requests made by AWS Data Exchange on behalf of the principal.

For example, you could create a policy to allow importing and exporting to AWS Data Exchange that includes these permissions.

JSON

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:GetObject",
            "Resource": "arn:aws:s3:::*aws-data-exchange*",
            "Condition": {
                "ForAnyValue:StringEquals": {
                    "aws:CalledVia": [
                         "dataexchange.amazonaws.com"
                    ]
                }
            }
        },
        {
            "Effect": "Allow",
            "Action": [
                "s3:PutObject",
                "s3:PutObjectAcl"
            ],
            "Resource": "arn:aws:s3:::*aws-data-exchange*",
            "Condition": {
                "ForAnyValue:StringEquals": {
                    "aws:CalledVia": [
```

These permissions allow providers to import and export with AWS Data Exchange. The policy includes the following permissions and restrictions:

- s3:PutObject and s3:PutObjectAcl These permissions are restricted only to S3 buckets that contain aws-data-exchange in their name. These permissions allows providers to write to AWS Data Exchange service buckets when importing from Amazon S3.
- **s3:GetObject** This permission is restricted to S3 buckets that contain aws-data-exchange in their name. This permission allows customers to read from AWS Data Exchange service buckets when exporting from AWS Data Exchange to Amazon S3.
- These permissions are restricted to requests made by using AWS Data Exchange with the IAM CalledVia condition. This allows the S3 PutObject permissions to only be used in the context of the AWS Data Exchange console or API.
- AWS Lake Formation and AWS Resource Access Manager (AWS RAM) To use AWS Lake
 Formation data sets you'll need to accept the AWS RAM share invitation for each net new
 provider that you have a subscription with. In order to accept the AWS RAM share invitation you
 will need to assume a role that has permission to accept a AWS RAM share invitation. To learn
 more about how AWS managed policies for AWS RAM, see Managed policies for AWS RAM.
- To create AWS Lake Formation data sets, you'll need to create the data set with an assumed role that allows IAM to pass a role to AWS Data Exchange. This will allow AWS Data Exchange to grant and revoke permissions to Lake Formation resources on your behalf. See an example policy below:

```
{
    "Effect": "Allow",
    "Action": "iam:PassRole",
    "Resource": "*",
    "Condition": {
        "StringEquals": {
            "iam:PassedToService": "dataexchange.amazonaws.com"
        }
}
```

} }



Note

Your users may also need additional permissions to read to or write from your own S3 buckets and objects that are not covered in this example.

For more information about users, groups, roles, and permissions, see Identities (Users, Groups, and Roles) in the IAM User Guide.

Resource-based policies

AWS Data Exchange does not support resource-based policies.

Other services, such as Amazon S3, do support resource-based permissions policies. For example, you can attach a policy to an S3 bucket to manage access permissions to that bucket.

Specifying policy elements: actions, effects, and principals

To use AWS Data Exchange, your user permissions must be defined in an IAM policy.

The following are the most basic policy elements:

- Resource In a policy, you use an Amazon Resource Name (ARN) to identify the resource to which the policy applies. All AWS Data Exchange API operations support resource level permissions (RLP), but AWS Marketplace actions don't support RLP. For more information, see AWS Data Exchange resources and operations.
- Action You use action keywords to identify resource operations that you want to allow or deny.
- Effect You specify the effect (allow or deny) when the user requests the specific action. If you don't explicitly grant access to (allow) a resource, access is implicitly denied. You can also explicitly deny access to a resource, which you might do to make sure that a user cannot access it, even if a different policy grants access.
- Principal In identity-based policies (IAM policies), the user that the policy is attached to is the implicit principal. For resource-based policies, you specify the user, account, service, or other entity that you want to receive permissions (applies to resource-based policies only). AWS Data Exchange doesn't support resource-based policies.

For more information about IAM policy syntax and descriptions, see <u>AWS Identity and Access</u> Management Policy Reference in the *IAM User Guide*.

Specifying conditions in a policy

When you grant permissions, you can use the IAM policy language to specify the conditions when a policy should take effect. With AWS Data Exchange, the CreateJob, StartJob, GetJob, and CancelJob API operations support conditional permissions. You can provide permissions at the JobType level.

AWS Data Exchange condition key reference

Condition key	Description	Туре
"dataexchange:JobType":"IMP ORT_ASSETS_FROM_S3"	Scopes permissions to jobs that import assets from Amazon S3.	String
<pre>"dataexchange:JobType":IMP0 RT_ASSETS_FROM_LAKE_FORMATI ON_TAG_POLICY" (Preview)</pre>	Scopes permissions to jobs that import assets from AWS Lake Formation (Preview)	String
"dataexchange:JobType":"IMP ORT_ASSET_FROM_SIGNED_URL"	Scopes permissions to jobs that import assets from a signed URL.	String
<pre>"dataexchange:JobType":"IMP ORT_ASSET_FROM_REDSHIFT_DAT A_SHARES"</pre>	Scopes permissions to jobs that import assets from Amazon Redshift.	String
<pre>"dataexchange:JobType":"IMP ORT_ASSET_FROM_API_GATEWAY_ API"</pre>	Scopes permissions to jobs that import assets from Amazon API Gateway.	String
"dataexchange:JobType":"EXP ORT_ASSETS_TO_S3"	Scopes permissions to jobs that export assets to Amazon S3.	String

Condition key	Description	Туре
"dataexchange:JobType":"EXP ORT_ASSETS_TO_SIGNED_URL"	Scopes permissions to jobs that export assets to a signed URL.	String
"dataexchange:JobType":EXPO RT_REVISIONS_TO_S3"	Scopes permissions to jobs that export revisions to Amazon S3.	String

For more information about specifying conditions in a policy language, see Condition in the IAM User Guide.

To express conditions, you use predefined condition keys. AWS Data Exchange has the JobType condition for API operations. However, there are AWS wide condition keys that you can use, as appropriate. For a complete list of AWS wide keys, see the IAM User Guide.

AWS Data Exchange API permissions: actions and resources reference

Use the following table as a reference when you are setting up Access control and writing a permissions policy that you can attach to an AWS Identity and Access Management (IAM) identity (identity-based policies). The table lists each AWS Data Exchange API operation, the actions for which you can grant permissions to perform the action, and the AWS resource for which you can grant the permissions. You specify the actions in the policy's Action field. You specify the resource value in the policy's Resource field.



Note

To specify an action, use the dataexchange: prefix followed by the API operation name (for example, dataexchange: CreateDataSet).

AWS Data Exchange API and required permissions for actions

AWS Data Exchange API operations	Required permissions (API actions)	Resources	Conditions
CreateDataSet	dataexchange:Creat eDataSet	N/A	aws:TagKeys
	Coucasee		aws:RequestTag
GetDataSet	dataexchange:GetDa taSet	Data set	aws:RequestTag
UpdateDataSet	dataexchange:Updat eDataSet	Data set	aws:RequestTag
PublishDataSet	<pre>dataexchange:Publi shDataSet</pre>	Data set	aws:RequestTag
DeleteDataSet	<pre>dataexchange:Delet eDataSet</pre>	Data set	aws:RequestTag
ListDataSets	<pre>dataexchange:ListD ataSets</pre>	N/A	N/A
CreateRevision	<pre>dataexchange:Creat eRevision</pre>	Data set	aws:TagKeys
	enevision		aws:RequestTag
GetRevision	<pre>dataexchange:GetRe vision</pre>	Revision	aws:RequestTag
DeleteRevision	<pre>dataexchange:Delet eRevision</pre>	Revision	aws:RequestTag
ListDataS etRevisions	<pre>dataexchange:ListD ataSetRevisions</pre>	Data set	aws:RequestTag
ListRevis ionAssets	<pre>dataexchange:ListR evisionAssets</pre>	Revision	aws:RequestTag

AWS Data Exchange API operations	Required permissions (API actions)	Resources	Conditions
CreateEve ntAction	<pre>dataexchange:Creat eEventAction</pre>	N/A	N/A
UpdateEve ntAction	<pre>dataexchange:Updat eEventAction</pre>	EventAction	N/A
GetEventAction	<pre>dataexchange:GetEv entAction</pre>	EventAction	N/A
ListEvent Actions	<pre>dataexchange:ListE ventActions</pre>	N/A	N/A
DeleteEve ntAction	<pre>dataexchange:Delet eEventAction</pre>	EventAction	N/A
CreateJob	dataexchange:Creat eJob	N/A	dataexcha nge:JobType
GetJob	dataexchange:GetJob	Job	dataexcha nge:JobType
StartJob**	dataexchange:StartJob	Job	dataexcha nge:JobType
CancelJob	dataexchange:Cance	Job	dataexcha nge:JobType
ListJobs	dataexchange:ListJobs	N/A	N/A
ListTagsF orResource	<pre>dataexchange:ListT agsForResource</pre>	Revision	aws:RequestTag
TagResource	dataexchange:TagRe source	Revision	aws:TagKeys aws:RequestTag

AWS Data Exchange API operations	Required permissions (API actions)	Resources	Conditions
UnTagResource	dataexchange:UnTag Resource	Revision	aws:TagKeys aws:RequestTag
UpdateRevision	dataexchange:Updat eRevision	Revision	aws:RequestTag
DeleteAsset	dataexchange:Delet eAsset	Asset	N/A
GetAsset	dataexchange:GetAsset	Asset	N/A
UpdateAsset	dataexchange:Updat eAsset	Asset	N/A
SendApiAsset	dataexchange:SendA piAsset	Asset	N/A

^{**} Additional IAM permissions might be needed depending on the type of the job you are starting. See the following table for the AWS Data Exchange job types and associated additional IAM permissions. For more information about jobs, see Jobs in AWS Data Exchange.

Note

Currently, the SendApiAsset operation is not supported for the following SDKs:

- · SDK for .NET
- AWS SDK for C++
- SDK for Java 2.x

AWS Data Exchange job type permissions for StartJob

Job type	Additional IAM permissions needed
IMPORT_ASSETS_FROM_S3	dataexchange:CreateAsset
<pre>IMPORT_ASSET_FROM_SIGNED_URL</pre>	dataexchange:CreateAsset
<pre>IMPORT_ASSETS_FROM_API_GATE WAY_API</pre>	dataexchange:CreateAsset
<pre>IMPORT_ASSETS_FROM_REDSHIFT _DATA_SHARES</pre>	<pre>dataexchange:CreateAsset ,redshift: AuthorizeDataShare</pre>
EXPORT_ASSETS_TO_S3	dataexchange:GetAsset
EXPORT_ASSETS_TO_SIGNED_URL	dataexchange:GetAsset
EXPORT_REVISIONS_TO_S3	<pre>dataexchange:GetRevision dataexcha nge:GetDataSet</pre>
	The IAM permission dataexcha nge:GetDataSet is only needed if you are using DataSet.Name as the dynamic reference for the EXPORT_RE VISIONS_TO_S3 job type.

You can scope data set actions to the revision or asset level through the use of wildcards, as in the following example.

arn:aws:dataexchange:us-east-1:123456789012:data-sets/99EXAMPLE23c7c272897cf1EXAMPLE7a/revisions/*/assets/*

Some AWS Data Exchange actions can only be performed on the AWS Data Exchange console. These actions are integrated with AWS Marketplace functionality. The actions require the AWS Marketplace permissions shown in the following table.

AWS Data Exchange console-only actions for subscribers

Console action	IAM permission
Subscribe to a product	aws-marketplace:Subscribe
	<pre>aws-marketplace:CreateAgree mentRequest</pre>
	aws-marketplace:AcceptAgree mentRequest
Send subscription verification request	aws-marketplace:Subscribe
	<pre>aws-marketplace:CreateAgree mentRequest</pre>
	<pre>aws-marketplace:AcceptAgree mentRequest</pre>
Enable subscription auto-renew	aws-marketplace:Subscribe
	<pre>aws-marketplace:CreateAgree mentRequest</pre>
	<pre>aws-marketplace:AcceptAgree mentRequest</pre>
View auto-renew status on a subscription	<pre>aws-marketplace:ListEntitle mentDetails</pre>
	<pre>aws-marketplace:ViewSubscri ptions</pre>
	<pre>aws-marketplace:GetAgreemen tTerms</pre>
Disable subscription auto-renew	aws-marketplace:Subscribe
	<pre>aws-marketplace:CreateAgree mentRequest</pre>

Console action	IAM permission
	<pre>aws-marketplace:AcceptAgree mentRequest</pre>
List active subscriptions	<pre>aws-marketplace:ViewSubscri ptions</pre>
	aws-marketplace:SearchAgreements
	<pre>aws-marketplace:GetAgreemen tTerms</pre>
View subscription	<pre>aws-marketplace:ViewSubscri ptions</pre>
	aws-marketplace:SearchAgreements
	aws-marketplace:GetAgreemen tTerms
	<pre>aws-marketplace:DescribeAgr eement</pre>
List subscription verification requests	<pre>aws-marketplace:ListAgreeme ntRequests</pre>
View subscription verification request	aws-marketplace:GetAgreemen tRequest
Cancel subscription verification request	<pre>aws-marketplace:CancelAgree mentRequest</pre>
View all offers targeted to the account	aws-marketplace:ListPrivate Listings
View details of a specific offer	<pre>aws-marketplace:GetPrivateL isting</pre>

AWS Data Exchange console-only actions for providers

Console action	IAM permission
Tag product	aws-marketplace:TagResource
	aws-marketplace:UntagResource
	aws-marketplace:ListTagsForResource
Tag offer	aws-marketplace:TagResource
	aws-marketplace:UntagResource
	aws-marketplace:ListTagsForResource
Publish product	aws-marketplace:StartChangeSet
	aws-marketplace:DescribeChangeSet
	dataexchange:PublishDataSet
Unpublish product	aws-marketplace:StartChangeSet
	aws-marketplace:DescribeChangeSet
Edit product	aws-marketplace:StartChangeSet
	aws-marketplace:DescribeChangeSet
Create custom offer	aws-marketplace:StartChangeSet
	aws-marketplace:DescribeChangeSet
Edit custom offer	aws-marketplace:StartChangeSet
	aws-marketplace:DescribeChangeSet
View product details	aws-marketplace:DescribeEntity
	aws-marketplace:ListEntities

Console action	IAM permission
View product's custom offer	aws-marketplace:DescribeEntity
View product dashboard	aws-marketplace:ListEntities
	aws-marketplace:DescribeEntity
List products to which a data set or revision has been published	aws-marketplace:ListEntities
	aws-marketplace:DescribeEntity
List subscription verificat ion requests	aws-marketplace:ListAgreementApprovalRequests
	aws-marketplace:GetAgreementApprovalRequest
Approve subscription verification requests	aws-marketplace:AcceptAgreementApprovalRequest
Decline subscription verification requests	aws-marketplace:RejectAgreementApprovalRequest
Delete information from subscription verification requests	aws-marketplace:UpdateAgreementApprovalRequest
View subscription details	aws-marketplace:SearchAgreements
	aws-marketplace:GetAgreementTerms

AWS managed policies for AWS Data Exchange

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you

reduce permissions further by defining <u>customer managed policies</u> that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see AWS managed policies in the IAM User Guide.

Topics

- AWS managed policy: AWSDataExchangeFullAccess
- AWS managed policy: AWSDataExchangeProviderFullAccess
- AWS managed policy: AWSDataExchangeReadOnly
- AWS managed policy: AWSDataExchangeServiceRolePolicyForLicenseManagement
- AWS managed policy: AWSDataExchangeServiceRolePolicyForOrganizationDiscovery
- AWS managed policy: AWSDataExchangeSubscriberFullAccess
- AWS managed policy: AWSDataExchangeDataGrantOwnerFullAccess
- AWS managed policy: AWSDataExchangeDataGrantReceiverFullAccess
- AWS Data Exchange updates to AWS managed policies

AWS managed policy: AWSDataExchangeFullAccess

You can attach the AWSDataExchangeFullAccess policy to your IAM identities.

This policy grants administrative permissions that allow full access to AWS Data Exchange and AWS Marketplace actions using the AWS Management Console and SDK. It also provides select access to Amazon S3 and AWS Key Management Service as needed to take full advantage of AWS Data Exchange.

To view permissions for this policy, see <u>AWSDataExchangeFullAccess</u> in the *AWS Managed Policy Reference*.

AWS managed policy: AWSDataExchangeProviderFullAccess

You can attach the AWSDataExchangeProviderFullAccess policy to your IAM identities.

This policy grants contributor permissions that provide data provider access to AWS Data Exchange and AWS Marketplace actions using the AWS Management Console and SDK. It also provides select access to Amazon S3 and AWS Key Management Service as needed to take full advantage of AWS Data Exchange.

To view permissions for this policy, see <u>AWSDataExchangeProviderFullAccess</u> in the *AWS Managed Policy Reference*.

AWS managed policy: AWSDataExchangeReadOnly

You can attach the AWSDataExchangeReadOnly policy to your IAM identities.

This policy grants read-only permissions that allow read-only access to AWS Data Exchange and AWS Marketplace actions using the AWS Management Console and SDK.

To view permissions for this policy, see <u>AWSDataExchangeReadOnly</u> in the *AWS Managed Policy Reference*.

AWS managed policy:

AWSDataExchangeServiceRolePolicyForLicenseManagement

You can't attach the AWSDataExchangeServiceRolePolicyForLicenseManagement to your IAM entities. This policy is attached to a service-linked role that allows AWS Data Exchange to perform actions on your behalf. It grants role permissions that allow AWS Data Exchange to retrieve information about your AWS organization and manage AWS Data Exchange data grants licenses. For more information, see Service-linked role for AWS Data Exchange license management later in this section.

To view permissions for this policy, see <u>AWSDataExchangeServiceRolePolicyForLicenseManagement</u> in the *AWS Managed Policy Reference*.

AWS managed policy:

AWSDataExchangeServiceRolePolicyForOrganizationDiscovery

You can't attach the AWSDataExchangeServiceRolePolicyForOrganizationDiscovery to your IAM entities. This policy is attached to a service-linked role that allows AWS Data Exchange to perform actions on your behalf. It grants role permissions that allow AWS Data Exchange to retrieve information about your AWS organization to determine eligibility for AWS Data Exchange data grants license distribution. For more information, see Service-linked roles for AWS Organization discovery in AWS Data Exchange.

To view permissions for this policy, see

<u>AWSDataExchangeServiceRolePolicyForOrganizationDiscovery</u> in the *AWS Managed Policy Reference*.

AWS managed policy: AWSDataExchangeSubscriberFullAccess

You can attach the AWSDataExchangeSubscriberFullAccess policy to your IAM identities.

This policy grants contributor permissions that allow data subscriber access to AWS Data Exchange and AWS Marketplace actions using the AWS Management Console and SDK. It also provides select access to Amazon S3 and AWS Key Management Service as needed to take full advantage of AWS Data Exchange.

To view permissions for this policy, see <u>AWSDataExchangeSubscriberFullAccess</u> in the *AWS Managed Policy Reference*.

AWS managed policy: AWSDataExchangeDataGrantOwnerFullAccess

You can attach the AWSDataExchangeDataGrantOwnerFullAccess policy to your IAM identities.

This policy gives a Data Grant owner access to AWS Data Exchange actions using the AWS Management Console and SDKs.

To view permissions for this policy, see <u>AWSDataExchangeDataGrantOwnerFullAccess</u> in the *AWS Managed Policy Reference*.

AWS managed policy: AWSDataExchangeDataGrantReceiverFullAccess

You can attach the AWSDataExchangeDataGrantReceiverFullAccess policy to your IAM identities.

This policy gives a Data Grant receiver access to AWS Data Exchange actions using the AWS Management Console and SDKs.

To view permissions for this policy, see <u>AWSDataExchangeDataGrantReceiverFullAccess</u> in the *AWS Managed Policy Reference*.

AWS Data Exchange updates to AWS managed policies

The following table provides details about updates to AWS managed policies for AWS Data Exchange since this service began tracking these changes. For automatic alerts about changes to

this page (and any other changes to this user guide), subscribe to the RSS feed on the <u>Document history for AWS Data Exchange</u> page.

Change	Description	Date
AWSDataExchangeDataGrantOwn erFullAccess – New policy	AWS Data Exchange added a new policy to grant Data Grant owners access to AWS Data Exchange actions.	October 24, 2024
AWSDataExchangeDataGrantRec eiverFullAccess – New policy	AWS Data Exchange added a new policy to grant Data Grant receivers access to AWS Data Exchange actions.	October 24, 2024
AWSDataExchangeReadOnly – Update to an existing policy	Added necessary permissions to the AWSDataExchangeReadOnly AWS managed policy for the new data grants feature.	October 24, 2024
<u>AWSDataExchangeServiceRoleP</u> <u>olicyForLicenseManagement</u> – New policy	Added a new policy to support service- linked roles to manage license grants in customer accounts.	October 17, 2024
AWSDataExchangeServiceRolePolicyForOrganizationDiscovery – New policy	Added a new policy to support service- linked roles to provide read access to account information in your AWS Organization.	October 17, 2024
AWSDataExchangeReadOnly	Added statement IDs to make the policy easier to read, expanded the wild carded permissions to the full list of read only ADX permissions, and added new actions: aws-marketplace:ListTagsForResource and aws-marketplace:ListPrivate Listings.	July 9, 2024

Change	Description	Date
AWSDataExchangeFullAccess	Removed action: aws-marke tplace:GetPrivateListing	May 22, 2024
AWSDataExchangeSubscriberFu llAccess	Added statement IDs to make the policy easier to read and added new action: aws-marketplace:ListPrivate Listings .	April 30, 2024
AWSDataExchangeFullAccess	Added statement IDs to make the policy easier to read and added new actions: aws-marketplace:TagResource , aws-marketplace:UntagResour ce , aws-marketplace:Li stTagsForResource , aws-marketplace:ListPrivateListings , aws-marketplace:GetPrivateListing , and aws-marketplace:De scribeAgreement .	April 30, 2024
AWSDataExchangeProviderFull Access	Added statement IDs to make the policy easier to read.	August 9, 2024
AWSDataExchangeProviderFull Access	Added dataexchange: SendD ataSetNotification , a new permission to send data set notificat ions.	March 5, 2024

Change	Description	Date
AWSDataExchangeSubscriberFu llAccess, AWSDataExchangePro viderFullAccess, and AWSDataEx changeFullAccess – Update to existing policies	Added granular actions across all managed policies. New actions added are aws-marketplace:Cr eateAgreementRequest , aws-marketplace:AcceptAgree mentRequest , aws-marke tplace:ListEntitlementDetails , aws-marketplace:ListEntitlementDetails , aws-marketplace:ListIngs , aws-marketplace:GetPrivateListing , license-manager:ListReceivedGrants aws-marketplace:TagResource , aws-marketplace:UntagResource , aws-marketplace:ListTagsForResource , aws-marketplace:ListTagsForResource , aws-marketplace:DescribeAgreement , aws-marketplace:DescribeAgreement , aws-marketplace:GetAgreementTermsaws-marketplace:GetLicense .	July 31, 2023
AWSDataExchangeProviderFull Access – Update to existing policy	Added dataexchange: Revok eRevision , a new permission to revoke a revision.	March 15, 2022
AWSDataExchangeProviderFull Access and AWSDataExchangeFul LAccess – Update to existing policies	Added apigateway: GET , a new permission to retrieve an API asset from Amazon API Gateway.	December 3, 2021
AWSDataExchangeProviderFull Access and AWSDataExchangeSub scriberFullAccess – Update to existing policies	Added dataexchange: SendA piAsset , a new permission to send a request to an API asset.	November 29, 2021

Change	Description	Date
AWSDataExchangeProviderFull Access and AWSDataExchangeFul LAccess – Update to existing policies	Added redshift: Authorize DataShare , redshift: DescribeDataSharesForProduc er , and redshift: DescribeD ataShares , new permissions to authorize access to and create Amazon Redshift data sets.	November 1, 2021
AWSDataExchangeSubscriberFu <u>llAccess</u> – Update to an existing policy	Added dataexchange:Creat eEventAction , dataexcha nge:UpdateEventAction , and dataexchange:DeleteEventAct ion , new permissions to control access to automatically export new revisions of data sets.	September 30, 2021
AWSDataExchangeProviderFull Access and AWSDataExchangeFul LAccess – Update to existing policies	Added dataexchange: Publi shDataSet , a new permission to control access to publishing new versions of data sets.	May 25, 2021
AWSDataExchangeReadOnly, AWSDataExchangeProviderFull Access, and AWSDataExchangeFul LAccess – Update to existing policies	Added aws-marketplace:Se archAgreements and aws-marke tplace:GetAgreementTerms to enable viewing subscriptions for products and offers.	May 12, 2021
AWS Data Exchange started tracking changes	AWS Data Exchange started tracking changes for its AWS managed policies.	April 20, 2021

Using service-linked roles for AWS Data Exchange

AWS Data Exchange uses AWS Identity and Access Management (IAM) <u>service-linked roles</u>. A service-linked role is a unique type of IAM role that is linked directly to AWS Data Exchange.

Using service-linked roles 350

Service-linked roles are predefined by AWS Data Exchange and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up AWS Data Exchange easier because you don't have to manually add the necessary permissions. AWS Data Exchange defines the permissions of its service-linked roles, and unless defined otherwise, only AWS Data Exchange can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting their related resources. This protects your AWS Data Exchange resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see <u>AWS services that work</u> with <u>IAM</u> and look for the services that have **Yes** in the **Service-linked roles** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Creating a service-linked role for AWS Data Exchange

You don't need to manually create a service-linked role. When you distribute a data grant using license manager, it creates the service-linked role for you.

To create a service-linked role

- 1. In the AWS Data Exchange console, sign in and choose **Data Grant settings**.
- 2. On the **Data Grant settings** page, choose **Configure integration**.
- 3. In the Create AWS Organizations integration section, select Configure integration.
- 4. On the **Create AWS Organizations integration** page, choose the appropriate trust level preference, and then choose **Create integration**.

You can also use the IAM console to create a service-linked role with a use case. In the AWS CLI or the AWS API, create a service-linked role with the <code>appropriate-service-name</code> . amazonaws . com service name. For more information, see Creating a service-linked role in the IAM User Guide. If you delete this service-linked role, you can use this same process to create the role again.

Editing a service-linked role for AWS Data Exchange

AWS Data Exchange does not allow you to edit the service-linked role. After you create a servicelinked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see Editing a service-linked role in the IAM User Guide.

Deleting a service-linked role for AWS Data Exchange

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.



Note

If the AWS Data Exchange service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes and try the operation again.

Before you can delete the service-linked role, you must:

- For the AWSServiceRoleForAWSDataExchangeLicenseManagement role, remove all AWS License Manager distributed grants for AWS Data Exchange data grants you received.
- For the AWSServiceRoleForAWSDataExchangeOrganizationDiscovery role, remove all AWS License Manager distributed grants for AWS Data Exchange data grants received by accounts in your AWS organization.

Manually deleting the service-linked role

Use the IAM console, the AWS CLI, or the AWS API to delete the service-linked role. For more information, see Deleting a service-linked role in the IAM User Guide.

Supported Regions for AWS Data Exchange service-linked roles

AWS Data Exchange supports using service-linked roles in all of the AWS Regions where the service is available. For more information, see AWS Regions and endpoints.

Service-linked role for AWS Data Exchange license management

AWS Data Exchange uses the service-linked role named

AWSServiceRoleForAWSDataExchangeLicenseManagement – this role allows AWS Data Exchange to retrieve information about your AWS organization and manage AWS Data Exchange data grants licenses.

The AWSServiceRoleForAWSDataExchangeLicenseManagement service-linked role trusts the following services to assume the role:

• license-management.dataexchange.amazonaws.com

The role permissions policy named

AWSDataExchangeServiceRolePolicyForLicenseManagement allows AWS Data Exchange to complete the following actions on the specified resources:

- Actions:
 - organizations:DescribeOrganization
 - license-manager:ListDistributedGrants
 - license-manager:GetGrant
 - license-manager:CreateGrantVersion
 - license-manager:DeleteGrant
- Resources:
 - All resources (*)

For more information about the

AWSDataExchangeServiceRolePolicyForLicenseManagement role, see <u>AWS managed</u> policy: AWSDataExchangeServiceRolePolicyForLicenseManagement.

For more information about using the

AWSServiceRoleForAWSDataExchangeLicenseManagement service-linked role, see <u>Using</u> service-linked roles for AWS Data Exchange.

You must configure permissions to allow your users, groups, or roles to create, edit, or delete a service-linked role. For more information, see Service-linked role permissions in the *IAM User Guide*.

Service-linked roles for AWS Organization discovery in AWS Data **Exchange**

AWS Data Exchange uses the service-linked role named

AWSServiceRoleForAWSDataExchangeOrganizationDiscovery - this role allows AWS Data Exchange to retrieve information about your AWS organization to determine eligibility for AWS Data Exchange data grants license distribution.



Note

This role is only needed in the AWS Organization's management account.

The AWSServiceRoleForAWSDataExchangeOrganizationDiscovery service-linked role trusts the following services to assume the role:

organization-discovery.dataexchange.amazonaws.com

The role permissions policy named

AWSDataExchangeServiceRolePolicyForOrganizationDiscovery allows AWS Data Exchange to complete the following actions on the specified resources:

- Actions:
 - organizations:DescribeOrganization
 - organizations:DescribeAccount
 - organizations:ListAccounts
- Resources:
 - All resources (*)

For more information about the

AWSDataExchangeServiceRolePolicyForOrganizationDiscovery role, see AWS managed policy: AWSDataExchangeServiceRolePolicyForOrganizationDiscovery.

For more information about using the

AWSServiceRoleForAWSDataExchangeOrganizationDiscovery service-linked role, see Using service-linked roles for AWS Data Exchange earlier in this section.

You must configure permissions to allow your users, groups, or roles to create, edit, or delete a service-linked role. For more information, see Service-linked role permissions in the *IAM User Guide*.

Compliance validation for AWS Data Exchange

To learn whether an AWS service is within the scope of specific compliance programs, see <u>AWS</u> services in Scope by Compliance Program and choose the compliance program that you are interested in. For general information, see AWS Compliance Programs.

You can download third-party audit reports using AWS Artifact. For more information, see Downloading Reports in AWS Artifact.

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- <u>Security Compliance & Governance</u> These solution implementation guides discuss architectural considerations and provide steps for deploying security and compliance features.
- HIPAA Eligible Services Reference Lists HIPAA eligible services. Not all AWS services are HIPAA eligible.
- <u>AWS Compliance Resources</u> This collection of workbooks and guides might apply to your industry and location.
- AWS Customer Compliance Guides Understand the shared responsibility model through the
 lens of compliance. The guides summarize the best practices for securing AWS services and map
 the guidance to security controls across multiple frameworks (including National Institute of
 Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and
 International Organization for Standardization (ISO)).
- <u>Evaluating Resources with Rules</u> in the AWS Config Developer Guide The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- <u>AWS Security Hub</u> This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS resources and to check your

Compliance validation 355

compliance against security industry standards and best practices. For a list of supported services and controls, see Security Hub controls reference.

- Amazon GuardDuty This AWS service detects potential threats to your AWS accounts, workloads, containers, and data by monitoring your environment for suspicious and malicious activities. GuardDuty can help you address various compliance requirements, like PCI DSS, by meeting intrusion detection requirements mandated by certain compliance frameworks.
- <u>AWS Audit Manager</u> This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

PCI DSS compliance

AWS Data Exchange supports the processing, storage, and transmission of credit card data by a merchant or service provider, and has been validated as being compliant with Payment Card Industry (PCI) Data Security Standard (DSS). For more information about PCI DSS, including how to request a copy of the AWS PCI Compliance Package, see PCI DSS Level 1.

Resilience in AWS Data Exchange

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

AWS Data Exchange has a single, globally available product catalog offered by providers. Subscribers can see the same catalog, regardless of which Region they are using. The resources underlying the product (data sets, revisions, assets) are regional resources that you manage programmatically or through the AWS Data Exchange console in supported Regions. AWS Data Exchange replicates your data across multiple Availability Zones within the Regions where the service operates. For information about supported Regions, see Global Infrastructure Region Table.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

PCI DSS compliance 356

Infrastructure security in AWS Data Exchange

As a managed service, AWS Data Exchange is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see AWS Cloud Security. To design your AWS environment using the best practices for infrastructure security, see Infrastructure Protection in Security Pillar AWS Well-Architected Framework.

You use AWS published API calls to access AWS Data Exchange through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the AWS Security Token Service (AWS STS) to generate temporary security credentials to sign requests.

AWS Data Exchange and interface VPC endpoints (AWS PrivateLink)

You can establish a private connection between your virtual private cloud (VPC) and AWS Data Exchange by creating an interface VPC endpoint. Interface endpoints are powered by AWS PrivateLink, a technology that enables you to privately access AWS Data Exchange API operations without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with AWS Data Exchange API operations. Traffic between your VPC and AWS Data Exchange does not leave the Amazon network.

Each interface endpoint is represented by one or more Elastic Network Interfaces in your subnets.



Note

Every AWS Data Exchange action, except for SendAPIAsset, is supported for VPC.

357 Infrastructure security

For more information, see <u>Interface VPC endpoints (AWS PrivateLink)</u> in the *Amazon VPC User Guide*.

Considerations for AWS Data Exchange VPC endpoints

Before you set up an interface VPC endpoint for AWS Data Exchange, ensure that you review Interface endpoint properties and limitations in the *Amazon VPC User Guide*.

AWS Data Exchange supports making calls to all of its API operations from your VPC.

Creating an interface VPC endpoint for AWS Data Exchange

You can create a VPC endpoint for the AWS Data Exchange service using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see <u>Creating an interface endpoint</u> in the *Amazon VPC User Guide*.

Create a VPC endpoint for AWS Data Exchange using the following service name:

• com.amazonaws.region.dataexchange

If you enable private DNS for the endpoint, you can make API requests to AWS Data Exchange using its default DNS name for the AWS Region, for example, com.amazonaws.us-east-1.dataexchange.

For more information, see <u>Accessing a service through an interface endpoint</u> in the *Amazon VPC User Guide*.

Creating a VPC endpoint policy for AWS Data Exchange

You can attach an endpoint policy to your VPC endpoint that controls access to AWS Data Exchange. The policy specifies the following information:

- · The principal that can perform actions
- The actions that can be performed
- The resources on which actions can be performed

For more information, see <u>Controlling access to services with VPC endpoints</u> in the *Amazon VPC User Guide*.

Example: VPC endpoint policy for AWS Data Exchange actions

The following is an example of an endpoint policy for AWS Data Exchange. When attached to an endpoint, this policy grants access to the listed AWS Data Exchange actions for all principals on all resources.

This example VPC endpoint policy allows full access only to the user bts in AWS account 123456789012 from vpc-12345678. The user readUser is allowed to read the resources, but all other IAM principals are denied access to the endpoint.

JSON

```
}
    "Id": "example-policy",
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "Allow administrative actions from vpc-12345678",
            "Effect": "Allow",
            "Principal": {
                "AWS": [
                    "arn:aws:iam::123456789012:user/bts"
                1
            },
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "aws:sourceVpc": "vpc-12345678"
                }
            }
        },
            "Sid": "Allow ReadOnly actions",
            "Effect": "Allow",
            "Principal": {
                "AWS": [
                    "arn:aws:iam::123456789012:user/readUser"
                1
            },
            "Action": [
                "dataexchange:list*",
```

```
"dataexchange:get*"
],
"Resource": "*"
}
]
```

Using AWS Data Exchange with the AWS Marketplace **Catalog API**

This chapter contains supplemental information for using AWS Data Exchange and the AWS Marketplace Catalog API. The AWS Marketplace Catalog API service provides an API interface for you as a provider to programmatically access the AWS Marketplace self-service publishing capabilities.

The API supports a wide range of operations for you to view and manage your products. You can extend your internal build or deployment pipeline to AWS Marketplace through API integration to automate your product update process. You can also create your own internal user interface on top of the API to manage your products on the AWS Marketplace.

You can use the AWS Marketplace Catalog API to update your AWS Data Exchange products. To view your products, you can use the ListEntities and DescribeEntity API operations. To update your AWS Data Exchange product, you need to create a new change set, which is the Catalog API resource that represents an asynchronous operation used to manage products. For more information, see the AWS Marketplace Catalog API Reference.

Keep the following in mind when working with the Catalog API:

- Each AWS Data Exchange product is represented in the Catalog API as an Entity.
- AWS Data Exchange products have DataProduct as the EntityType.
- Each product can have only one concurrently running change set at a time. This means that you can't create a second change set until the first one has finished running.

Topics

Add data sets to AWS Data Exchange

Add data sets to AWS Data Exchange



Note

Data sets added via the Catalog API change set of type AddDataSets default to the publishing method of the product.

AddDataSets 361 To add data sets to your AWS Data Exchange product, start a change set of type AddDataSets. To do so, you can use the StartChangeSet API operation and specify the change type, the product identifier, the product type, and the details including the data set Amazon Resource Name (ARN).

Tutorial: Adding new data sets to a published data product

This tutorial walks you through detailed steps to add new AWS Data Exchange data sets to a published product. The tutorial has the following high-level steps.

Topics

- Set up IAM permissions
- Access the AWS Marketplace Catalog API
- Get your product ID from the AWS Data Exchange console
- Start a change request
- · Check the status of your change set

Set up IAM permissions

Before you begin, you need AWS Identity and Access Management (IAM) permissions for using the AWS Marketplace Catalog API. These permissions are in addition to the permissions you need for using AWS Data Exchange.

- 1. Navigate your browser to the IAM console and sign in using an AWS account that can manage IAM permissions.
- 2. From the left navigation pane, choose **Policies**.
- 3. Choose Create policy.
- 4. Choose the **JSON** tab, and provide the following permissions. This provides full access to the AWS Marketplace Catalog API. You can restrict access as appropriate for your use case.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
       "Effect": "Allow",
       "Action": [
```

```
"aws-marketplace:CancelChangeSet",
    "aws-marketplace:ListChangeSets",
    "aws-marketplace:DescribeEntity",
    "aws-marketplace:StartChangeSet",
    "aws-marketplace:ListEntities",
    "aws-marketplace:DescribeChangeSet",
    "dataexchange:PublishDataSet"
    ],
    "Resource": "*"
    }
]
```

- Choose Next: Review.
- 6. Provide a name for the policy (for example, **CatalogAPIFullAccess**), and then choose **Create Policy**.
- 7. Using the IAM console, choose the users, groups, or roles that you want to attach the policy to.

Access the AWS Marketplace Catalog API

To access the AWS Marketplace Catalog API, use the following HTTP client endpoint.

```
catalog.marketplace.us-east-1.amazonaws.com
```

Get your product ID from the AWS Data Exchange console

Before you can use the AWS Marketplace Catalog API to publish new data sets, get your product ID from the AWS Data Exchange console. Navigate to the **Product Dashboard**, and then copy the product ID you would like to publish data sets for. You may also use the <u>AWS Marketplace Catalog API</u> to find your product ID, using the ListEntities action with the **DataProduct@1.0** entity type.

Start a change request

To start a change request to add a data set in your test product

- 1. Copy the entity ID that you get by following the instructions in <u>Get your product ID from the</u> AWS Data Exchange console.
- 2. Make a StartChangeSet request with an AddDataSets change type.



Note

For information about working with change sets in the AWS Marketplace Catalog API, see Working with change sets. For more information about working with the identifier for entities, see Identifier.

Example request

```
https://catalog.marketplace.us-east-1.amazonaws.com/StartChangeSet
```

Example request body

```
{
    "Catalog": "AWSMarketplace",
    "ChangeSetName": "Adding Data Set to my test Data Product",
    "ChangeSet": [
        {
            "ChangeType": "AddDataSets",
            "Entity": {
                "Identifier": "entity-id@1",
                "Type": "DataProduct@1.0"
            },
            "Details": "{ \"DataSets\": [ { \"Arn\": \"data-set-arn\" } ] }"
        }
    ]
}
```

Example response

```
{
  "ChangeSetId": "cs-bnEXAMPLE4mkz9oh",
  "ChangeSetArn": "arn:aws:aws-marketplace:us-east-1:account-id:AWSMarketplace/
ChangeSet/cs-bnEXAMPLE4mkz9oh"
}
```

Check the status of your change set

After you use the StartChangeSet API operation to start the change request, you can use the DescribeChangeSet operation to check its status. Provide the change set ID returned in the StartChangeSet API response.

Example request

```
https://catalog.marketplace.us-east-1.amazonaws.com/DescribeChangeSet?
catalog=AWSMarketplace&changeSetId=cs-bnEXAMPLE4mkz9oh
```

Example request body

```
{
"changeSetId":"cs-bnEXAMPLE4mkz9oh"
}
```

Example response

```
{
    "ChangeSetId": "cs-bnEXAMPLE4mkz9oh",
    "ChangeSetArn": "arn:aws:aws-marketplace:us-east-1:account-id:AWSMarketplace/
ChangeSet/cs-bnEXAMPLE4mkz9oh",
    "ChangeSetName": "Adding Data Set to my test Data Product",
    "StartTime": "2018-09-20T19:45:03.115+0000",
    "EndTime": "2018-09-20T19:48:12.517+0000",
    "Status": "SUCCEEDED",
    "FailureDescription": null,
    "ChangeSet": [
        {
            "ChangeType": "AddDataSets",
            "Entity": {
                "Type": "DataProduct@1.0",
                "Identifier": "entity-id@1"
            },
            "ErrorList": []
        }
    ]
}
```

AddDataSets exceptions

The following exceptions can occur when you use the AWS Marketplace Catalog API with AWS Data Exchange:

DATA_SET_NOT_FOUND

AddDataSets exceptions 365

This happens when the requested data set was not found. To resolve this issue, ensure that there's not a typo in the data set ARN and that your AWS account owns the data set, and try again.

INVALID_INPUT

The request couldn't be processed due to input that isn't valid. To resolve this issue, ensure that there's not a typo in the request and that the product does not exceed the maximum number of allowed data sets.

DATA_SET_ALREADY_PUBLISHED

This happens when the data set has already been previously added to the product.

DATA_SET_DUPLICATE_PROVIDED

This happens when the same data set is provided more than once in the request.

AddDataSets exceptions 366

Document history for AWS Data Exchange

The following table describes the documentation for this release of the AWS Data Exchange User Guide. For notification about updates to this documentation, you can subscribe to the RSS feed.

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

Change	Description	Date
New managed policy and update to existing policy	AWS Data Exchange added the new AWSDataEx changeDataGrantOwn erFullAccess and AWSDataExchangeDat aGrantReceiverFull Access AWS managed policies. Edited the AWSDataExchangeRea dOnly AWS managed policies. For more informati on, see AWS managed policies.	October 24, 2024
Added new AWS managed policies and service-linked roles	Added the AWSDataEx changeServiceRoleP olicyForLicenseMan agement and AWSDataEx changeServiceRoleP olicyForOrganizati onDiscovery AWS managed policies. For more information, see AWS managed policies for AWS Data Exchange. Also added the AWSServiceRoleForA WSDataExchangeLice nseManagement and	October 17, 2024

AWSServiceRoleForA
WSDataExchangeOrga
nizationDiscovery
service-linked roles. For
more information, see <u>Using</u>
service-linked roles for AWS
Data Exchange.

Update to existing policies

Expanded the wild carded permissions to the full list of read only ADX permissions and added aws-marke tplace:ListTagsForResource and aws-marke tplace:ListPrivate Listings to AWSDataEx changeReadOnly.

July 9, 2024

Update to existing policies

Removed action aws-marke tplace:GetPrivateL isting from AWSDataEx changeFullAccess and AWSDataExchangePro viderFullAccess .

May 22, 2024

Update to existing policies

Statement IDs and the following new actions have been added to these policies: AWSDataExchangeRea dOnly managed policy: aws-marketplace:Li stTagsForResource and aws-marketplaceLis tPrivateListings ; AWSDataExchangeSub scriberFullAccess managed policy: aws-marke tplace:ListPrivate Listings ; AWSDataEx changeFullAccess managed policy: aws-marke tplace:TagResource ,aws-marketplace:Un tagResource ,aws-marke tplace:ListTagsFor Resource , aws-marke tplace:ListPrivate Listings , aws-marke tplace:GetPrivateL isting , and aws-marke tplace:DescribeAgr eement . For more informati on, see AWS managed

policies.

April 30, 2024

Update to existing policy

The following new permission has been added to the AWSDataExchangeProviderFullAccess managed policy: dataexchange: SendDataSetNotification . For more information, see AWS managed policies.

March 5, 2024

Ability to create data grants is now available

Data owners can now share data using AWS Data Exchange without registering as a AWS Marketplace seller. For more information see <u>Creating data grants on AWS</u> Data Exchange.

December 14, 2023

<u>Provider-generated notificat</u> ions are now available

Providers can send notificat ions corresponding to data updates, data delays, schema changes, and deprecations. Subscribers receive these notifications as events in Amazon EventBridge they can use to build automated workflows or deliver human-readable notifications to emails and chat programs using AWS User Notifications. For more information see Provider-generated notifications.

October 31, 2023

New subscriber event notifications available for auto-export jobs

Subscribers can receive notifications for two new events: Auto-export Job Completed and Auto-export Job Failed. For more informati on, see Events for an auto-export job completed and Events for an auto-export job failed.

October 4, 2023

Ability for subscribers to download files directly from an Amazon S3 bucket

Subscribers can now download files directly from a provider's Amazon S3 bucket from the AWS Data Exchange Console. For more informati on, see Publishing a new product containing Amazon S3 data access.

September 5, 2023

Changes in AWS Data

ExchangeAWS CloudTrail and migration to AWS Marketpla ce Agreement Service (MPAS)

AWS Data Exchange is migrating to AWS Marketpla ce Agreement Service causing changes in AWS Data Exchange CloudTrail events. For more information, see Upcoming changes in AWS Data Exchange CloudTrail logging.

June 1, 2023

Ability to use AWS User Notifications

AWS User Notifications provides users with a single place in the AWS Managemen t Console to set up and view all relevant AWS notifications across accounts, AWS Regions, and services. Users can configure delivery channels for notifications, such as email, chat, and mobile push notifications. For more information, see AWS User Notifications for AWS Data Exchange.

May 18, 2023

Ability to publish and subscribe to products containing Amazon S3 data access

Subscribing to and publishin g data products containing Amazon S3 data access is now generally available. For more information, see Publishin g a new product containin g Amazon S3 data access and Subscribing to a product containing Amazon S3 data access.

March 14, 2023

Ability for subscribers to receive notifications for Amazon S3 data access data set resources

Subscribers can now receive notifications when a provider performs actions on Amazon S3 resources. For more information, see <u>Amazon</u>
EventBridge events.

February 10, 2023

Updated tutorials to include

AWS Data Exchange for

Amazon S3 (Test Product)

(Preview)

The following tutorial shows how to browse the AWS
Data Exchange catalog to find and subscribe to AWS
Data Exchange for Amazon
S3 (Test Product) (Preview)
: Tutorial: Subscribe to AWS
Data Exchange for Amazon S3
(Test Product) (Preview).

February 6, 2023

Ability to publish and subscribe to products containing Amazon S3 data access (Preview) Ability for data subscribers to access AWS Glue tables through AWS Lake Formation (Preview)

Providers can now create products that contain Amazon S3 data access. For more information, see Publishin g a new product containin g Amazon S3 data access (Preview). Subscribers can now find, subscribe to, and use data from the data provider's Amazon S3 data sets. For more information, see Subscribing to a product containing Amazon S3 data access (Preview). Subscribe rs can find and subscribe to live, ready-to-use, third-par ty AWS Glue tables through AWS Lake Formation that they can query and analyze without extracting, transform ing, and loading the underlyin g files. For more information see, Subscribe to and access a product containing AWS Lake Formation data sets (Preview) Subscribing to a product containing Amazon S3 data access (Preview). Subscribe rs can find and subscribe to live, ready-to-use, third-par ty AWS Glue tables through AWS Lake Formation that they can query and analyze without extracting, transform ing, and loading the underlyin g files. For more information,

November 30, 2022

see <u>Subscribe to and access a</u> <u>product containing AWS Lake</u> Formation data sets (Preview).

Israel is now an eligible jurisdiction

Residents in Israel are now eligible to become sellers on AWS Data Exchange. For more information, see <u>Getting</u> started as a provider.

August 29, 2022

Extended Provider Program

The Extended Provider
Program (EPP) is now
generally available. For more
information, see Extended
Provider Program (EPP).

August 9, 2022

Export file size limit increase

The file size limit of a cross-Region revision export to Amazon S3 has increased from 100 GB to 1,000 GB. The number of assets that can be exported to Amazon S3 in a single cross-Region ExportRevision job has increased from 2,000 to 10,000. For more informati on, see Export and import job quidelines.

August 4, 2022

Similar products

Subscribers can now see a list of similar products at the bottom of a product detail page, which they can use to continue their browse journey without needing to return to the search results page. For more information, see Subscribe to and access a product.

July 28, 2022

Post-subscription enhanceme nt

After subscribing to a product that contains different types of data sets, subscribers can now view see separate cards with icons that display the different types of data sets. Subscribers can learn more information about the data sets and go directly to their entitled data from the postsubscription page. In addition, subscribers to products that contain S3 data sets can set up manual or automatic exports directly on the postsubscription page. For more information, see Subscribe to and access a product.

July 25, 2022

Ability to export data set

ID and data set name when
exporting revisions

When exporting revisions to an Amazon S3 bucket, both providers and subscribers can now export the data set ID and the name of the data set being exported. For more information, see Exporting revisions.

July 14, 2022

Integration with Open Data on AWS

Anyone, with or without an AWS account, can now search and find data products from the Open Data on AWS project. For more information, see Using Open Data on AWS data sets.

June 21, 2022

Exporting data sets video

Documentation-only update to add a video: One-Time Exports of Third-Party Data Sets from AWS Data Exchange. For more informati on, see Exporting assets and Exporting revisions.

May 27, 2022

Ability to specify metered costs for API products and subscribe to API products with Pay As You Go pricing

Providers can now specify metered costs for their API products. For more informati on, see <u>Publishing a product containing APIs</u>. Subscribers can now find and subscribe to third-party APIs with Pay As You Go pricing. This feature reduces upfront subscriber costs relative to monthly data file subscriptions. For more information, see <u>Subscribe</u> to and access a product containing APIs.

May 19, 2022

Ability to view and edit subscription verification requests as a provider

Documentation-only update to clarify how to view and edit subscription verificat ion requests as a provider. For more information, see Subscription verification for providers.

May 6, 2022

<u>Updated tutorials to include</u> data dictionaries and samples The following tutorials now include data dictionaries and samples: Tutorial: Subscribe to AWS Data Exchange
Heartbeat on AWS Data
Exchange, Tutorial: Subscribe to AWS Data Exchange for APIs (Test Product) on AWS
Data Exchange, and Tutorial:
Subscribe to Worldwide Event Attendance (Test Product) on AWS Data Exchange.

April 13, 2022

Ability to provide and subscribe to products containing data dictionaries and samples

Providers can now create and update data products that contain data dictionaries and samples. For more informati on, see Data dictionaries and Samples. Subscribers can evaluate the products containing data dictionaries and samples before subscribi ng. For more information, see Data dictionaries and samples. Subscribers can learn more about how to manage their subscriptions in the new topic Managing subscriptions.

March 31, 2022

Publishing products video

Documentation-only update to add a video: Publish products on AWS Data Exchange. For more informati on, see <u>Publishing a new</u> product.

March 18, 2022

Update to existing policy

The following new permissio n has been added to the AWSDataExchangePro viderFullAccess managed policy: dataexchange: RevokeRevision. For more information, see AWS managed policies.

March 15, 2022

Ability to revoke revisions

Providers can revoke subscribers' access to a revision and delete the assets of the revision. For more information, see Revoking revisions. Subscribers will get an Amazon EventBridge event notifying them that their access to the revision was revoked and the reason for the revocation. For more information, see Amazon EventBridge events.

March 15, 2022

Added tutorial for subscribi ng to products containing API data sets The following new tutorial has been added: <u>Tutorial</u>: <u>Subscribe to AWS Data</u>
<u>Exchange for APIs (Test Product) on AWS Data</u>
Exchange.

January 14, 2022

Ability to publish and subscribe to products containing Amazon Redshift data sets

Providers can now create and license products that contain Amazon Redshift data sets. For more informati on, see Publishing a product containing Amazon Redshift data sets. Subscribers can now find, subscribe to, and use data from the data provider's Amazon Redshift data sets. For more information, see Subscribi ng to a product containin g Amazon Redshift data sets. Subscribers can also receive notifications when a provider performs actions on an Amazon Redshift resource. For more informati on, see Amazon EventBridge events. The following tutorial has been added: Tutorial: Subscribe to Worldwide Event Attendance (Test Product) on AWS Data Exchange.

January 4, 2022

Update to existing policies

The following new permission to retrieve an API from Amazon API Gateway has been added to the AWS managed policies:
AWSDataExchangeProviderFullAccess and AWSDataExchangeFul lAccess: apigateway: GET . For more information, see AWS managed policies.

December 3, 2021

Update to existing policies

The following new permission n to send a request to an API asset has been added to the AWS managed policies:

AWSDataExchangeProviderFullAccess and AWSDataExchangeSubscriberFullAccess:

dataexchange:SendApiAsset.For more information, see AWSmanaged policies.

November 29, 2021

Ability to provide and subscribe to third-party APIs

Providers can now create API data products using AWS Data Exchange and use AWS Data Exchange to manage subscriber authentication, pricing, billing, and pay-asyou-go access to their REST APIs. For more information, see Publishing a new API data product. Subscribers can now find and subscribe to API-based data from thirdparty REST APIs in the AWS Cloud. They can use AWS native authentication and governance and use AWSgenerated SDKs to make API calls. For more information, see Subscribing to an API data product.

November 29, 2021

<u>Update to existing policies</u> (Public Preview)

The following new permissio ns to authorize access to and create Amazon Redshift data sets have been added to the AWS managed policies (Public Preview): AWSDataEx changeProviderFull Access and AWSDataEx changeFullAccess : redshift:Authorize DataShare , redshift: DescribeDataShares ForProducer , and redshift:DescribeD ataShares . For more information, see AWS managed policies.

November 1, 2021

Ability to publish and subscribe to Amazon Redshift data products (Public Preview)

Providers can now create and license Amazon Redshift data products using AWS Data Exchange. For more information, see Publishing a new Amazon Redshift data product (Preview). Subscribe rs can now find, subscribe to, and use data from the data provider's Amazon Redshift data sets. For more information, see Subscribing to an Amazon Redshift data product (Preview)

October 19, 2021

Update to an existing policy

The following new permissions to control access to automatically export new revisions of data sets have been added to the AWS managed policy AWSDataEx changeSubscriberFullAccess: dataexchange:Creat eEventAction,

dataexchange:Updat
eEventAction , and
dataexchange:Delet
eEventAction .For

more information, see AWS

managed policies.

September 30, 2021

Ability to automatically export revisions

Subscribers can now automatically export revisions . For more information, see Automatically exporting revisions to an S3 bucket as a subscriber (console).

September 30, 2021

<u>Updated procedure for how</u> to use jobs

The Jobs in AWS Data

Exchange section has been updated to clarify how to import and export assets and export revisions through jobs.

September 7, 2021

Added procedure for how to unsubscribe from a data product

The <u>Subscribing to data</u>
<u>products on AWS Data</u>
<u>Exchange</u> section has been reorganized and a new subsection has been added to clarify how to unsubscribe from a product. For more information, see <u>Unsubscribe from a product on AWS Data Exchange</u>.

August 11, 2021

Support for sharing licenses through AWS License Manager You can share licenses to products that you purchase with other accounts in your AWS organization. For more information, see Sharing license subscriptions in an organization.

August 4, 2021

Ability to automatically publish revisions

Providers can now automatic ally publish revisions to data sets. For more information, see Publishing a new data set revision using automatic revision publishing. For information about how to migrate an existing data set to automatic revision publishing, see Migrating an existing product to automatic revision publishing.

July 22, 2021

Updated p	roduct description
templates	

The following product description templates have been updated: Media and entertainment long description template and Retail and location long description template.

July 19, 2021

More eligible jurisdictions

The following are now eligible to become sellers on AWS Data Exchange: Hong Kong SAR and Qatar. For more information, see Eligible jurisdictions for AWS Data Exchange products.

June 24, 2021

Ability to view changes to managed policies

You can now see the changes made to AWS managed policies for AWS Data Exchange. They are tracked in

the AWS managed policies for

AWS Data Exchange topic.

May 25, 2021

Added payment scheduler

You can now use a payment schedule to invoice subscribe rs for private or renewed private offers. For more information, see Create private offers.

May 24, 2021

Added ability to add data sets programmatically

You can now add data sets using the AWS Marketplace Catalog API service. For more information, see <u>Using AWS</u>
<u>Data Exchange with the AWS</u>
<u>Marketplace Catalog API.</u>

August 23, 2020

Support for	preferred
currency	

You can pay for AWS Data Exchange subscriptions using your preferred currency. For more information see Pricing. July 27, 2020

More eligible jurisdictions

The following are now eligible to become sellers on AWS Data Exchange: Bahrain, Norway, Switzerland, and the United Arab Emirates (UAE). For more information, see Eligible jurisdictions for AWS Data Exchange products.

June 16, 2020

Added encryption support for exporting data sets

AWS Data Exchange now supports configurable encryption parameters when exporting data sets to Amazon S3. For more information, see Exporting assets to an Amazon S3
Bucket.

April 27, 2020

AWS Data Exchange is now generally available

AWS Data Exchange is a service that makes it easy for AWS customers to create, update, maintain, and securely exchange file-based data sets in the AWS Cloud.

November 13, 2019