
Data Transfer Hub

Implementation Guide

Data Transfer Hub: Implementation Guide

Copyright © 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

Welcome	1
Cost	2
Cost of an Amazon S3 transfer task	2
Cost of an Amazon ECR transfer task	3
Architecture overview	5
Amazon S3 plugin	6
Amazon ECR plugin	7
Solution components	8
Web console	8
Amazon S3 transfer engine	8
Amazon ECR transfer engine	8
Security	9
IAM roles	9
Amazon CloudFront	9
Design considerations	10
Regional deployments	10
AWS CloudFormation template	11
Automated deployment	12
Deployment overview	12
Step 1. (Option 1) Launch the stack in AWS Regions	12
Prerequisite	12
Deploy the AWS CloudFormation template for Option 1 – AWS Regions	12
Step 1. (Option 2) Launch the stack in AWS China Regions	13
Prerequisites	13
Prerequisite 1: Create an OIDC user pool	14
Prerequisite 2: Configure domain name service resolution	15
Deploy the AWS CloudFormation template for Option 2 – AWS China Regions	15
Step 2. Launch the web console	16
Log in to the web console	16
Log in using Amazon Cognito user pool for AWS Regions	17
(Optional 2) OpenID authentication for AWS China Regions	17
Step 3. Create a transfer task	17
Create an Amazon S3 transfer task	18
Create an Amazon ECR transfer task	19
Additional resources	21
Uninstall the solution	22
Using the AWS Management Console	22
Using AWS Command Line Interface	22
Deleting the Amazon S3 buckets	22
Source code	23
Document revisions	24
Contributors	25
Notices	26
AWS glossary	27

Framework for secure, scalable, and trackable data transfer for Amazon Simple Storage Service (Amazon S3) objects and Amazon Elastic Container Registry (Amazon ECR) images

Publication date: January 2022

This solution provides secure, scalable, and trackable data transfer for Amazon Simple Storage Service (Amazon S3) objects and Amazon Elastic Container Registry (Amazon ECR) images. This data transfer helps customers expand their businesses globally by easily moving data in and out of AWS China Regions.

The solution's web console provides an interface for managing the following tasks:

- Transferring Amazon S3 objects between AWS China Regions and AWS Regions
- Transferring data from other cloud providers' object storage services (including Alibaba Cloud OSS, Tencent COS, and Qiniu Kodo) to Amazon S3
- Transferring objects from Amazon S3 compatible object storage service to Amazon S3
- Transferring Amazon ECR images between AWS China Regions and AWS Regions
- Transferring container images from public container registries (for example, Docker Hub, Google gcr.io, Red Hat Quay.io) to Amazon ECR

Note

If you need to transfer Amazon S3 objects between AWS Regions, we recommend that you use [Cross-Region Replication](#); if you want to transfer Amazon S3 objects within the same AWS Region, we recommend using [Same-Region Replication](#).

This implementation guide describes architectural considerations and configuration steps for deploying Data Transfer Hub in the Amazon Web Services (AWS) Cloud. It includes links to an AWS CloudFormation template that launches and configures the AWS services required to deploy this solution using AWS best practices for security and availability.

The guide is intended for IT architects, developers, DevOps, data analysts, and marketing technology professionals who have practical experience architecting in the AWS Cloud.

Cost

You are responsible for the cost of the AWS services used while running this solution, which can vary based on whether you are transferring Amazon S3 objects or Amazon ECR images.

The solution automatically deploys an additional Amazon CloudFront distribution and an Amazon S3 bucket for storing the static website assets in your account. You are responsible for the incurred variable charges from these services. For full details, refer to the pricing webpage for each AWS service you will be using in this solution.

The following examples demonstrate how to estimate the cost. Two example estimates are for transferring S3 objects, and one is for transferring ECR images.

Cost of an Amazon S3 transfer task

For an Amazon S3 transfer task, the cost can vary based on the total number of files and the average file size.

Example 1: As of January 2022, transfer 1 TB of S3 files from AWS Oregon Region (us-west-2) to AWS Beijing Region (cn-north-1), the average file size is 50MB.

- Total files: ~2,048
- Average speed per EC2 instance: ~1GB/min
- Total EC2 instance hours: ~17 hours

AWS service	Dimensions	Total Cost
Amazon EC2	\$0.0084 per hour (t4g.micro)	\$0.14
Amazon S3	~ 12 GET requests + 10 PUT request per file GET: \$0.0004 per 1000 request PUT: \$0.005 per 1000 request	\$0.11
Amazon DynamoDB	~2 write requests per file \$1.25 per million write	\$0.01
Amazon SQS	~2 request per file \$0.40 per million requests	\$0.01
Data Transfer Out	0.09 per GB	\$92.16
Others (For example, CloudWatch, Secrets Manager, etc.)		~ \$1
	TOTAL:	~ \$93.43

Example 2: As of January 2022, transfer 1 TB of S3 files from AWS Oregon region (us-west-2) to Mainland China Beijing Region (cn-north-1), the average file size is 10KB.

- Total files: ~107,400,000
- Average speed per EC2 instance: ~6MB/min (~10 files per sec)
- Total EC2 instance hours: ~3000 hours

AWS Service	Dimensions	Total Cost
Amazon EC2	\$0.0084 per hour (t4g.micro)	\$25.20
Amazon S3	~ 2 GET requests + 1 PUT request per file GET: \$0.0004 per 1000 request PUT: \$0.005 per 1000 request	\$622.34
Amazon DynamoDB	~2 write requests per file \$1.25 per million write	\$268.25
Amazon SQS	~2 request per file \$0.40 per million requests	\$85.92
Data Transfer Out	0.09 per GB	\$92.16
Others (For example, CloudWatch, Secrets Manager, etc.)		\$20
	TOTAL:	~ \$1113.87

Cost of an Amazon ECR transfer task

For an Amazon ECR transfer task, the cost can vary based on network speed and total size of ECR images.

Example 3: As of January 2022, transfer 27 Amazon ECR images (~3 GB in total size) from AWS Ireland Region (eu-west-1) to AWS Beijing Region (cn-north-1). The total runtime is about 6 minutes.

AWS Service	Dimensions	Total Cost
AWS Lambda	\$0.0000004 per 100ms	\$0.000072 (35221.95 ms)
AWS Step Functions	\$0.000025 per State Transition (~ 60 state transitions per run in this case)	\$0.0015
Fargate	\$0.04048 per vCPU per hour \$0.004445 per GB per hour (0.5 vCPU 1GB Memory)	\$0.015 (~ 2200s)

Data Transfer Hub Implementation Guide
Cost of an Amazon ECR transfer task

AWS Service	Dimensions	Total Cost
Data Transfer Out	0.09 per GB	\$0.27
Others (For example, CloudWatch, Secrets Manager, etc.)	Almost 0 costs	\$0
	TOTAL:	~ \$0.287

Architecture overview

Deploying the Data Transfer Hub solution with the default parameters builds the following environment in the AWS Cloud.

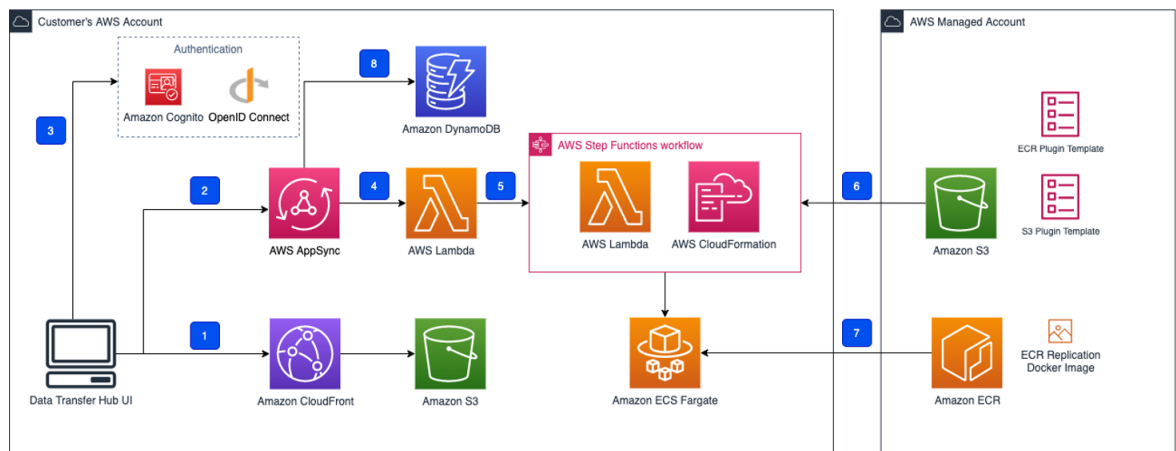


Figure 1: Data Transfer Hub architecture on AWS

The solution automatically deploys and configures a serverless architecture with the following services:

1. The solution's static web assets (front end user interface) are stored in [Amazon S3](#) and made available through [Amazon CloudFront](#).
2. The backend APIs are provided via [AWS AppSync](#) GraphQL.
3. Users are authenticated by either [Amazon Cognito User Pool](#) (in AWS Regions) or by an OpenID connect provider (in AWS China Regions) such as Authing, Auth0, etc.
4. AWS AppSync runs [AWS Lambda](#) to call backend APIs.
5. Lambda starts an [AWS Step Functions](#) workflow that uses [AWS CloudFormation](#) to start or stop/delete the ECR or S3 plugin template.
6. The plugin templates are hosted in a centralized Amazon S3 bucket managed by AWS
7. The solution also provisions an [Amazon ECS](#) cluster that runs the container images used by the plugin template, and the container images are hosted in [Amazon ECR](#).
8. The data transfer task information is stored in [Amazon DynamoDB](#).

Important

If you deploy this solution in AWS (Beijing) Region operated by Beijina Sinnet Technology Co., Ltd. (Sinnet), or the AWS (Ningxia) Region operated by Ningxia Western Cloud Data Technology Co., Ltd., you are required to provide a domain with [ICP Recordal](#) before you can access the web console.

The web console is a centralized place to create and manage all data transfer jobs. Each data type (for example, Amazon S3 or Amazon ECR) is a plugin for Data Transfer Hub, and is packaged as an AWS CloudFormation template hosted in an S3 bucket that AWS owns. When you create a transfer task, an AWS Lambda function initiates the Amazon CloudFormation template, and state of each task is stored and displayed in the DynamoDB tables.

As of December 2021, the solution supports two data transfer plugins: an Amazon S3 plugin and an Amazon ECR plugin.

Amazon S3 plugin

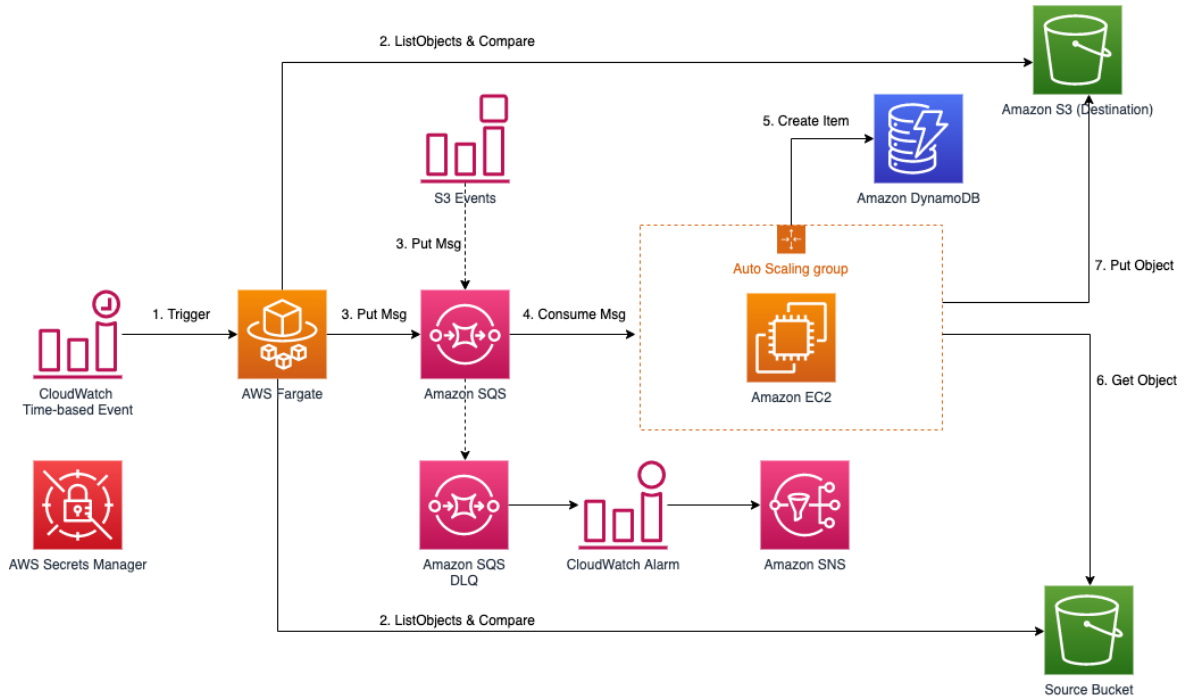


Figure 2: Data Transfer Hub Amazon S3 plugin architecture

The Amazon S3 plugin runs the following workflows:

1. A time-based Event Bridge rule triggers the AWS Fargate task to run on an hourly basis.
2. The Fargate task lists all the objects in the source and destination buckets and determines which objects should be transferred.
3. Fargate sends a message for each object that will be transferred to Amazon Simple Queue Service (Amazon SQS). Amazon S3 event messages can also be supported for more real-time data transfer; whenever there is object uploaded to source bucket, the event message is sent to the same SQS queue.
4. A JobWorker running in EC2 consumes the messages in SQS and transfers the object from the source bucket to the destination bucket. You can use an Auto Scaling Group to control the number of EC2 instances to transfer the data based on business need.
5. A record with transfer status for each object is stored in Amazon DynamoDB.
6. The Amazon EC2 instance will get (download) the object from the source bucket based on the SQS message.
7. The EC2 instance will put (upload) the object to the destination bucket based on the SQS message.

Note

If an object (or part of an object) failed to transfer, the JobWorker releases the message in the queue, and the object is transferred again after the message is visible in the queue (default visibility timeout is set to 15 minutes). If the transfer fails again, the message is sent to the dead letter queue and a notification alarm is sent.

Amazon ECR plugin

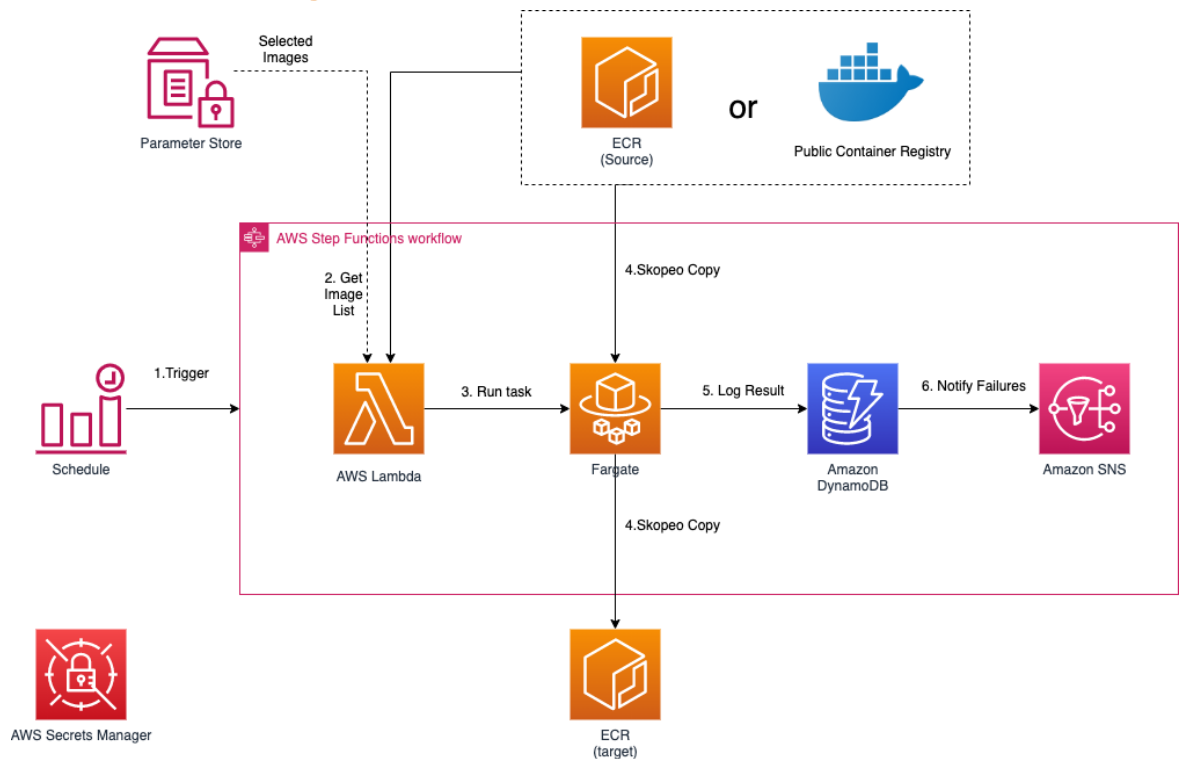


Figure 3: Data Transfer Hub Amazon ECR plugin architecture

The Amazon S3 plugin runs the following workflows:

1. An EventBridge rule runs an AWS Step Functions workflow on a regular basis (by default, it runs daily).
2. Step Functions invokes AWS Lambda to retrieve the list of images from the source.
3. Lambda will either list all the repository content in the source Amazon ECR, or get the stored image list from System Manager Parameter Store.
4. The transfer task will run within Fargate in a maximum concurrency of 10. If a transfer task failed for some reason, it will automatically retry three times.
5. Each task uses [skopeo copy](#) copy to copy the images into the target ECR.
6. After the copy completes, the status (either success or fail) is logged into DynamoDB for tracking purpose.

Solution components

This solution has three components: 1) a web console, 2) the Amazon S3 transfer engine, and 3) the Amazon ECR transfer engine.

Web console

This solution provides a simple web console which allows you to create and manage transfer tasks for Amazon S3 and Amazon ECR.

Amazon S3 transfer engine

Amazon S3 transfer engine runs the Amazon S3 plugin and is used for transferring objects from their sources into S3 buckets. The S3 plugin supports the following features:

- Transfer Amazon S3 objects between AWS China Regions and AWS Regions
- Transfer objects from Alibaba Cloud OSS / Tencent COS / Qiniu Kodo to Amazon S3
- Transfer objects from S3 Compatible Storage service to Amazon S3
- Support near real time transfer via S3 Event
- Support Transfer with object metadata
- Support incremental data transfer
- Auto retry and error handling

Amazon ECR transfer engine

Amazon ECR engine runs the Amazon ECR plugin and is used for transferring container images from other container registries. The ECR plugin supports the following features:

- Transfer Amazon ECR images between AWS China Regions and AWS Regions
- Transfer from public container registry (such as Docker Hub, GCR.io, Quay.io) to Amazon ECR
- Transfer selected images to Amazon ECR
- Transfer all images and tags from Amazon ECR

The ECR plugin leverages [skopeo](#) for the underlying engine. The AWS Lambda function lists images in their sources and uses Fargate to run the transfer jobs.

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This [shared model](#) reduces your operational burden because AWS operates, manages, and controls the components including the host operating system, the virtualization layer, and the physical security of the facilities in which the services operate. For more information about AWS security, see [AWS Cloud Security](#).

IAM roles

AWS Identity and Access Management (IAM) roles allow customers to assign granular access policies and permissions to services and users on the AWS Cloud. This solution creates IAM roles that grant the solution's AWS Lambda functions, Amazon API Gateway and Amazon Cognito access to create regional resources.

Amazon CloudFront

This solution deploys a web console [hosted](#) in an Amazon Simple Storage Service (Amazon S3) bucket. To help reduce latency and improve security, this solution includes an Amazon CloudFront distribution with an origin access identity, which is a CloudFront user that provides public access to the solution's website bucket contents. For more information, refer to [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) in the *Amazon CloudFront Developer Guide*.

Design considerations

Regional deployments

This solution uses services which may not be currently available in all AWS Regions. Launch this solution in an AWS Region where required services are available. For the most current availability by Region, refer to the [AWS Regional Services List](#).

Supported regions for deployment in AWS Regions

Region ID	Region Name
us-east-1	US East (N. Virginia)
us-east-2	US East (Ohio)
us-west-1	US West (N. California)
us-west-2	US West (Oregon)
ap-south-1	Asia Pacific (Mumbai)
ap-northeast-2	Asia Pacific (Seoul)
ap-southeast-1	Asia Pacific (Singapore)
ap-southeast-2	Asia Pacific (Sydney)
ap-northeast-1	Asia Pacific (Tokyo)
ca-central-1	Canada (Central)
eu-central-1	Europe (Frankfurt)
eu-west-1	Europe (Ireland)
eu-west-2	Europe (London)
eu-west-3	Europe (Paris)
eu-north-1	Europe (Stockholm)
sa-east-1	South America (São Paulo)

Supported regions for deployment in AWS China Regions

Region ID	Region Name
cn-north-1	China (Beijing) Region Operated by Sinnet
cn-northwest-1	China (Ningxia) Regions operated by NWCD

AWS CloudFormation template

To automate deployment, this solution uses the following AWS CloudFormation templates, which you can download before deployment:

[View template](#)

DataTransferHub-cognito.template: Use this template to launch the solution and all associated components in **AWS Regions** where Amazon Cognito is available. The default configuration deploys Amazon S3, Amazon CloudFront, AWS AppSync, Amazon DynamoDB, AWS Lambda, Amazon ECS, and Amazon Cognito, but you can customize the template to meet your specific needs.

[View template](#)

DataTransferHub-openid.template: Use this template to launch the solution and all associated components in **AWS China Regions** where Amazon Cognito is **not** available. The default configuration deploys The default configuration deploys Amazon S3, Amazon CloudFront, AWS AppSync, Amazon DynamoDB, AWS Lambda, and Amazon ECS, but you can customize the template to meet your specific needs.

Automated deployment

Before you launch the solution, [review the cost \(p. 2\)](#), architecture, network security, and other considerations discussed in this guide. Follow the step-by-step instructions in this section to configure and deploy the solution into your account.

Time to deploy: Approximately 15 minutes

Deployment overview

Use the following steps to deploy this solution on AWS. For detailed instructions, follow the links for each step.

Step 1. Launch the stack

- (Option 1) [Deploy the AWS CloudFormation template in AWS Regions \(p. 12\)](#)
- (Option 2) [Deploy the AWS CloudFormation template in AWS China Regions \(p. 13\)](#)

Step 2. [Launch the web console \(p. 16\)](#)

Step 3. [Create a transfer task \(p. 17\)](#)

Step 1. (Option 1) Launch the stack in AWS Regions

Important

The following deployment instructions apply to AWS Regions only. For deployment in AWS China Regions refer to [Option 2 \(p. 13\)](#).

Prerequisite

If you are deploying in a new account for the first time, the service-linked role for AppSync may not exist. You must run the following command in [CloudShell](#) to create the role.

```
aws iam create-service-linked-role
--aws-service-name appsync.amazonaws.com
```

Deploy the AWS CloudFormation template for Option 1 – AWS Regions

This automated AWS CloudFormation template deploys Data Transfer Hub in the AWS Cloud. You must complete the prerequisite before launching the stack.

Note

You are responsible for the cost of the AWS services used while running this solution. For more details, visit the [Cost \(p. 2\)](#) section in this guide, and refer to the pricing webpage for each AWS service used in this solution.

1. Sign in to the AWS Management Console and use the button below to launch the `DataTransferHub-cognito.template` AWS CloudFormation template. Alternatively, you can [download the template](#) as a starting point for your own implementation.



2. The template launches in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar.
3. On the **Create stack** page, verify that the correct template URL is in the **Amazon S3 URL** text box and choose **Next**.
4. On the **Specify stack details** page, assign a name to your solution stack. For information about naming character limitations, refer to [IAM and STS Limits](#) in the *AWS Identity and Access Management User Guide*.
5. Under **Parameters**, review the parameters for this solution template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
AdminEmail	<Requires input>	The email of the Admin user.

6. Choose **Next**.
7. On the **Configure Stack Options** page, keep the default values and choose **Next**.
8. On the **Review** page, review and confirm the settings. Check the box acknowledging that the template will create AWS Identity and Access Management (IAM) resources.
9. Choose **Create stack** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation console in the **Status** column. You should receive a `CREATE_COMPLETE` status in approximately 15 minutes.

Step 1. (Option 2) Launch the stack in AWS China Regions

Important

The following deployment instructions apply to AWS China Regions only. For deployment in AWS Regions refer to [Option 1 \(p. 12\)](#).

Prerequisites

1. [Create an OIDC User Pool \(p. 14\)](#).
2. [Configure domain name service \(DNS\) resolution \(p. 15\)](#).

Prerequisite 1: Create an OIDC user pool

In AWS Regions where Amazon Cognito is not yet available, you can use OIDC to provide authentication. The following procedure uses AWS Partner [Authing](#) as an example, but you can also choose any available provider.

1. Sign up for an Authing developer account. For more information, see [How to register an account](#).
2. Sign in to Authing.
3. Select **Create new user pool**, enter a name, and choose **Confirm**.
4. After the user pool is created, you can then create an application for OIDC authentication.
 - a. From **Application**, select **Configuration**, then choose **Basic Settings**.
 - b. Ensure that the protocol type is OIDC.
5. Update the authorization configuration.
 - a. From **Application**, select **Configuration**, then choose **Authorization configuration**.
 - b. Update the authorization flow to **implicit** and the return type to **id_token**.
 - c. For the id_token signature algorithm, select **RS256**.
 - d. Choose **Save**.
6. Update the callback URL.
 - a. From **Application**, select **Configuration**, then choose **Auth Config**.
 - b. Modify the login callback URL to `https://<your-custom-domain>/authentication/callback`.
 - c. Choose **Save**.

Note

Verify that the domain name has completed ICP registration in China.

7. Update login control.
 - a. From **Application**, select **Login control**, then choose **Registration and Login**.
 - b. Uncheck the checkboxes for Email and Phone to deactivate all ways to register.
 - c. Choose **Save**.
8. Create an admin user.
 - a. From **Users & Roles**, select **Users**, then choose **Create user**.
 - b. Enter the email for the user.
 - c. Choose **OK**.
 - d. Check the email for a temporary password.
 - e. Reset the user password.

Note

Because this Solution does not support application roles, all the users will receive admin rights.

Prerequisite 2: Configure domain name service resolution

Configure domain name service (DNS) resolution to point the ICP licensed domain to the CloudFront default domain name. Optionally, you can use your own DNS resolver.

The following is an example for configuration an Amazon Route 53.

1. Create a hosted zone in Amazon Route 53. For more information refer to the [Amazon Route 53 Developer Guide](#).
2. Create a CNAME record for the console URL.
 - a. From the hosted zone, choose **Create Record**.
 - b. In the **Record name** input box, enter the host name.
 - c. From **Record type** select **CNAME**.
 - d. In the value field, Enter the CloudFormation output PortalUrl.
 - e. Select **Create records**.
3. Add alternative domain names to the CloudFront distribution.
 - a. Configure the corresponding domain name in CloudFront to open the CloudFront console by finding the distribution ID for PortalURL in the list and selecting **ID** (or check the check box, and then select **Distribution Settings**).

Edit the distribution and add the Route 53 record to the alternative domain Names (CNAMEs).

Deploy the AWS CloudFormation template for Option 2 – AWS China Regions

This automated AWS CloudFormation template deploys Data Transfer in the AWS Cloud. You must Create an ODIC User Pool and Configure DNS resolution before launching the stack.

Note

You are responsible for the cost of the AWS services used while running this solution. For more details, visit the [Cost \(p. 2\)](#) section in this guide, and refer to the pricing webpage for each AWS service used in this solution.

1. Sign in to the AWS Management Console and use the button below to launch the `DataTransferHub-openid.template` AWS CloudFormation template. Alternatively, you can [download the template](#) as a starting point for your own implementation.



2. The template launches in your console's default Region. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar.
3. On the **Create stack** page, verify that the correct template URL is in the **Amazon S3 URL** text box and choose **Next**.

4. On the **Specify stack details** page, assign a name to your solution stack. For information about naming character limitations, refer to [IAM and STS Limits](#) in the *AWS Identity and Access Management User Guide*.
5. Under **Parameters**, review the parameters for this solution template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
OidcProvider	<Requires input>	Refers to the Issuer shown in the OIDC application configuration.
OidcClientId	<Requires input>	Refers to the App ID shown in the OIDC application configuration.
OidcCustomerDomain	<Requires input>	Refers to the customer domain that has completed ICP registration in China, not the subdomain provided by Authing. It must start with <code>https://</code> .

6. Choose **Next**.
7. On the **Configure Stack Options** page, keep the default values and choose **Next**.
8. On the **Review** page, review and confirm the settings. Check the box acknowledging that the template will create AWS Identity and Access Management (IAM) resources.
9. Choose **Create Stack** to deploy the stack.

You can view the status of your stack in the AWS CloudFormation console in the **Status** column. You should receive a `CREATE_COMPLETE` status in approximately 15 minutes.

Step 2. Launch the web console

After the stack is successfully created, navigate to the CloudFormation **Outputs** tab and select the **PortalUrl** value to access the Data Transfer Hub web console.

Log in to the web console

There are two options for you to log in and configure the Data Transfer Hub, depending on where the solution was deployed.

1. Log in with Amazon Cognito User Pool (for AWS Regions)
2. Log in with OpenID using Authing.cn (for AWS China Regions)

After successful deployment, an email containing the temporary login password is sent to the provided email address.

Log in using Amazon Cognito user pool for AWS Regions

1. Using a web browser, enter the **PortalURL** from the CloudFormation **Output** tab, then navigate to the Amazon Cognito console.
2. Sign in with the **AdminEmail** and the temporary password.
 - a. Set a new account password.
 - b. (Optional) Verify your email address for account recovery.
3. After the verification is complete, the system opens the Data Transfer Hub web console.

(Option 2) OpenID authentication for AWS China Regions

1. Using a web browser, enter the Data Transfer Hub domain name.
 - a. If you are logging in for the first time, the system will open the Authing.cn login interface.
2. Enter the username and password you registered when you deployed the solution, then choose **Login**. The system opens the Data Transfer Hub web console.
3. Change your password and then sign in again.

Step 3. Create a transfer task

Use the web console to create a transfer task for Amazon S3 or Amazon ECR.

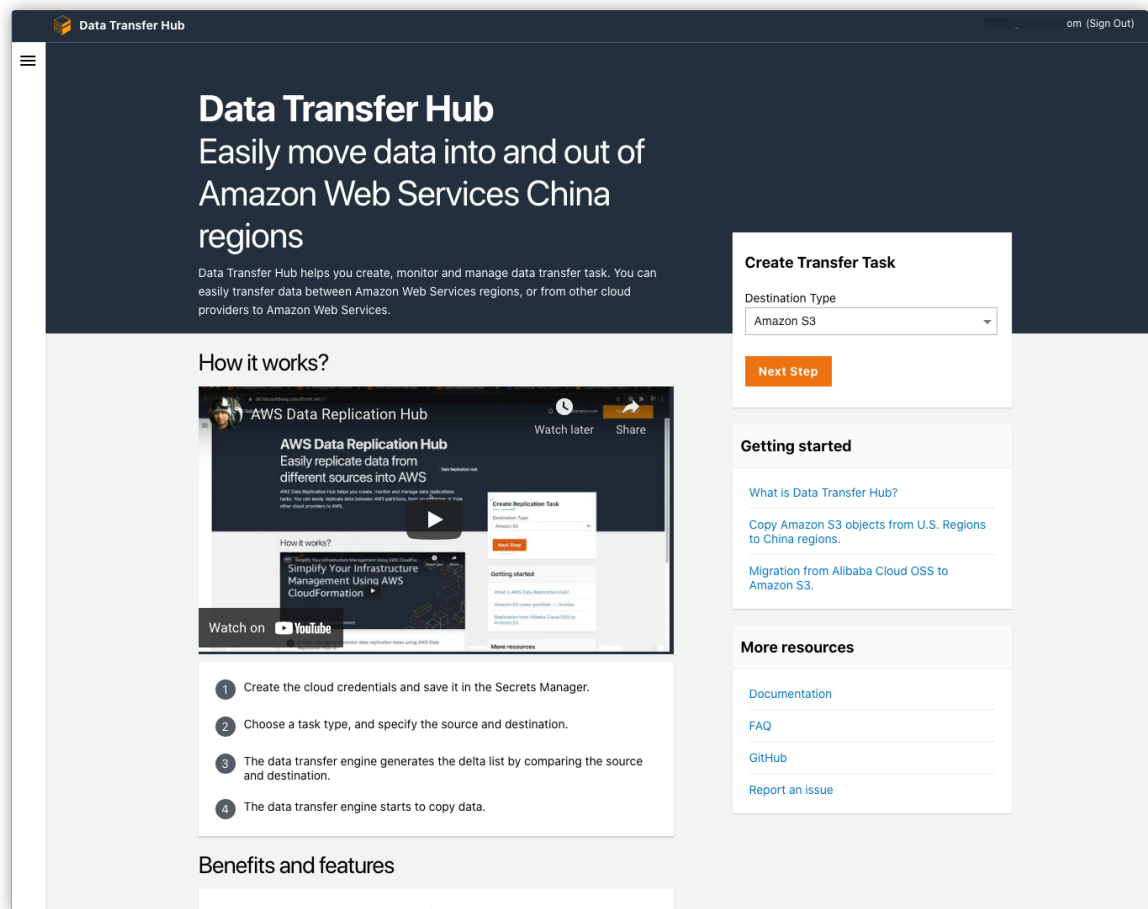


Figure 4: Data Transfer Hub web console

Create an Amazon S3 transfer task

1. From the **Create Transfer Task** page, select **Amazon S3**, and then choose **Next Step**.
2. From the **Engine options** page, under engine, select **Amazon S3**, and then choose **Next Step**.
3. Specify the transfer task details.
 - a. Under **Source Type**, select the data source, for example, **Amazon S3**.
4. Enter the **Source settings** for your request. Provide the **Input bucket name** and **object prefix** (optional).

If the data source bucket is in the account where Data Transfer Hub was deployed, select **Yes**, and choose to allow the S3 Event to trigger data transfer.

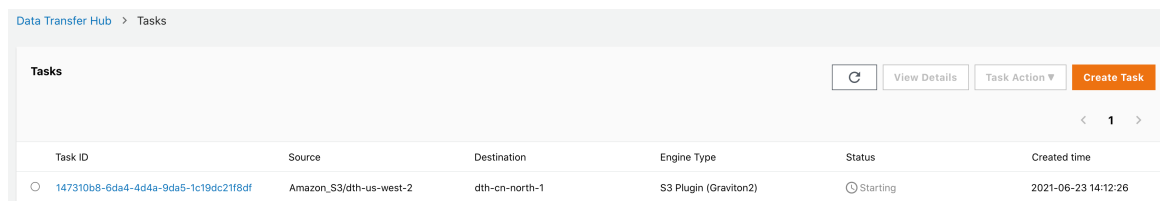
If the source bucket is *not* in the same account where Data Transfer Hub was deployed, select **No**, then specify the credentials for the source bucket.

To create credential info, choose **Secrets Manager** to navigate to the current Region's AWS Secrets Manager console.

- a. From the left menu, select **Secrets**, then choose **Store a new secret** and select the **other type of secrets** key type.

- b. Fill in the `access_key_id` and `secret_access_key` information in the **Plaintext** input box according to the displayed format. For more information, refer to [IAM features](#) in the *IAM User Guide*. Choose **Next**.
 - c. (Optional) Enter the key name and description. Choose **Next**.
 - d. In the configuration of automatic rotation, select **Disable automatic rotation**. Choose **Next**.
 - e. Keep the default value and choose **Save** to complete the creation of the key.
 - f. Navigate back to the Data Transfer Hub task creation interface and refresh the interface. Your new secret is displayed in the drop-down list.
 - g. Select the certificate (Secret).
5. Provide destination settings for the S3 buckets.
- If the source S3 bucket is in the same account where Data Transfer Hub was deployed, then in **destination settings**, you must create or provide credential information for the S3 destination bucket. Otherwise, no credential information is needed. Use the following steps to update the destination settings.
- a. From **Engine settings**, verify the values and modify them if necessary. We recommend to have the minimum capacity set to at least 1 if for incremental data transfer.
 - b. From **Advanced Options**, keep the default values.
 - c. From the **More** tab, provide any additional information. An email address must be provided in **Alarm Email**.
4. Choose **Next** and review your task parameter details.
5. Choose **Create Task**.

After the task is created successfully, it will appear on the **Tasks** page.



The screenshot shows the 'Data Transfer Hub > Tasks' page. At the top right, there are buttons for 'View Details', 'Task Action', and 'Create Task'. Below the buttons is a table with the following columns: Task ID, Source, Destination, Engine Type, Status, and Created time. One task is listed with the following details:

Task ID	Source	Destination	Engine Type	Status	Created time
147310b8-6da4-4d4a-9da5-1c19dc21f8df	Amazon_S3/dth-us-west-2	dth-cn-north-1	S3 Plugin (Graviton2)	Starting	2021-06-23 14:12:26

Figure 5: Transfer task details and status

Select the **Task ID** to go to the task **Details** page, and then choose **CloudWatch Dashboard** to monitor the task status.

Create an Amazon ECR transfer task

1. From the **Create Transfer Task** page, select **Amazon ECR**, and then choose **Next Step**.
2. Specify the transfer task details.
 - a. From **Source Type**, select the container registry type.
 - b. From **Source settings**, fill in the source Region name and the account ID. To create credential info, choose **Secrets Manager** to navigate to the current Region's AWS Secrets Manager console.
 - i. From the left menu, select **Secrets**, then choose **Store a new secret** and select the **other type of secrets** key type.
 - ii. Fill in the `access_key_id` and `secret_access_key` information in the **Plaintext** input box according to the displayed format. For more information, refer to [IAM features](#) in the *IAM User Guide*. Choose **Next**.

- iii. (Optional) Enter the key name and description. Choose **Next**.
- iv. In the configuration of automatic rotation, select **Disable automatic rotation**. Choose **Next**.
- v. Keep the default value and choose **Save** to complete the creation of the key.
- vi. Navigate back to the Data Transfer Hub task creation interface and refresh the interface. Your new secret is displayed in the drop-down list.
- vii. Select the certificate (Secret).
- viii. Provide destination settings.

If the source is in the same account with Data Transfer Hub deployment, you need to create/ provide credential info for the destination. Otherwise, no credential information is needed.

- 3. In the **More** section, provide additional information. From the **More** tab, provide any additional information. An email address must be provided in **Alarm Email**.
- 3. Choose **Next** and review your task parameter details.
- 4. Choose **Create Task**.

After the task is created successfully, you can view the task status in **Tasks** page.

Additional resources

- | | |
|---|---|
| <ul style="list-style-type: none">• AWS CloudFormation• Amazon S3• AWS Lambda• AWS Step Functions• Amazon CloudFront• Amazon ECR | <ul style="list-style-type: none">• Amazon DynamoDB• AWS AppSync• Amazon Cognito• AWS IAM• Amazon EC2• Amazon Route 53 |
|---|---|

Uninstall the solution

You can uninstall the Data Transfer Hub solution from the AWS Management Console or by using the AWS Command Line Interface. You must manually stop any active transfer tasks before uninstalling.

Using the AWS Management Console

1. Sign in to the AWS CloudFormation console.
2. On the **Stacks** page, select this solution's installation stack.
3. Choose **Delete**.

Using AWS Command Line Interface

Determine whether the AWS Command Line Interface (AWS CLI) is available in your environment. For installation instructions, refer to [What Is the AWS Command Line Interface](#) in the *AWS CLI User Guide*. After confirming that the AWS CLI is available, run the following command.

```
$ aws cloudformation delete-stack --stack-name  
<installation-stack-name>
```

Deleting the Amazon S3 buckets

This solution is configured to retain the solution-created Amazon S3 bucket (for deploying in an opt-in Region) if you decide to delete the AWS CloudFormation stack to prevent accidental data loss. After uninstalling the solution, you can manually delete this S3 bucket if you do not need to retain the data. Follow these steps to delete the Amazon S3 bucket.

1. Sign in to the [Amazon S3 console](#)
1. Choose **Buckets** from the left navigation pane.
2. Locate the *<stack-name>* S3 buckets.
3. Select the S3 bucket and choose **Delete**.

To delete the S3 bucket using AWS CLI, run the following command:

```
$ aws s3 rb s3://<bucket-name> --force
```

Source code

Visit our [GitHub repository](#) to download the source files for this solution and to share your customizations with others. The Data Transfer Hub templates are generated using the [AWS Cloud Development Kit \(AWS CDK\) \(AWS CDK\)](#). Refer to the [README.md file](#) for additional information.

Document revisions

Date	Change
January 2022	Initial release

Contributors

- Aiden Dai
- Kervin Hu
- Haiyun Chen
- Joe Shi
- Ashwini Rudra
- Jyoti Tyagi

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS current product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. AWS responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Data Transfer Hub is licensed under the terms of the of the Apache License Version 2.0 available at [Classless Inter-Domain Routing \(CIDR\)](#).

AWS glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS General Reference*.