Alibaba Cloud

Data Online Migration Migrate data from Amazon Simple Storage Service (Amazon S3) to OSS

Document Version: 20220628

C-J Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
- 2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
- 3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
- 4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
- 5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud and/or its affiliates Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
- 6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
A Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger: Resetting will result in the loss of user configuration data.
O Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
C) Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.
⑦ Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click OK.
Courier font	Courier font is used for commands	Run the cd /d C:/window command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	bae log listinstanceid Instance_ID
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

Table of Contents

1.Background information	05
2.Prepare for a migration job	06
3.Migrate data	08
4.Manage migration jobs	17
5.Regions and endpoints	19

1.Background information

This tutorial describes how to migrate data from Amazon Simple Storage Service (Amazon S3) to Alibaba Cloud Object Storage Service (OSS).

Alibaba Cloud Data Online Migration is a data channel between different storage services. You can use Data Online Migration to migrate data from third-party data stores to OSS or between OSS buckets.

When you use Data Online Migration, you need only to log on to the console, specify a source data address and a destination OSS data address, and then create a migration job. After the migration job starts, you can manage the job in the console. For example, you can view the progress and traffic information of the job. You can also generate a migration report to check the list of migrated files and the list of files that fail to be migrated.

♥ Notice

- When you read data from the source data address during a migration job, a fee is incurred for the outbound Internet traffic. You are charged by the storage service provider of the source data address.
- By default, Data Online Migration does not support cross-border data migration. For example, you cannot migrate data from a bucket that is deployed in the China (Beijing) region to a bucket that is deployed in the US (Silicon Valley) region. Before you create a migration job to migrate data across borders, you must submit a ticket to apply for the permissions that are required for creating cross-border migration jobs. You must commit that your business is legitimate, data transit conforms to local rules and regulations, and your data does not contain illegal information.

This tutorial includes the following topics:

- Preparations
- Implement migration
- Subsequent operations

2.Prepare for a migration job

This topic describes how to prepare for a migration job.

AWS S3

• Estimate the amount of data

Estimate the amount and the number of files that you want to migrate. You can log on to the Amazon CloudWatch console, and select an Amazon S3 bucket to view the amount and the number of files that you want to migrate.

CloudWatch Dashboards Alarms	CloudWatch: S3 ~ All resources	Time range 1h 3h 12h 1d 3d Tw custom • Actions • 📿 •
ALARM O	ALARM O INSUFFICIENT DATA O OK O	
Billing Logs	Bucket metrics	
Log groups Insights Metrics Events Rules Event Buses ServiceLens Service Map Traces	65.4G	Number Of Objects Average Count 51.9k 26.0k 1 09/20 03/21 09/20 03/21 09/20 03/21 09/20 03/21
Synthetics Canaries Contributor Insights Settings Favorites		operational issues. The metrics are available at 1 minute intervals after some latency to process. These metrics are billed at w to opt-in to getting these metrics, see Metrics Configurations for Buckets documentation. If you have enabled Request
	S3 buckets Bucket Name	Date Created

(?) Note To ensure a successful migration, you must enter an appropriate amount and number of files when you create a migration job.

• Restore the data in the bucket

Data Online Migration does not restore data in archived files during migration. Before you create a migration job to migrate data of the Archive storage class, you must manually restore the data. Otherwise, the files that are not restored or are being restored cannot be migrated. For more information about how to restore data of the Archive storage class, see <u>Solutions to migrate archived data</u>.

? Note

- Before you create a source data address and a migration job, make sure that data of the Archive storage class is restored.
- The time required to restore data of the Archive storage class is based on the amount of data. Make sure that you specify an appropriate value for the restoration duration. This can prevent data from entering the archived state during migration.
- You may be charged for the restoration operation. For information about the billing methods, contact the service provider that offers the source bucket.
- Create an AccessKey pair

Log on to the AWS IAM console, create an Identity and Access Management console (IAM) user, attach the AmazonS3ReadOnlyAccess permission policy to the IAM user, and then create an AccessKey pair for the IAM user.

? Note

Data Online Migration does not support AWS S3 whitelists. If a whitelist exists, delete it or submit a ticket.

Alibaba Cloud OSS

• Create a destination OSS bucket.

Create a destination bucket to store the migrated data. For more information, see Create buckets.

- Create a Resource Access Management (RAM) user and grant required permissions to the RAM user.
 - i. Log on to the RAM console.
 - ii. In the left-side navigation pane, choose Identities > Users.
 - iii. On the Users page, click Create User.
 - iv. On the **Create User** page, specify Logon Name and Display Name in the User Account Information section.
 - v. In the Access Mode section, select Console Access and OpenAPI Access. Then, save the generated logon name, password, AccessKey ID, and AccessKey secret.
 - a. Console Access: If you select this option, you must configure the console password, password reset settings, and multi-factor authentication settings.
 - b. OpenAPI Access: If you select this option, an AccessKey pair is automatically created for the RAM user. The RAM user can call API operations or use other development tools to access Alibaba Cloud resources.
 - vi. After the RAM user is created, return to the **Users** page. Find the RAM user and click **Add Permissions** in the Actions column. In the panel that appears, select the AliyunOSSFullAccess and AliyunMGWFullAccess policies and click OK. This way, the RAM user is granted the permissions to read and write files from OSS and perform online data migrations.
 - vii. In the left-side navigation pane, click **Overview**.
 - viii. On the page that appears, navigate to the Account Management section and click the link under RAM user logon. On the page that appears, enter the logon name and password of the RAM user to log on to the Alibaba Cloud Management Console.

3.Migrate data

This topic describes how to migrate data by using Data Transport.

Precautions

When you perform an online migration, take note of the following items:

- A migration job occupies the network resources of the source and destination data addresses. To ensure business continuity, we recommend that you specify a rate limit for your migration job or run your migration job during off-peak hours.
- Before an online migration job starts, the system checks the files at the source data address and the destination data address. If the source and destination have files with the same name, and the last modification time of the source file is later than the last modification time of the destination file, the destination file is overwritten. If both files are important and contain different information, we recommend that you change the name of one file or back up the destination file.
- Data Online Migration allows you to migrate only the data of a single bucket per job. You cannot migrate all data that belongs to your account at a time.

Step 1: Create a source data address

- 1. Log on to the Alibaba Cloud Data Transport console.
- 2. In the left-side navigation pane, choose **Data Online Migration > Data Address**. On the page that appears, click **Create Data Address**.
- 3. In the **Create Data Address** panel, configure the parameters and click **OK**. The following table describes the parameters.

Parameter	Required	Description
Data Type	Yes	Select AWS S3.
Data Name	Yes	Enter a name. The name must be 3 to 63 characters in length. The name cannot contain special characters, except for hyphens (-) and underscores (_).
Endpoint	Yes	Enter a region where the source data address is located. For more information, see Regions and endpoints.
Bucket	Yes	The name of an AWS Simple Storage Service (AWS S3) bucket. The bucket name cannot contain invalid characters, such as spaces, line breaks, and tab characters.

Parameter	Required	Description
Prefix	Yes	 Migrate All Data: All the data in the bucket is migrated. If you select this option, you do not need to specify a prefix. Migrate Partial Data: Only files in a specified directory are migrated. If you select this option, you must specify the prefix of the directory. A prefix cannot start with a forward slash (/) but must end with a forward slash (/). For example, you can specify docs/ as the prefix.
Access Key Id and Secret Access Key	Yes	Specify the AccessKey pair that is used to migrate data. Delete the key pair after the migration job is completed.

- 4. (Optional) If the system prompts that this feature is in public preview, click **Application** to apply to use this feature.
- 5. (Optional)Enter the required information and submit the application for the feature in public preview. After the application is approved, the system sends a text message to notify you.

Step 2: Create a destination data address

- 1. In the left-side navigation pane, choose **Data Online Migration > Data Address**. On the page that appears, click **Create Data Address**.
- 2. In the **Create Data Address** panel, configure the parameters and click **OK**. The following table describes the parameters.

Parameter	Required	Description
Data Type	Yes	Select OSS.
Data Name	Yes	Enter a name. The name must be 3 to 63 characters in length. The name cannot contain special characters, except for hyphens (-) and underscores (_).
Data Region	Yes	Select the region where the destination data address is located.

Parameter	Required	Description
Enable Transfer Acceleration	No	For Data Online Migration to use the transfer acceleration feature provided by Object Storage Service (OSS), you must enable the transfer acceleration feature for buckets. The transfer acceleration feature takes effect within 30 minutes after it is enabled. After you enable the feature, wait for 30 minutes before you create a migration job.
		Notice If you use the transfer acceleration feature to access your OSS buckets, you are charged additional fees for transfer acceleration. For more information, see Transfer acceleration .
	Yes	Select an endpoint based on the region where your data is stored. For more information, see Regions and endpoints.
		 http://oss-cn-endpoint.aliyuncs.com: indicates that you use an HTTP-based endpoint to access OSS over the Internet.
OSS Endpoint		 https://oss-cn-endpoint.aliyuncs.com: indicates that you use an HTTPS-based endpoint to access OSS over the Internet.
		Note If you want to migrate data from third-party cloud services to OSS, you must access OSS from a public endpoint over the Internet.
Access Key Id and Access Key Secret	Yes	Specify the AccessKey pair that is used to migrate data. For more information, see Prepare for a migration job.
OSS Bucket	Yes	Select a bucket to store the migrated data.

	No	An OSS prefix cannot start with a forward slash (/) but must end with a forward slash (/). For example, you can specify data/to/oss/ as the OSS prefix. If you want to migrate data to the root directory of the specified bucket, do not specify this parameter.
OSS Prefix No		Notice If the name of a source file starts with a forward slash (/), you must specify an OSS prefix when you configure the destination data address. If no OSS prefix is specified, the migration job fails. For example, if the name of a file to be migrated is <i>/test/test.png</i> , you must specify an OSS prefix, such as <i>oss/</i> . After the migration job is completed, the file name changes from _/t est/test.png to <i>oss//test/test.png</i> .

Step 3: Create a migration job

- 1. In the left-side navigation pane, choose **Data Online Migration > Migration Jobs**. On the page that appears, click **Create Job**.
- 2. In the **Create Job** panel, read the terms of the migration service and select **I understand the above terms and conditions, and apply for opening data migration service**. Then, click **Next**.
- 3. In the Fee Reminder dialog box, click Yes, Go Ahead.
- 4. In the **Job Config** step, configure the parameters and click **Next**. The following table describes the parameters.

Parameter	Required	Description
Job Name	Yes	Enter a job name. The name must be 3 to 63 characters in length and can contain only lowercase letters, digits, and hyphens (-). It cannot start or end with a hyphen (-).
Source Data Address	Yes	Select the source data address that you created.
Destination Data Address	Yes	Select the destination data address that you created.

Parameter	Required	Description
		 Specify directories in which the files and subdirectories are migrated or not migrated. Do not filter: All data from the source data address is migrated. Exclude: The files and subdirectories in the specified directories are not migrated. Contain: Only the files and subdirectories in the specified directories are migrated. Mote
Specified Directory	No	 Before you configure the Specified Directory parameter, you must configure a whitelist. For more information, see Advanced features. A directory name cannot start with a forward slash (/) or a backslash (\), and cannot contain double slashes (//), double periods (), or double quotation marks ("). The character string that consists of all the specified directory names cannot exceed 10 KB in length.
		 A directory name must end with a forward slash (/). For example, you can specify <i>docs/</i> as the directory name. You can specify up to 20 directories of the Exclude or Contain type.
		Select the data migration mode.
		 Full: The system migration mode. Full: The system migrates all data from the source data address to the destination data address. You must configure the Start Time Point of File parameter. The system migrates all data whose last modification time is later than the specified time. After all data is migrated, the migration job ends.
		If you change the data at the source data address after the migration job is completed, you can submit another job to migrate all data. In this case, the system migrates only the changed data.
		 Incremental: The system runs a migration job based on the values that you specified for the Migration Interval and Migration Times parameters.

Parameter	Required	• You must configure the Start Time Point Description of File parameter. During the first migration,
		 the system migrates all data whose last modification time is later than the specified time. After the first migration job is completed, incremental migration jobs are performed based on the value that you specified for the Migration Interval parameter. For each incremental migration, the system migrates only the data that is created or modified after the previous migration starts and before the current migration starts. If you set the Migration Times parameter to N, full migration is performed once and then
		incremental migration is performed (N - 1)
Migration Type	Yes	incremental migration is performed (N - 1) times. For example, the current time is 08:00, March 10, 2019 and you set the Migration Interval parameter to 1, the Migration Times parameter to 5, and the Start Time Point of File parameter to 08:00, March 5, 2019. When the first migration starts, the system migrates all files whose last modification time is between 08:00, March 5, 2019 and 08:00, March 10, 2019. In this example, the first migration requires one hour to complete and the second migration starts at 10:00, March 10, 2019. The duration from 08:00 to 10:00 consists of one hour for the first migration process and one hour for the migration interval. During the second migration, the system migrates the files whose last modification time is between 08:00, March 10, 2019 and 10:00, March 10, 2019. The migration job consists of one full migration and four incremental migrations.

Parameter	Required	Description Before you start a full or
		 incremental migration job, the system compares the files at the source data address with the files at the destination data address. If a source file has the same name as a destination file, the destination file is overwritten when one of the following conditions is met: The Content-Type values of the source file and the destination file are different. The last modification time of the source file is later than that of the destination file. The size of the source file is different from that of the destination file.
Start Time Point of File	Yes	 Specify a time to filter data to be migrated. Data will be migrated if it is created or modified after the specified time. All: All files are migrated. Assign: Only the files that are created or modified after the specified time are migrated. For example, if you set the time to 08:00:00, November 1, 2018, only the files that are created or modified after 08:00:00, November 1, 2018 are migrated.
Migration Interval	This parameter is required when you set the Migration Type parameter to Incremental.	The default value is 1 and the maximum value is 24. Unit: hours.
Migration Times	This parameter is required when you set the Migration Type parameter to Incremental.	The default value is 1 and the maximum value is 30.

Parameter	Required	Description
		The method that is used to process the source file whose name is the same as that of a destination file. Valid values:
		 LastModified: If a source file has the same name as a destination file, the system compares the last modification time that is indicated by the LastModified property of the files.
		 If the last modification time of the source file is later than that of the destination file, the source file is migrated and the destination file is overwritten.
		If the last modification time of the source file is earlier than that of the destination file, the source file is not migrated and the destination file is retained.
		 If the two files have the same last modification time, the system checks the values of the Size and Content-Type properties of the files.
File Overwrite Method	Yes	If the files have the same Size and Content- Type values, the source file is not migrated. If the two files have different values of at least one property, the source file is migrated and the destination file is overwritten.
		 Condition: If a source file has the same name as a destination file, the system compares the values of the LastModified, Size, and Content- Type properties of the two files.
		 If all property values of the two files are the same, the source file is not migrated.
		 If the two files have different values of at least one property, the source file is migrated and the destination file is overwritten.
		 All: If a source file has the same name as a destination file, the system performs no comparison and directly overwrites the destination file with the source file.
		 No: If a source file has the same name as a destination file, the system performs no comparison and directly skips the source file in the migration.

5. In the **Performance** step, navigate to the **Data Prediction** section and configure the **Data Size** and **File Count** parameters.

? Note To ensure a successful migration, we recommend that you estimate the amount of data that the system will migrate as accurately as possible. For more information, see Estimate the amount of data that you want to migrate.

6. (Optional)In the **Performance** step, navigate to the **Flow Control** section, specify the **Time Range** and **Max Flow(MB/s)** parameters, and then click **Add**.

Onte To ensure business continuity, we recommend that you configure the Time Range and Max Flow(MB/s) parameters based on the traffic fluctuations.

7. Click Create and wait until the migration job is completed.

4.Manage migration jobs

This topic describes several subsequent operations that you can perform after you create a migration job.

View the status of a migration job

After you create a migration job, one of the following states is displayed for the job:

- Migrating: Data is being migrated. Wait until the migration job is completed.
- Create Failed: The migration job has failed to be created. You can view the cause of the failure and create another migration job.
- Finished: The migration job is completed. You can view the migration report.
- Failed: The migration job has failed. You can view the migration report and migrate the files that have failed to be migrated.

Modify the throttling settings

During migration, you can modify the throttling settings of a migration job based on your needs.

- 1. Log on to the Data Transport console.
- 2. Choose **Data Online Migration > Migration Jobs**. On the migration job list, find a migration job and click **Manage**.
- 3. Click **Stop** and verify that the job is stopped.
- 4. On the Flow Control Time Schedule chart, click Reset.
 - To add a throttling setting, select an appropriate time range and maximum flow and then click Add.
 - To delete a throttling setting, click 📊 next to the throttling setting.
 - To modify a throttling setting, delete the current setting and then add a new throttling setting.
- 5. Click Submit and then click Start to rerun the migration job.

View a migration report

- 1. On the migration job list, find the migration job and click Manage.
- 2. Click **Generate Migration Report**. After a report is generated, click **Export** to export the migration report.

In a migration report, the following files appear in the File Lists section:

- The file whose name ends with _total_list . This file lists all files to be migrated.
- The file whose name ends with __completed_list . This file lists the files that have been migrated.
- The file whose name ends with __error_list . This file lists the files that have failed to be migrated.
- 3. In the Object Storage Service (OSS) console, find the automatically generated folder *aliyun_mgw_i mport_report/*. The folder includes the three files that appear in the migration report. You can download these files and view the detailed list that is included in each file. We recommend that you use the ossbrowser tool to view these files.

Each list consists of the following columns:

- The columns of the file whose name ends with _total_list include the source data address, file name, source version ID, and destination version ID. The columns also display whether the job version is the latest at the source data address and whether the job version is deleted at the source data address. The last two columns display the file size (measured in bytes) and the time when the file was last modified. The format of a source data address is <a href="https://www.verdors.com/weatherthestinations.c
- The columns of the file whose name ends with _completed_list include the file name, the source version ID, and the destination version ID. The columns also display whether the job version is the latest at the source data address and whether the job version is deleted at the source data address. The last four columns display the file size (measured in bytes), CRC-64 checksum, time when the migration job was completed, and whether the migration was skipped.

(?) Note Skip migration: indicates whether the file that has the same name at the source data address and the destination data address has been skipped during migration. Valid values: 1 and 0. The value 1 indicates that the file has been skipped during migration. The value 0 indicates that the file has not been skipped during migration. The value of this field depends on the value of the File Overwrite Method parameter that you specify when you create the migration job.

• The columns of the file whose name ends with _error_list include the file name, the source version ID, and the destination version ID. The columns also display whether the job version is the latest at the source data address and whether the job version is deleted at the source data address. The last three columns display the time when the migration job started, the time when the migration job ended, and the error description.

? Note The following four fields apply only if multi-version migration is enabled:

- Source version ID: the version ID of the file at the source data address.
- Destination version ID: the version ID of the file at the destination data address.
- Latest at source data address: indicates whether the current version is the latest version of the file at the source data address. Valid values: true and false.
- Deleted at source data address: indicates whether the current version is deleted at the source data address. Valid values: true and false.

Rerun a failed migration job

If a migration job fails, you can check the generated file whose name ends with _error_list to identify the cause of the failure and resolve the issue. Then, on the migration job list, find the failed migration job, click **Manage**, and click **Retry** to migrate the failed files.

5.Regions and endpoints

Supported regions and endpoints for Amazon Simple Storage Service (Amazon S3) are listed in the following table.

RegionName	Region	Endpoint	Protocol
US East (Ohio)	us-east-2	http://s3.us-east- 2.amazonaws.com http://s3-us-east- 2.amazonaws.com	HTTP/HTTPS
US East (N. Virginia)	us-east-1	http://s3.us-east- 1.amazonaws.com http://s3-us-east- 1.amazonaws.com	HTTP/HTTPS
US West (N. California)	us-west-1	http://s3.us-west- 1.amazonaws.com http://s3-us-west- 1.amazonaws.com	HTTP/HTTPS
US West (Oregon)	us-west-2	http://s3.us-west- 2.amazonaws.com http://s3-us-west- 2.amazonaws.com	HTTP/HTTPS
Canada (Central)	ca-central-1	http://s3.ca-central- 1.amazonaws.com http://s3-ca-central- 1.amazonaws.com	HTTP/HTTPS
Asia Pacific (Seoul)	ap-northeast-2	http://s3.ap-northeast- 2.amazonaws.com http://s3-ap- northeast- 2.amazonaws.com	HTTP/HTTPS
Asia Pacific (Osaka- Local)	ap-northeast-3	http://s3.ap-northeast- 3.amazonaws.com http://s3-ap- northeast- 3.amazonaws.com	HTTP/HTTPS

Migrate data from Amazon Simple S torage Service (Amazon S3) to OSS-Regions and endpoints

RegionName	Region	Endpoint	Protocol
Asia Pacific (Singapore)	ap-southeast-1	http://s3.ap- southeast- 1.amazonaws.com http://s3-ap- southeast- 1.amazonaws.com	HTTP/HTTPS
Asia Pacific (Sydney)	ap-southeast-2	http://s3.ap- southeast- 2.amazonaws.com http://s3-ap- southeast- 2.amazonaws.com	HTTP/HTTPS
Asia Pacific (Tokyo)	ap-northeast-1	http://s3.ap-northeast- 1.amazonaws.com http://s3-ap- northeast- 1.amazonaws.com	HTTP/HTTPS
China (Beijing)	cn-north-1	http://s3.cn-north- 1.amazonaws.com.cn	HTTP/HTTPS
China (Ningxia)	cn-northwest-1	http://s3.cn- northwest- 1.amazonaws.com.cn	HTTP/HTTPS
EU (Frankfurt)	eu-central-1	http://s3.eu-central- 1.amazonaws.com Http://s3-eu-central- 1.amazonaws.com	HTTP/HTTPS
EU (Ireland)	eu-west-1	http://s3.eu-west- 1.amazonaws.com http://s3-eu-west- 1.amazonaws.com	HTTP/HTTPS
EU (London)	eu-west-2	http://s3.eu-west- 2.amazonaws.com http://s3-eu-west- 2.amazonaws.com	HTTP/HTTPS
EU (Paris)	eu-west-3	http://s3.eu-west- 3.amazonaws.com http://s3-eu-west- 3.amazonaws.com	HTTP/HTTPS

Data Online Migration

RegionName	Region	Endpoint	Protocol
South America (São Paulo)	sa-east-1	http://s3.sa-east- 1.amazonaws.com http://s3-sa-east- 1.amazonaws.com	HTTP/HTTPS