ALIBABA CLOUD

Alibaba Cloud

Data Online Migration

Case Study

Document Version: 20211118

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.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized 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 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
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.
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.
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.
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 list --instanceid` *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.Seamlessly migrate data of a web-based service provider to OSS

This topic describes how to migrate the data of a web-based service provider from a cloud service to Alibaba Cloud Object Storage Service (OSS).

Background information

Enterprise A is a web-based service provider that offers services for editing media files, such as images and videos. The main business applications are deployed in a cloud-based architecture provided by Enterprise B. The data that is stored on the servers of Enterprise B includes 100,000,000 files and has a total size of approximately 320 TB. The size of the data increases by 20 GB each day. The bandwidth for the storage service of Enterprise B 250 MB/s. The bandwidth for OSS is also 250 MB/s. The business applications require a maximum bandwidth of 50 MB/s.

To accelerate business growth, Enterprise A wants to move the business applications to OSS. Both the existing data and incremental data need to be migrated to OSS. To ensure the successful migration of large amounts of historical data and business continuity, the following requirements must be met:

- During the migration, users can read data as normal.
- After the migration, the data is complete and the business applications can run as expected without service interruptions.

Migration solutions

The following data migration solution is based on the customer requirements and background information:

1. Use Data Online Migration to migrate the existing data of Enterprise A from a cloud service provided by Enterprise B to OSS. Make sure that the customer updates no data during the entire migration process.
2. After the existing data is migrated, create back-to-origin rules in OSS for users to access the incremental data that is migrated.
3. Switch the business applications to OSS.
4. After the business applications are switched to OSS, use Data Online Migration to migrate incremental data to OSS.
5. After all data is migrated and validated, delete the data from the original source.

Step 1: Migrate the existing data

1. Create an OSS bucket to store the migrated data. For more information, see Create buckets.
2. Obtain the AccessKey pairs that are used to migrate data:
   - To obtain the AccessKey pair provided by Enterprise B, log on to the cloud service console and view the AccessKey pair.
   - Obtain the AccessKey pair of the OSS Resource Access Management (RAM) user. For more information, see Create a RAM user and grant permissions to the RAM user.
3. Create data addresses and a full migration job. For more information, see the topics about data migration tutorials in the documentation of Data Online Migration.

   In the **Job Config** step, configure the migration job. The following figure provides an example of

the configurations.

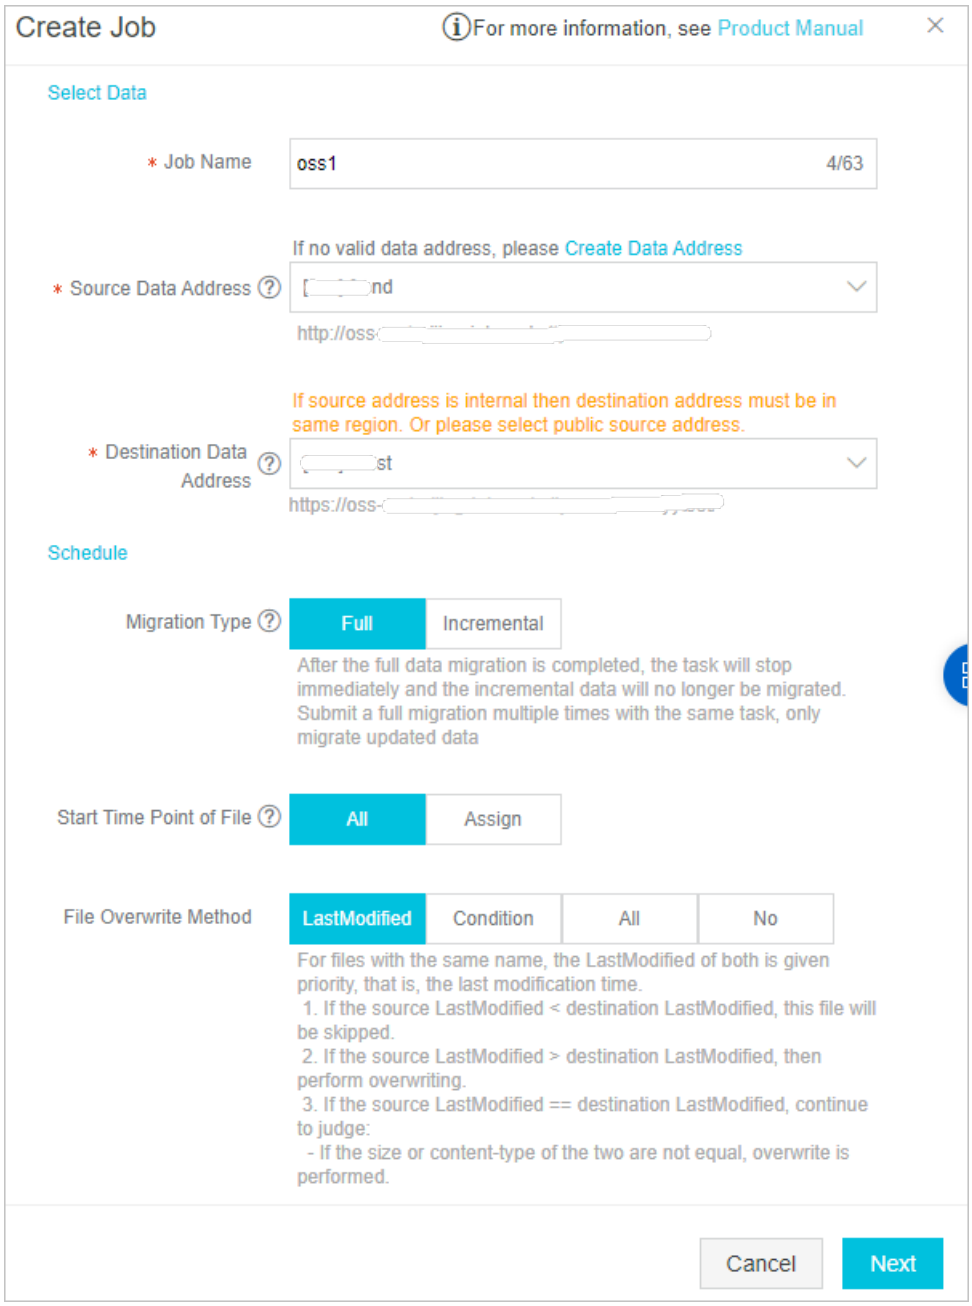

In the **Performance** step, configure the performance parameters. The following figure provides an example of the configurations.

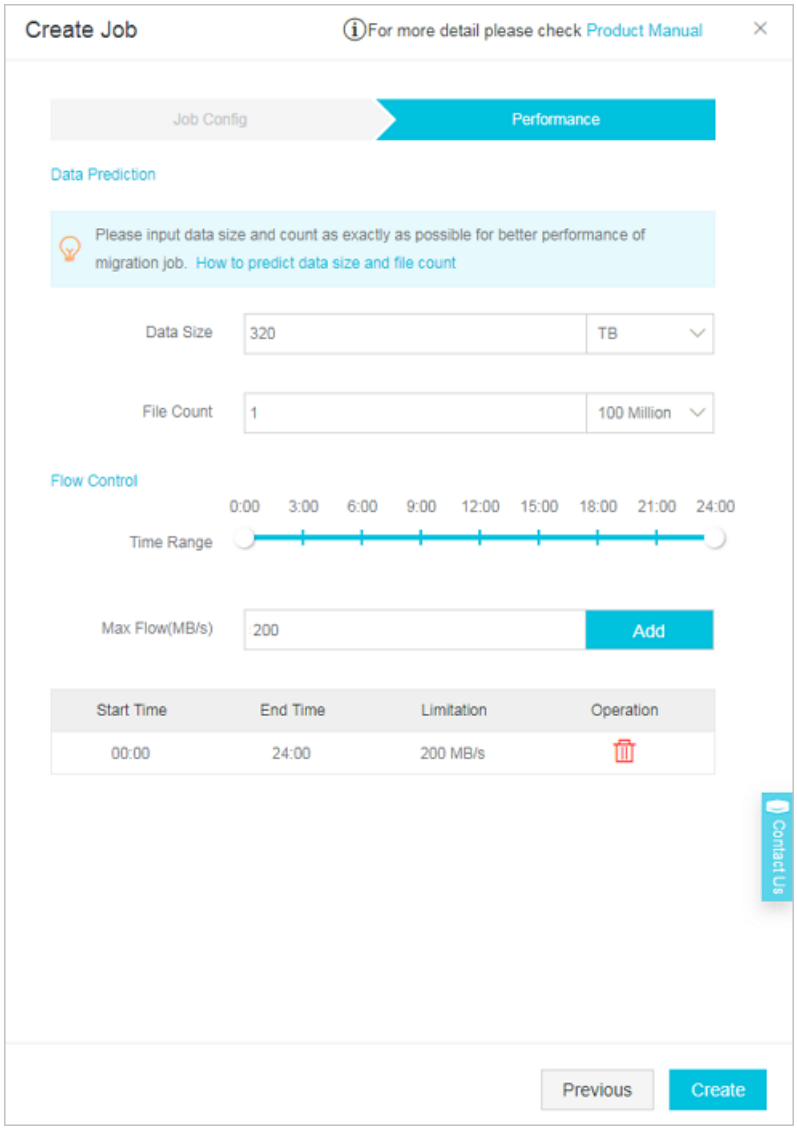

4. After the migration, view the migration report and compare data between the source data address and the destination data address to ensure that all data is migrated.

> **Note** If a file fails to be migrated, troubleshoot the failure. For information about migration failures, see Common causes of a migration failure and solutions.

Step 2: Create back-to-origin rules

It takes approximately 25 days to migrate the existing data. During the migration process, data at the source data address is continuously growing. To ensure business continuity and a seamless business switchover, create back-to-origin rules. If a file that you request does not exist in OSS, OSS fetches the file from its source data address and returns it to you based on the back-to-origin rules.

1. Log on to the OSS console.
2. In the list of buckets, select the bucket in which the migrated data is stored.
3. In the left-navigation pane, choose **Basic Settings** > **Back-to-Origin**. Then, click **Configure**.

4. Click **Create Rule**. In the **Create Rule** dialog box, configure the parameters.

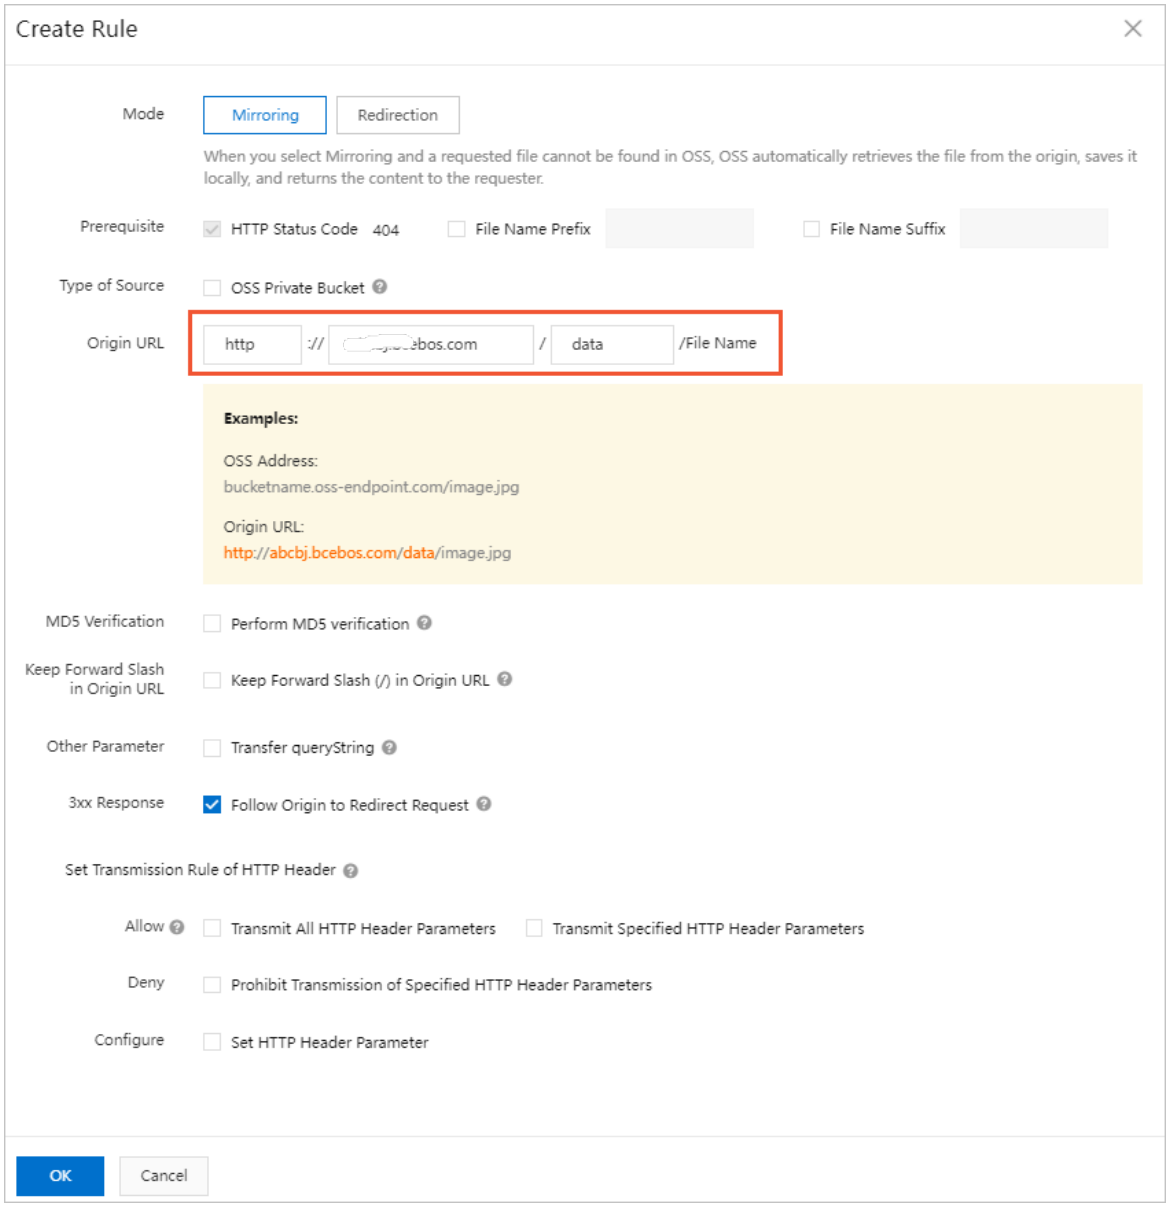

- **Mode**: Select **Mirroring**.
- **Prerequisite**: **HTTP Status Code 404** is selected by default. Configure **File Name Prefix** based on your needs. File Name Prefix can be left empty.
- **Origin URL**: Enter the endpoint of the cloud service where the source data resides.
- For more information about the parameter configuration, see Create back-to-origin rules.

> **Note** You can create a maximum of five back-to-origin rules. The five rules can be in effect at the same time. For multiple source data addresses, you can create multiple back-to-origin rules. You can set different rules by specifying different values for **File Name Prefix** so that OSS can fetch various types of data.

5. Click **OK**.

Step 3: Switch the business to OSS

Change the data address from which Enterprise A obtains data to an OSS data address.

Step 4: Migrate the incremental data

During the migration of the existing data, approximately 100,000 files with a total size of about 500 GB are generated at the source data address. You must migrate these incremental data to OSS.

1. Create an incremental migration job based on the instructions in Step 1.

   In the **Job Config** step, configure the migration job. The following figure provides an example of the configurations.

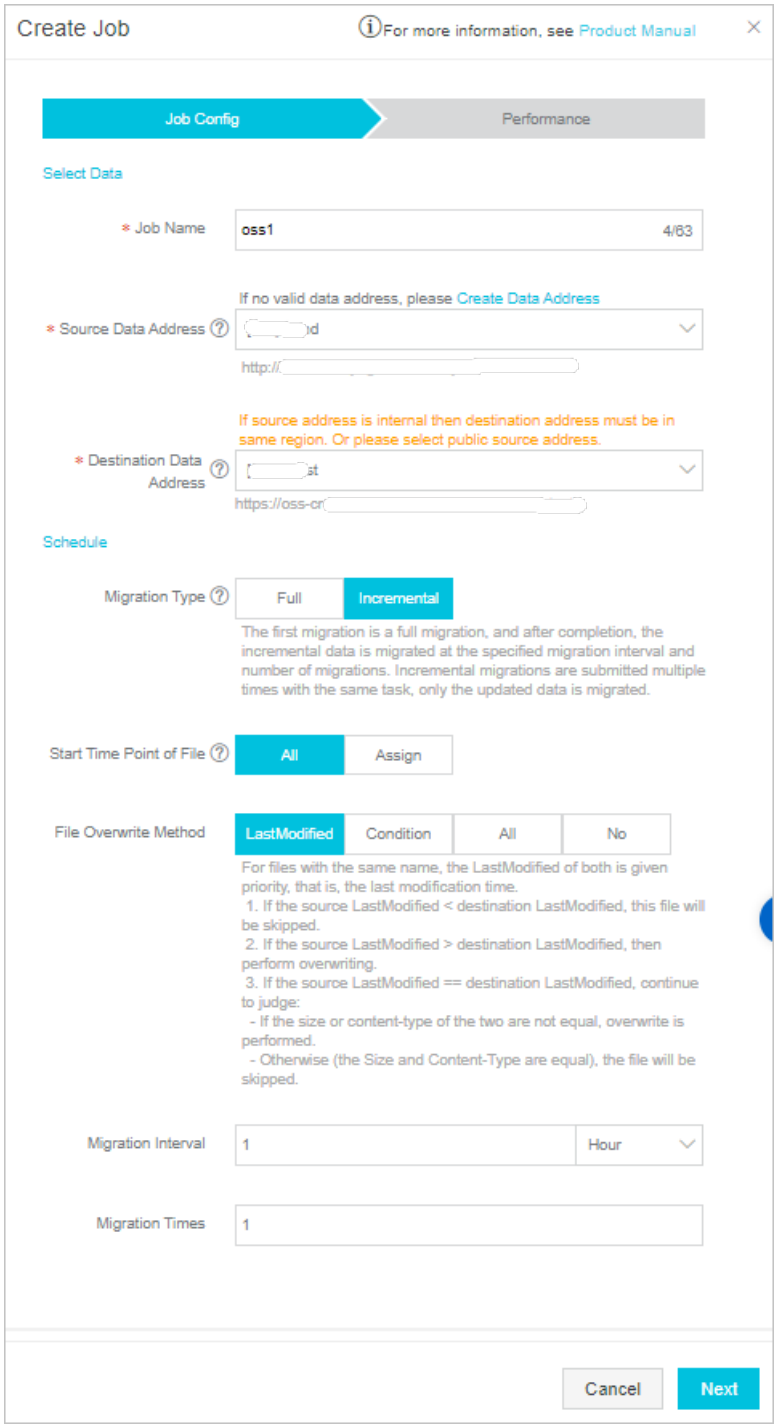

   In the **Performance** step, configure the performance parameters. The following figure provides an

example of the configurations.

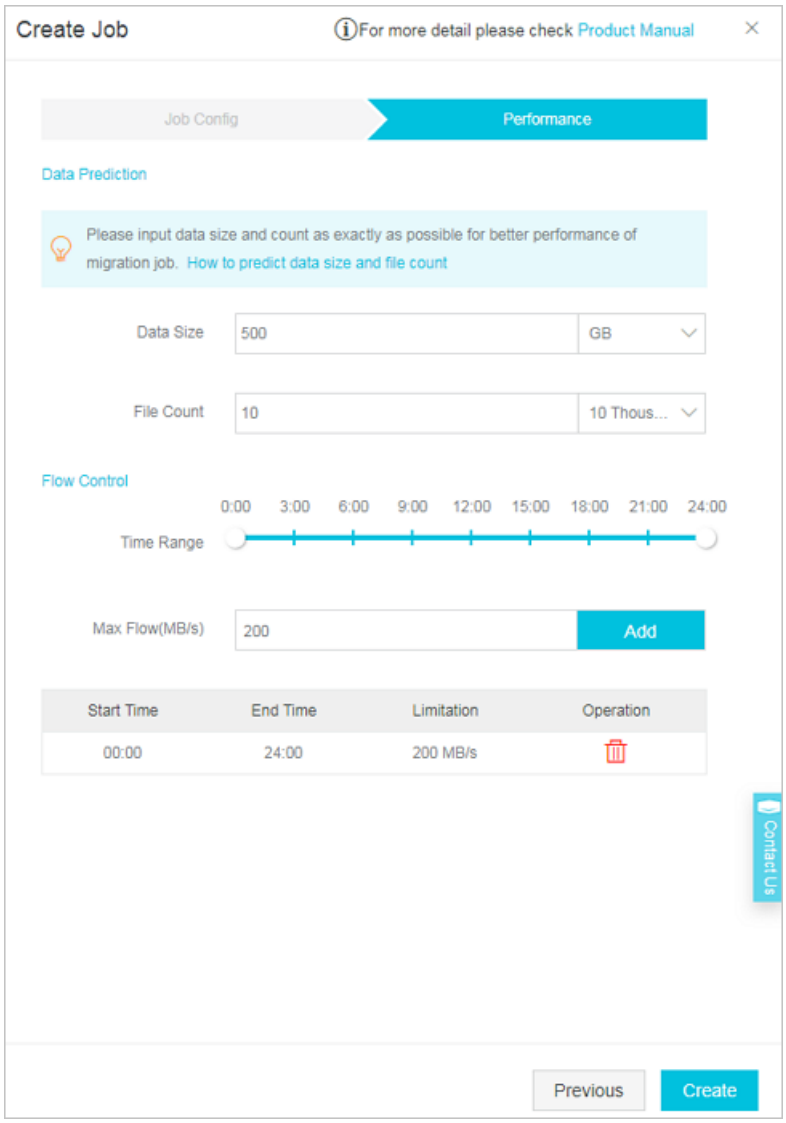

2. Click **Create** to migrate data.
3. After the migration, view the migration report and compare data between the source data address and the destination data address to ensure that all data is migrated.

> **Note** If a file fails to be migrated, troubleshoot the failure. For information about migration failures, see Common causes of a migration failure and solutions.

Step 5: Delete the data at the source data address

You can create a lifecycle rule to delete the files at the source data address one day after all data is migrated. This way, you are no longer charged for storing the data at the source data address one day after the data is migrated.

2.Migrate data from an on-premises NAS file system to OSS for an entertainment company

This topic describes how to migrate data from an internal Network Attached Storage (NAS) file system of an entertainment company in Hangzhou to Alibaba Cloud Object Storage Service (OSS) for long-term storage.

Background information

An entertainment company in Hangzhou stores data, such as media files and documents, on its internal NAS file server. The data includes 5,000,000 files, which are about 20 TB in size. The NAS server is located in the data center of the company. The server uses the Server Message Block (SMB) protocol and has a firewall installed. The server cannot be accessed over the Internet but provides an internal IP address of 10.0.0.254.

To meet the requirements for subsequent maintenance and online application development, the company needs to migrate data from the NAS server to OSS for long-term storage.

Migration scheme

Based on the user requirements and background information, the following migration scheme is formulated:

1. Create an OSS bucket in the China (Hangzhou) region and change the default storage location to the endpoint of this bucket.
2. Use a dedicated leased line to connect the on-premises NAS server to an Alibaba Cloud virtual private cloud (VPC). Modify the firewall settings of the NAS server and enable access to the NAS server from all the IP addresses in the VPC.
3. Use Data Online Migration to migrate data from the NAS server to OSS.

Step 1: Create a bucket and change the storage location

1. In the China (Hangzhou) region, create a bucket to store data. For more information, see Create buckets.
2. Set the bucket policy to allow the employees of the company to access the bucket. For more information about the configurations, see Configure bucket policies to authorize other users to access OSS resources.
3. Inform the internal employees to change the default storage location to the bucket.

Step 2: Connect the NAS server to an Alibaba Cloud VPC

1. Use a dedicated leased line that provides a transmission rate of 1 Gbit/s to connect the NAS server to an Alibaba Cloud VPC. For more information, see Create a dedicated connection over an Express Connect circuit.
2. Modify the firewall settings of the NAS server to enable access to the NAS server from all the IP addresses in the VPC.

Step 3: Use Data Online Migration to migrate data from the NAS server to OSS

1. Create a Resource Access Management (RAM) user in the Alibaba Cloud Management Console, grant the RAM user the permissions to create migration jobs, and obtain the AccessKey pair of the RAM user. For more information, see Create and authorize a RAM user.
2. Create a NAS data address. For more information about the parameters, see Migrate data from NAS to OSS.

   The following figure shows the configuration details.

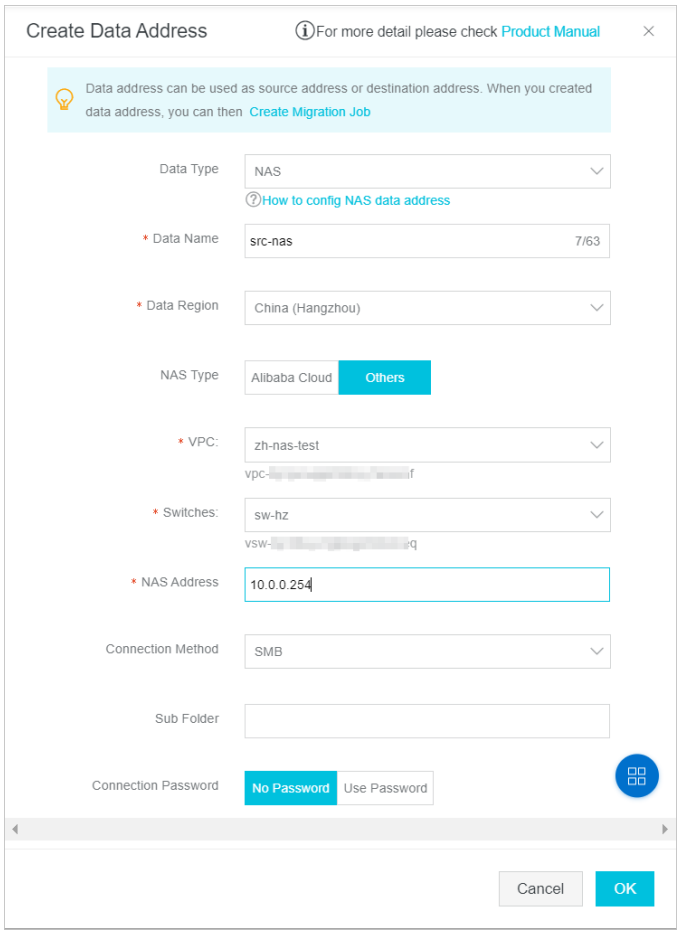

3. Create an OSS data address. For more information about the parameters, see Migrate data from NAS to OSS.

   The following figure shows the configuration details.

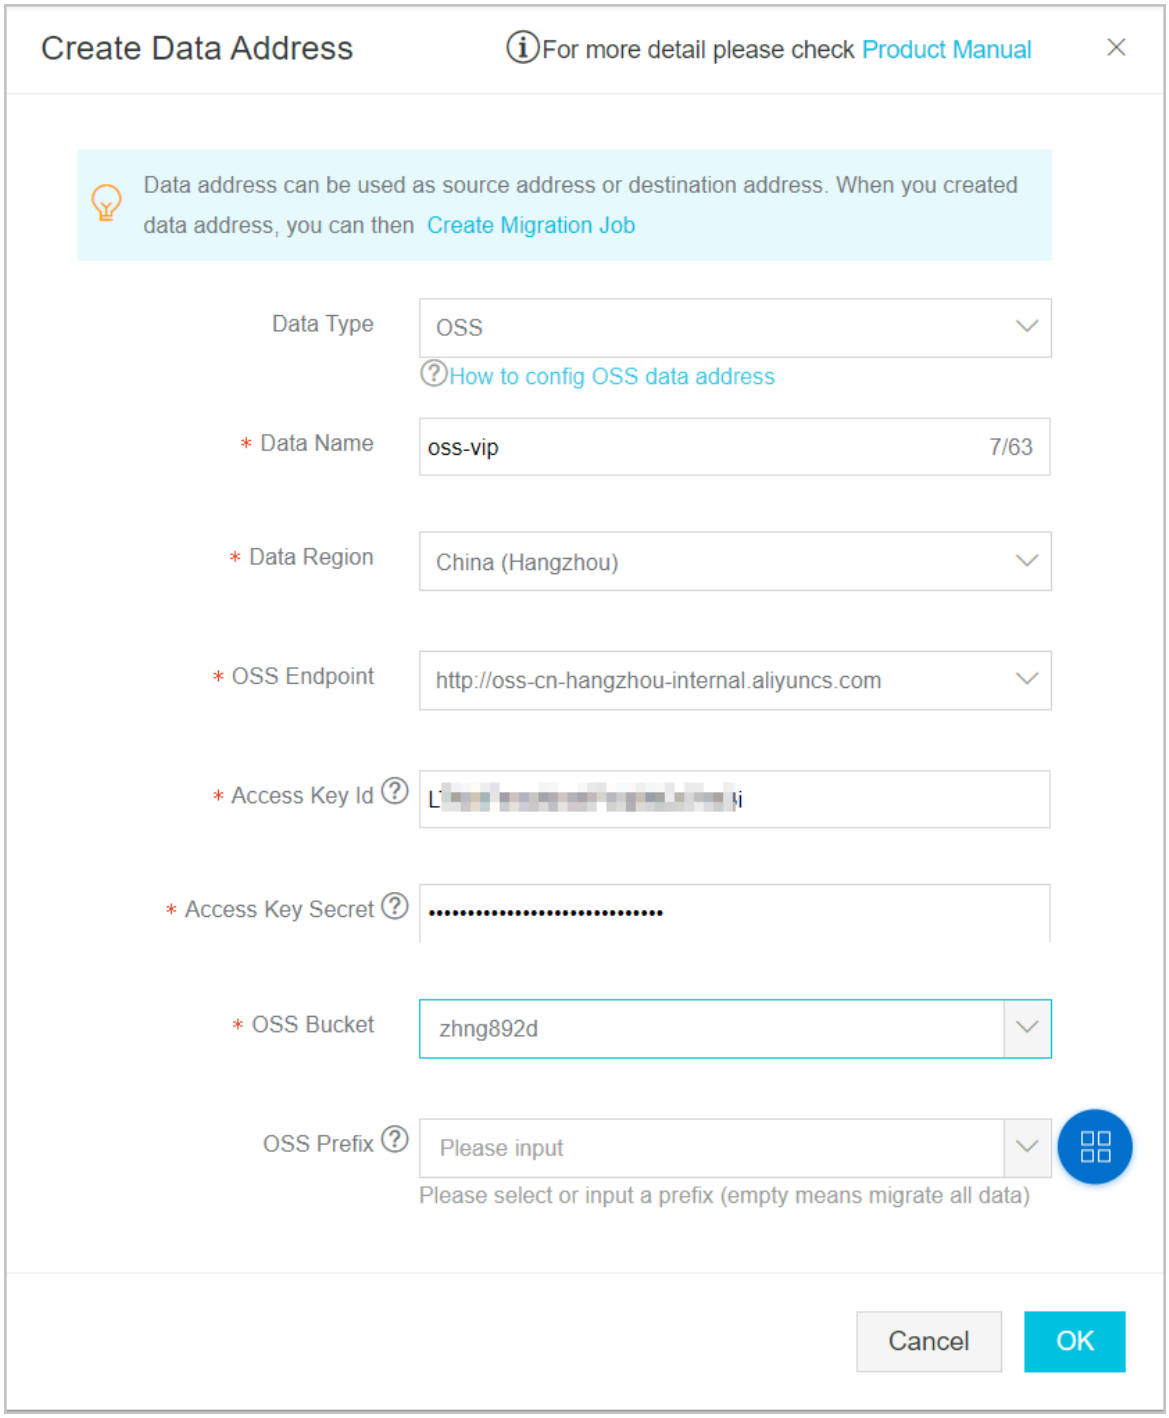

4. Create a full migration job and configure performance optimization. For more information about the parameters, see Migrate data from NAS to OSS.

   > **Notice** In this example, the entertainment company has no bandwidth needs for other applications during data migration. Therefore, no flow limit is set. In actual practice, you can set appropriate flow limits based on the usage of the bandwidth.

   The following figure shows how to configure performance optimization.

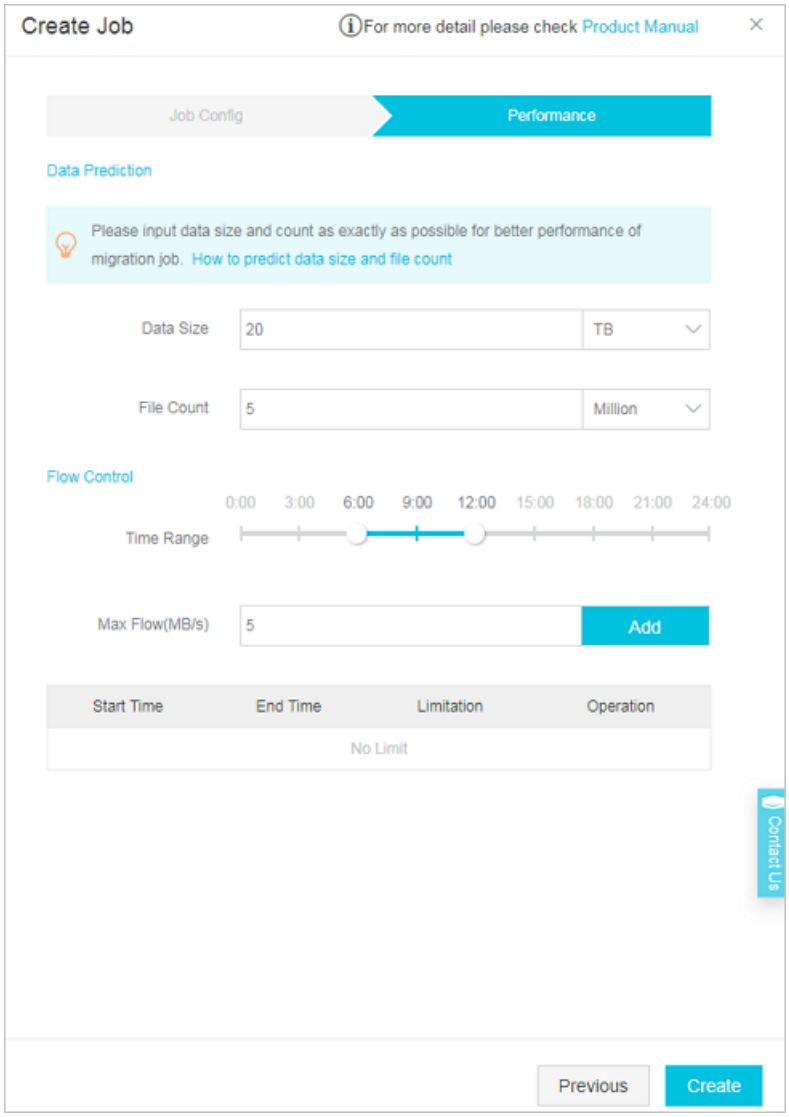

5. A migration job requires about two days to complete. After the migration job is completed, you must verify that all data is migrated. To do so, view the migration report and compare the data at the source data address with the data at the destination data address.

   > **Note** If a file fails to be migrated, troubleshoot the failure. For more information, see Common causes of a migration failure and solutions.

   After the data is migrated, you can store and manage data on OSS.

3.Migrate data between NAS file systems that are located in different VPCs for a company

This topic describes how to migrate data of a company between Network Attached Storage (NAS) file systems that are located in different virtual private clouds (VPCs).

Background information

A Shenzhen company is named Company A. As Company A grows and expands, it establishes a subsidiary in Hangzhou. The subsidiary is named Branch B. Branch B stores data in a separate Apsara File Storage NAS file system. Branch B must synchronize the data to the Apsara File Storage NAS file system of Company A on a daily basis. Each day, Branch B generates about 100,000 files whose size is about 100 GB.

The Apsara File Storage NAS file systems of Company A and Branch B are located in different Alibaba Cloud VPCs. The CIDR block of the VPC where the NAS file system of Company A is located is 172.16.1.0/24. The CIDR block of the VPC where the NAS file system of Branch B is located is 10.0.0.0/24.

> **Note** If you are using a third-party NAS file system, you must use a dedicated leased line to connect your NAS server to an Alibaba Cloud VPC. For more information, see Create a dedicated connection over an Express Connect circuit.

Migration scheme

1. Use Cloud Enterprise Network (CEN) to establish a connection between the two VPCs of Company A and Branch B and configure permission groups. Make sure that all the addresses within the VPC of Branch B have the following permissions: the read-only access to the NAS file system of Branch B and the read/write access to the NAS file system of Company A.
2. Create a migration job to synchronize the data of Branch B to Company A on a regular basis.

Step 1: Connect the VPCs of Company A and Branch B by using CEN

1. Use CEN to connect the VPCs of Company A and Branch B. For more information, see Connect VPCs that are located in multiple regions and owned by different accounts.
2. Modify the NAS permission groups of Company A and Branch B. This allows all the devices in the 10.0.0.0/24 CIDR block to read data from the NAS file system of Branch B and write data to the NAS file system of Company A. For more information, see Manage a permission group.

Step 2: Create a migration job

1. Create a Resource Access Management (RAM) user in the Alibaba Cloud Management Console and grant the RAM user the permissions to create migration jobs. For more information, see Create and authorize a RAM user.
2. Create a source NAS data address. For more information about the parameters, see Create a source data address. The following figure shows the configuration details.

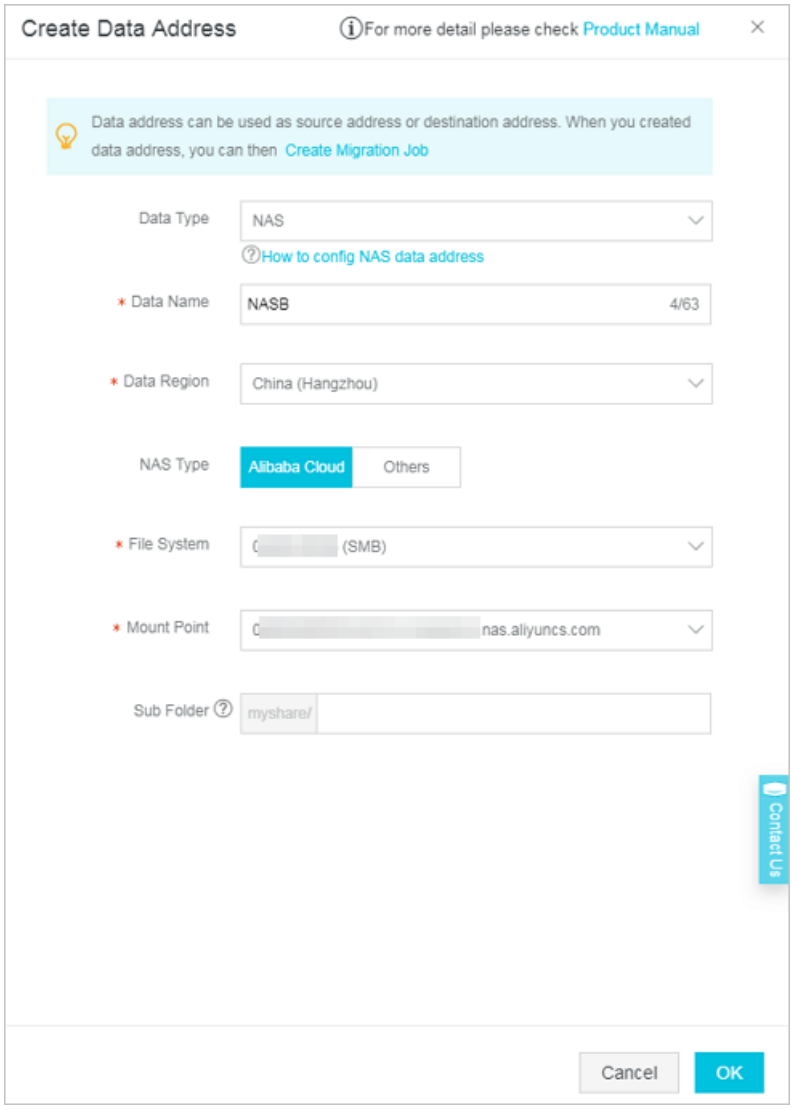

3. Create a destination NAS data address. For more information about the parameters, see Create a destination data address. The following figure shows the configuration details.

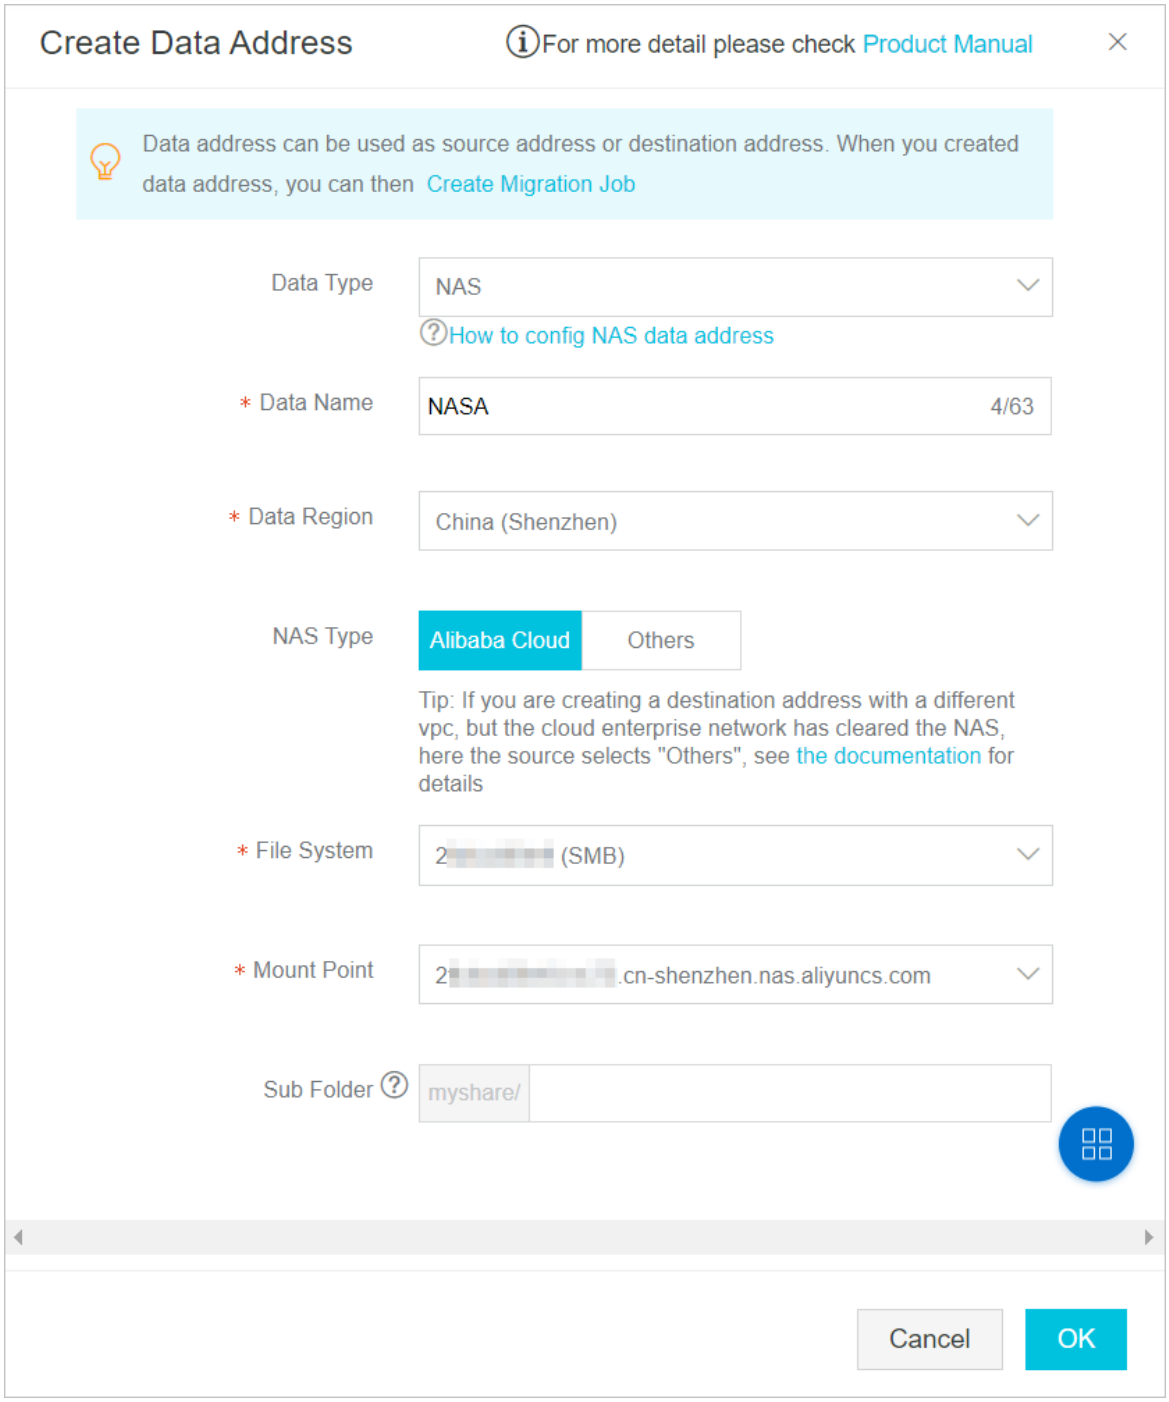

4. Create a migration job of the **Sync** type. To ensure business continuity, set the daily start time of a synchronization job to 22:00:00. For more information about the parameters, see Create a migration job. The following figure shows the configuration details.

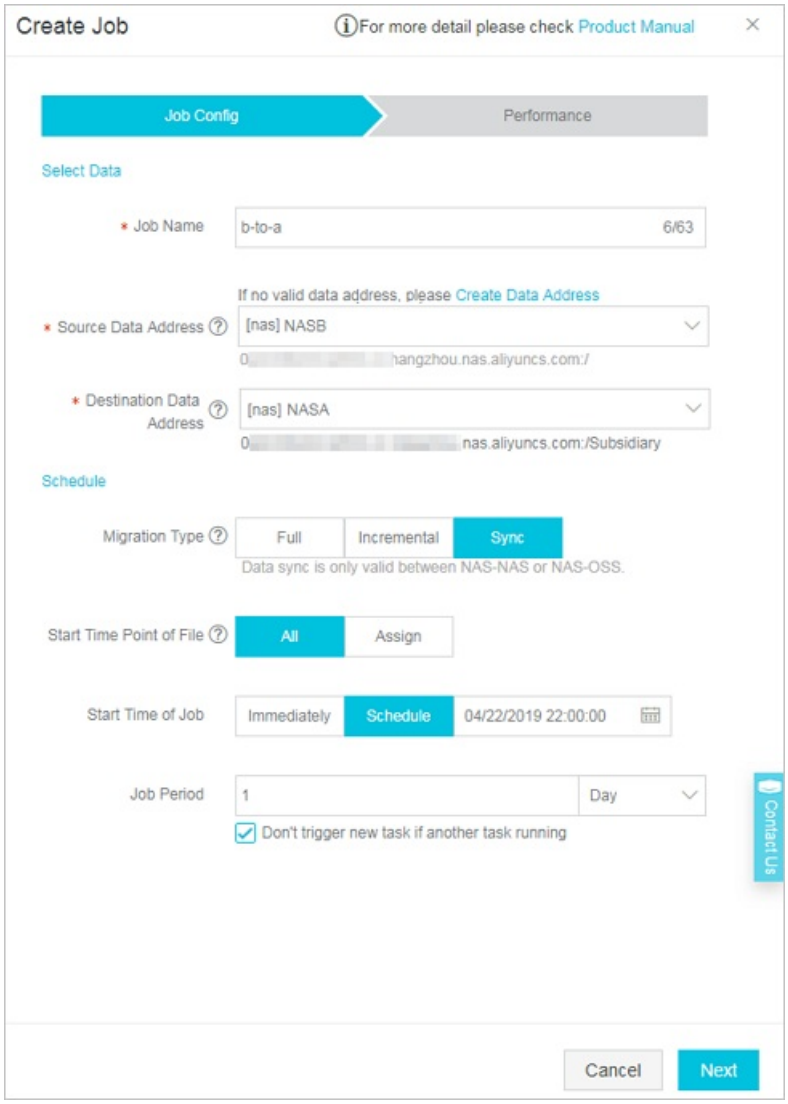

Notice

- A synchronization job keeps running until you stop the job. Therefore, to synchronize data on a regular basis, you need only to create one synchronization job.
- In this example, the default settings in the Performance step are used because the customer synchronizes a small amount of data during off-peak hours. In actual practice, you need to specify appropriate performance parameters based on your needs.

5. After each synchronization job is completed, you can check the status of the job and compare the data at the source data address and the data at the destination data address. This allows you to verify that all data is migrated. For more information about how to view the status of synchronization jobs, see Manage synchronization jobs.

4.Migrate data from an on-premises NAS file system to Apsara File Storage NAS for a pharmaceutical company

This topic describes how to migrate data from an on-premises Network Attached Storage (NAS) server to Apsara File Storage NAS for long-term storage.

Background information

A Hangzhou pharmaceutical company stores data, such as product documents and experimental data, on its internal NAS file server. The data includes 10,000,000 files, which are about 10 TB in size. The NAS server is located in the data center of the company. The server uses the Network File System (NFS) protocol and has a firewall installed. The server cannot be accessed over the Internet but provides an internal IP address of 10.0.0.254.

For data security and cost saving, the company needs to migrate data from the NAS server to Apsara File Storage NAS.

Migration scheme

Based on the user requirements and background information, the following migration scheme is formulated:

1. Create an Apsara File Storage NAS file system in the China (Hangzhou) region and attach the file system to an Alibaba Cloud virtual private cloud (VPC).
2. Use a dedicated leased line to connect the on-premises NAS server to the VPC. Modify the firewall settings of the NAS server and enable access to the NAS server from all the IP addresses in the VPC.
3. Use Data Online Migration to migrate data from the on-premises NAS server to Apsara File Storage NAS.

Step 1: Create an Apsara File Storage NAS file system

1. In the China (Hangzhou) region, create an Apsara File Storage NAS file system that uses the NFS protocol type. For more information, see Mount an NFS file system on a Linux ECS instance.
2. Attach the Apsara File Storage NAS file system to a VPC. For more information, see Mount an NFS file system on a Linux ECS instance.
3. Modify the security group of the VPC to enable read/write access to the NAS file system from all the IP addresses in the VPC. For more information, see Manage a permission group.

Step 2: Connect the NAS server to the Alibaba Cloud VPC

1. Use a dedicated leased line that provides a transmission rate of 1 Gbit/s to connect the NAS server to the VPC where the Apsara File Storage NAS file system is located. For more information, see Create a dedicated connection over an Express Connect circuit.
2. Modify the firewall settings of the NAS server to enable access to the NAS server from all the IP addresses in the VPC.

Step 3: Create a migration job

1. Create a Resource Access Management (RAM) user in the Alibaba Cloud Management Console and grant the RAM user the permissions to create migration jobs. For more information, see Create and authorize a RAM user.
2. Use the RAM user to log on to the Data Transport console.
3. Create a source data address based on the information about the on-premises NAS server. For more information about the parameters, see Create a source data address. The following figure shows the configuration details.

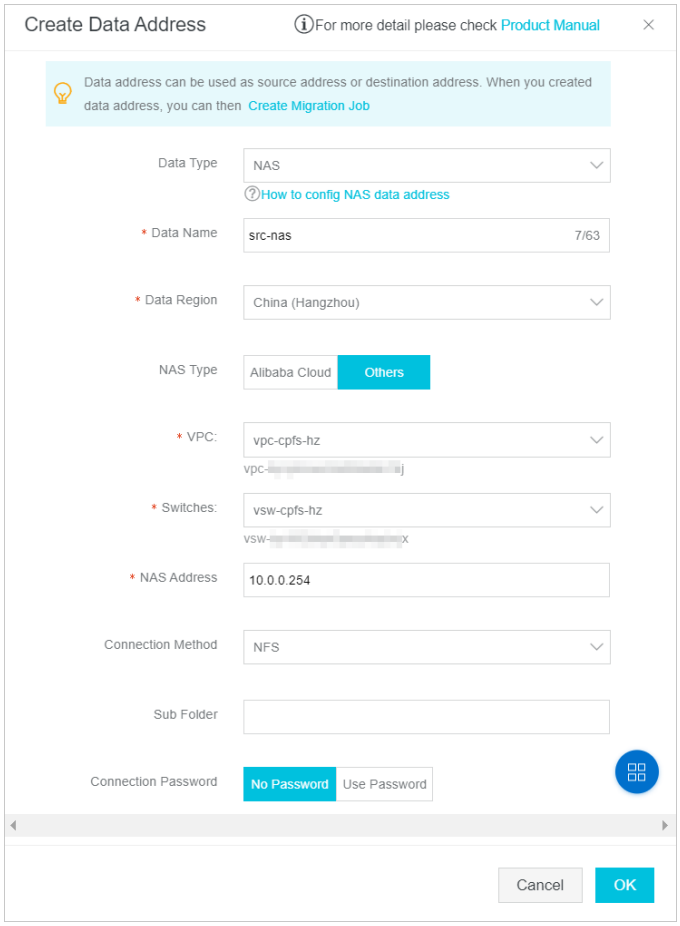

4. Create a destination data address based on the information about the Apsara File Storage NAS file system. The following figure shows the configuration details.

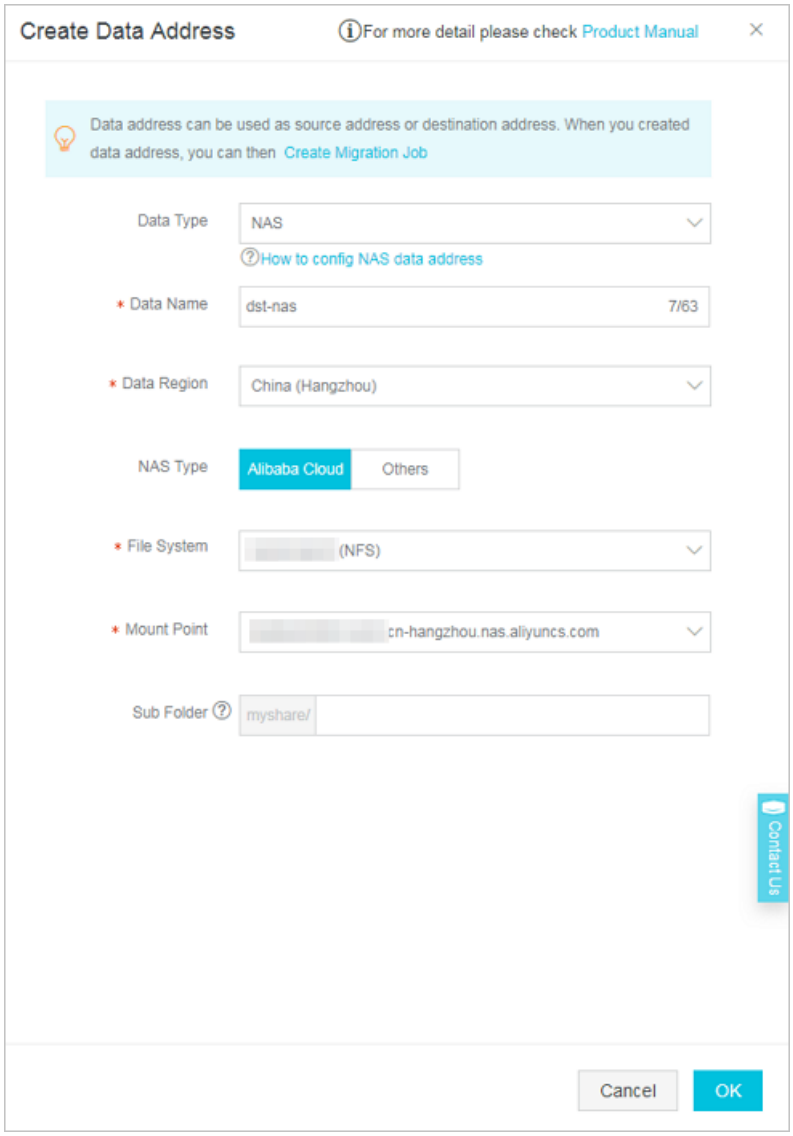

5. Create a migration job of the **Full** type to migrate data from the on-premises NAS server to Apsara File Storage NAS. For more information about the parameters, see Create a migration job. The following figure shows the configuration details.

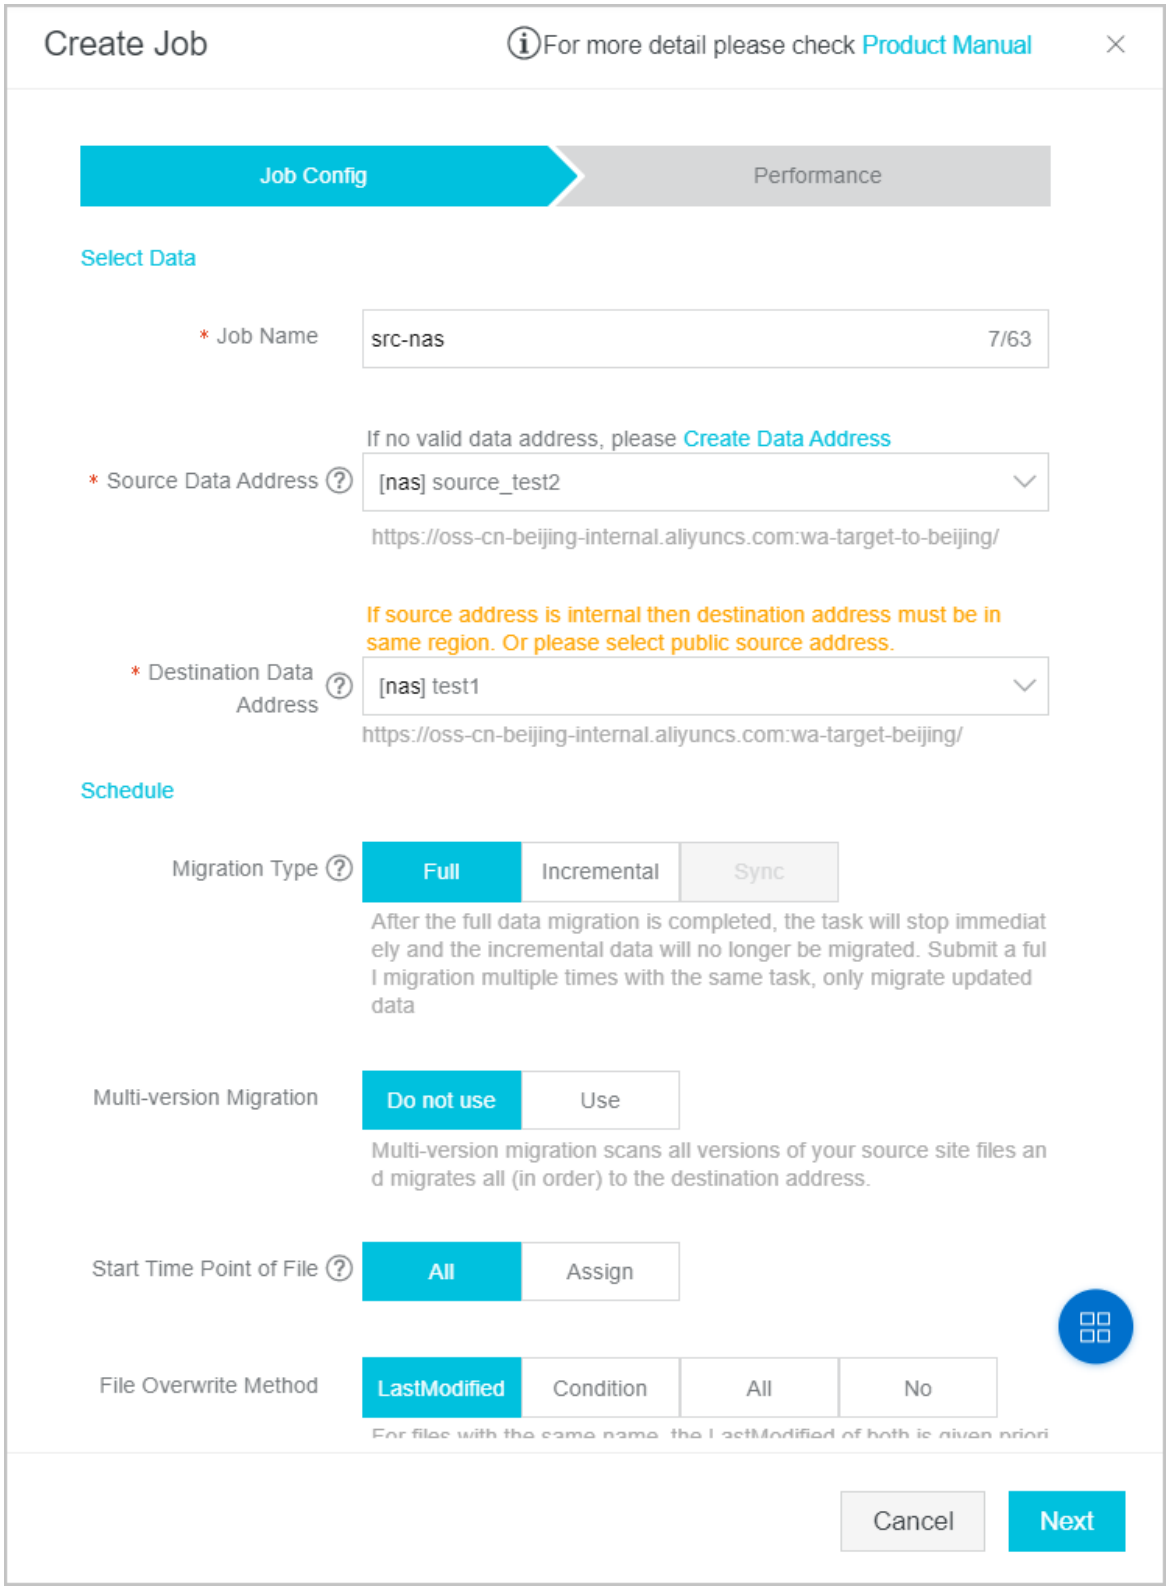
Create Job
For more detail please check Product Manual
Job Config
Performance
Select Data
* Job Name
src-nas
7/63
If no valid data address, please Create Data Address
* Source Data Address
[nas] source_test2
https://oss-cn-beijing-internal.aliyuncs.com:wa-target-to-beijing/
If source address is internal then destination address must be in same region. Or please select public source address.
* Destination Data Address
[nas] test1
https://oss-cn-beijing-internal.aliyuncs.com:wa-target-beijing/
Schedule
Migration Type
Full
Incremental
Sync
After the full data migration is completed, the task will stop immediately and the incremental data will no longer be migrated. Submit a full migration multiple times with the same task, only migrate updated data
Multi-version Migration
Do not use
Use
Multi-version migration scans all versions of your source site files and migrates all (in order) to the destination address.
Start Time Point of File
All
Assign
File Overwrite Method
LastModified
Condition
All
No
Cancel
Next

The following figure shows you how to configure performance optimization.

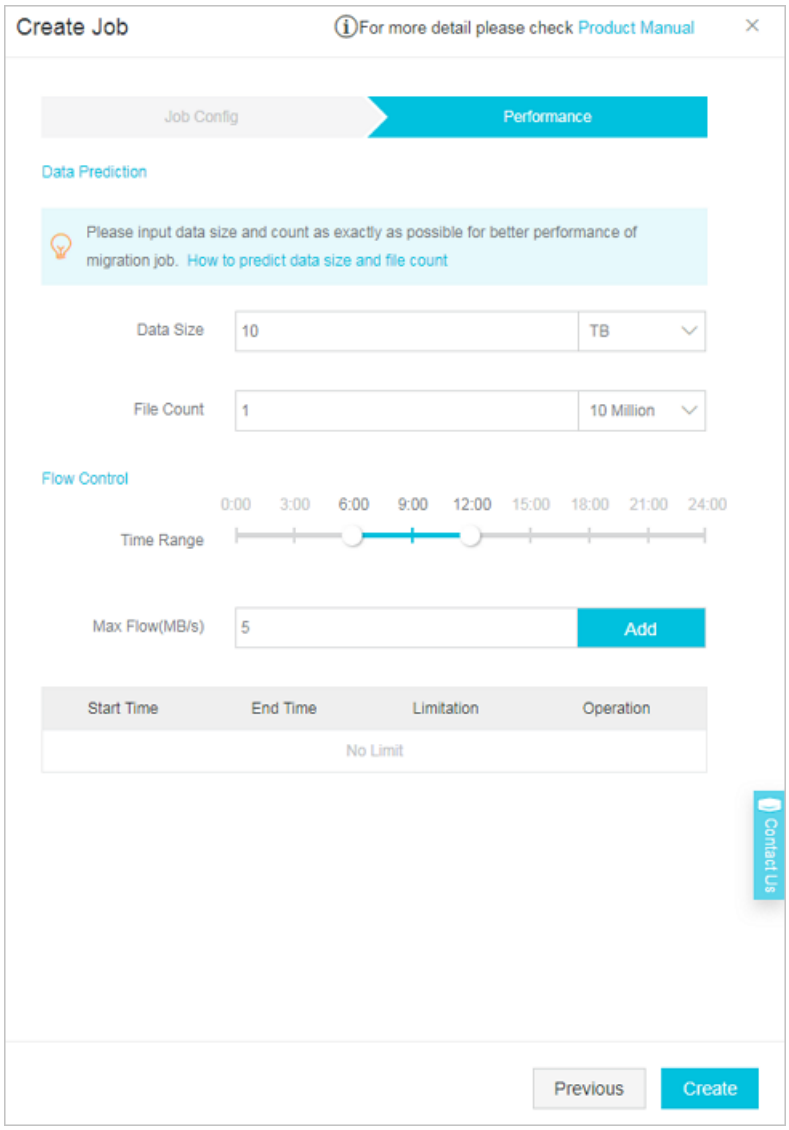

6. A migration job requires about one day to complete. After the migration job is completed, you must verify that all data is migrated. To do so, view the migration report and compare the data at the source data address with the data at the destination data address.

> **Note** If a file fails to be migrated, troubleshoot the failure. For more information, see Common causes of a migration failure and solutions.

5.Migrate data from a local IDC to OSS

This topic describes how to migrate data from a local IDC to OSS.

Context

An e-commerce enterprise uses the Fast Distributed File System (FastDFS) to store data in the user-created IDC. The data includes about 30,000,000 files and has a total size of about 300 TB. The enterprise has connected the local IDC to a VPC in China (Shenzhen) by using Alibaba Cloud Express Connect.

Enterprise A wants to switch over the businesses to OSS due to development needs. To ensure business continuity, the following needs must be met for this business switchover:

- During migration, you must avoid the impacts on normal data access from users.
- After the migration job is complete, you must check data integrity to ensure a seamless switchover of the businesses to OSS.

Migration solution

Based on the background information, you can migrate data as follows:

1. Create an OSS bucket in the China (Shenzhen) region and change the default storage location to the data address of this bucket.
2. Use the built-in FastDFS NGINX module to export all files to be migrated to a list of HTTP URLs, and you can access these URLs in the VPC.

   HTTP URLs are separated by line and each line indicates a file. Separate multiple URLs with line feeds (\n). For more information about the format, see Migrate data from HTTP and HTTPS sources to OSS.
3. Use Data Transport to migrate data from the local IDC to OSS.
4. Switch over your businesses to OSS after the migration is complete.

Step 1: Create a bucket and modify the storage location

1. In the China (Shenzhen) region, create a bucket to store data. For more information, see Create buckets.
2. Configure the bucket policy and only enable access to the bucket from enterprise employees. For more information, see Configure bucket policies to authorize other users to access OSS resources.
3. Inform employees of changing the default storage location to the bucket.

Step 2: Create a migration job

1. Create a RAM user and grant the RAM user the permission to create migration jobs. For more information, see Create and authorize a RAM user.
2. Log on to the Data Transport console as the new RAM user.
3. Create an HTTP source data address. For more information, see Create a source data address.

When creating an HTTP source data address, you must select **Use** for **Whether to Use VPC** and specify the VPC. To ensure a successful migration, you must specify the VPC to access HTTP URLs. The following figure shows the configuration details.

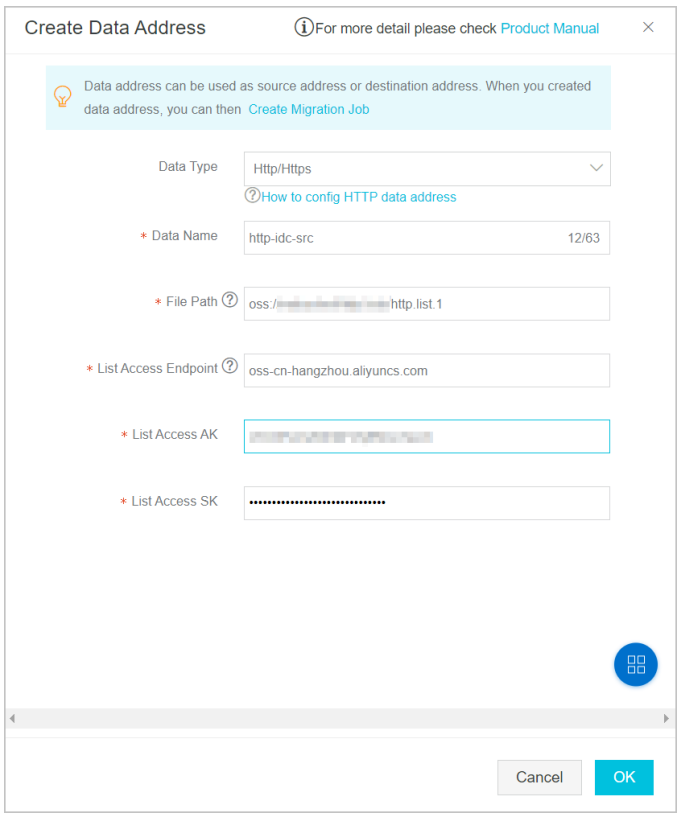

4. Create an OSS destination data address. For more information, see Create a destination data address.

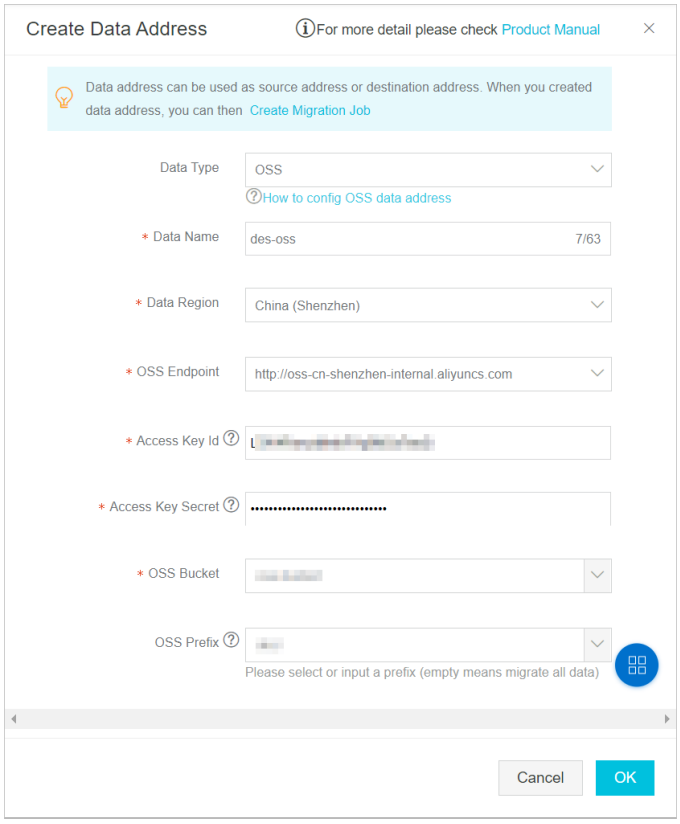

5. Create a full migration job and configure the parameters in the Performance step. For more information, see Create a migration job.

   > **Note** Based on the available bandwidth provided by the enterprise, the migration process requires about two days. In actual practice, you can set an appropriate flow limit based on the usage of the bandwidth.

6. To ensure that all data is migrated after migration, you need to view the migration report and compare data at both the source data address and the destination data address.

   > **Note** For more information about how to troubleshoot migration issues, see Common causes of a migration failure and solutions.

Step 3: Switch over businesses to OSS

After the migration is complete, you can change the data address where the business applications retrieve data to OSS. Then, you can store and manage data on OSS.