Scalable File Service

Best Practices

Issue	01
Date	2025-12-10

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: https://www.huaweicloud.com/intl/en-us/

Contents

1 SFS Best Practices Summary

This section summarizes the SFS best practices and provides a brief description for each practice to help you use SFS easily.

Table 1-1 SFS best practices summary

Practice	Description
Migrating Data Between a General-Purpose File System and an SFS Turbo File System	This practice describes how to migrate data between a general-purpose file system and an SFS Turbo file system.
Migrating Data from SFS Capacity-Oriented to General-Purpose File System Using fpart	This practice describes how to use fpart to migrate data from an SFS Capacity-Oriented file system to a general-purpose file system.
Migrating Data from SFS Capacity-Oriented to SFS Turbo Using fpart	This practice describes how to use fpart to migrate data from an SFS Capacity-Oriented file system to an SFS Turbo file system.
Migrating Data from SFS Capacity-Oriented to Other File Systems Using rclone	This practice describes how to use rclone to migrate data from an SFS Capacity-Oriented file system to a general-purpose or an SFS Turbo file system.
Migrating Data from SFS Capacity-Oriented to SFS Turbo (CIFS Protocol)	This practice describes how to migrate data from a CIFS SFS Capacity-Oriented file system to an SMB SFS Turbo file system.
Configuring DR for SFS General-Purpose File Systems	This practice describes how to configure DR for SFS general-purpose file systems in a given region.

Table 1-2 Best practices for using SFS with other services

Practice	Description	Related Services
Migrating Data from SFS 1.0 to SFS 3.0	This practice describes how to use MgC to migrate data from SFS Capacity-Oriented file systems to SFS general-purpose file systems. The migration is simple, efficient, and secure.	SFS Capacity-Oriented, General-Purpose File System, and MgC
Moving Data Between NAS Systems and Switching Services Seamlessly	This practice describes how to use MgC migration plan to perform a NAS-to-NAS data migration while ensuring service continuity.	General-Purpose File System or SFS Capacity-Oriented, and MgC
Migrating File Systems in Batches	This practice describes how to use MgC to efficiently migrate file systems in batches.	General-Purpose File System or SFS Turbo, and MgC

2 Migrating Data Between a General-Purpose File System and an SFS Turbo File System

Solution Overview

You can migrate data from an SFS general-purpose file system to an SFS Turbo file system or the other way around.

This solution creates a Linux ECS to connect a general-purpose file system with an SFS Turbo file system.

Constraints

- Only Linux ECSs can be used for data migration.
- The Linux ECS, general-purpose file system, and SFS Turbo file system must be in the same VPC.
- Incremental migration is supported, so you can only migrate the changed data.

Prerequisites

- You have created a Linux ECS.
- You have created a general-purpose file system and an SFS Turbo file system and have obtained their addresses.

Resource Planning

Table 2-1 describes the resource planning in this solution.

Table 2-1 Resource planning

Resource	Example Configuration	Description
ECS	Specifications: 8 vCPUs | 16 GB | c7.2xlarge.2 OS: Linux Region: CN-Hong Kong VPC: VPC1	Ensure that the **/mnt/src** and **/mnt/dst** directories have been created.

Procedure

Step 1 **Log in to the ECS console**.

Step 2 Log in to the Linux ECS.

Step 3 Mount file system 1 (either the general-purpose or SFS Turbo file system). After that, you can access file system 1 on the Linux ECS.

```
mount -t nfs -o vers=3,timeo=600,noresvport,nolock <address-of-file-system-1> /mnt/src
```

Step 4 Mount file system 2 (the other file system that is not mounted in the previous step). After that, you can access file system 2 on the Linux ECS.

```
mount -t nfs -o vers=3,timeo=600,noresvport,nolock <address-of-file-system-2> /mnt/dst
```

Step 5 Download and install rclone. For the download address, see **https://rclone.org/downloads/**.

Step 6 Synchronize data.

```
rclone copy /mnt/src /mnt/dst -P --transfers 32 --checkers 64 --links --create-empty-src-dirs
```

NOTE

The following describes the parameters. Set **transfers** and **checkers** based on the system specifications.

- **/mnt/src**: source path
- **/mnt/dst**: destination path
- **--transfers**: number of files that can be transferred concurrently
- **--checkers**: number of local files that can be scanned concurrently
- **-P**: data copy progress
- **--links**: replicates the soft links from the source. They are saved as soft links in the destination.
- **--copy-links**: replicates the content of files to which the soft links point. They are saved as files rather than soft links in the destination.
- **--create-empty-src-dirs**: replicates the empty directories from the source to the destination.

After data synchronization is complete, go to the target file system to check whether data is migrated.

----End

Verification

Step 1 Log in to the Linux ECS.

Step 2 Check the file synchronization results on the destination server.

```
cd /mnt/dst
ls | wc -l
```

Step 3 If the data volume is the same as that on the source server, data is migrated successfully.

----End

3 Migrating Data of SFS Capacity-Oriented File Systems

3.1 Overview

You can use Migration Center (MGC) or fpart to migrate data of SFS Capacity-Oriented file systems. You are advised to use MGC. If MGC fails to meet the requirements in your region, use fpart instead.

- Use MGC to migrate data of SFS Capacity-Oriented file systems to general-purpose file systems or SFS Turbo file systems.

  For details, see **Performing a NAS-to-NAS Migration and Service Cutover**.
- Use mount and fpart to migrate data of SFS Capacity-Oriented file systems to general-purpose file systems or SFS Turbo file systems.

  For details, see the following:
  - **Migrating Data from SFS Capacity-Oriented to General-Purpose File System Using fpart**
  - **Migrating Data from SFS Capacity-Oriented to SFS Turbo Using fpart**

3.2 Migrating Data from SFS Capacity-Oriented to General-Purpose File System Using fpart

Solution Overview

This section describes how to use fpart to migrate data from SFS Capacity-Oriented to General-Purpose File System. The migration aims to efficiently and securely migrate data while ensuring data integrity and consistency. The entire process includes a full migration (no impact on services), a first incremental migration (no impact on services), unmounting the SFS Capacity-Oriented file system (services are interrupted), a second incremental migration (services are interrupted), and changing the mount point (services are interrupted).

Constraints

- This solution is only applicable for the data migration using Linux ECSs.
- This solution is not applicable for SFS Capacity-Oriented file systems using CIFS or those mounted on Windows clients.
- The Linux ECS, SFS Capacity-Oriented file system, and general-purpose file system must be in the same VPC.
- If the general-purpose file system is used by a CCE cluster, the CCE cluster must be upgraded to version 1.19.10, and the everest add-on version must be later than 2.0.9.

Solution Architecture

Figure 3-1 Process of migrating data from SFS Capacity-Oriented to General-Purpose File System using fpart

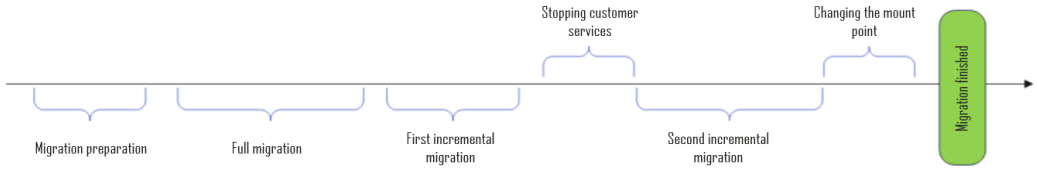

Resource Planning

Table 3-1 describes the resource planning in this solution.

Table 3-1 Resource planning

Resource	Example Configuration	Description
ECS	Specifications: 8 vCPUs | 16 GiB | c7.2xlarge.2 OS: Linux Region: CN-Hong Kong VPC: VPC1	Ensure that the **/mnt/src** and **/mnt/dst** directories have been created.

Preparations

1. Create a general-purpose file system on the Huawei Cloud console in the VPC where the SFS Capacity-Oriented file system resides. If no VPC is available, create one in the same region and then create the general-purpose file system. For how to create a VPC, see **Creating a VPC with a Subnet**.
2. Buy a VPC endpoint (used for free) in the desired VPC.

   The following table shows the available VPC endpoint service names.

Table 3-2 VPC endpoint service names

Region	Service Name
CN North-Beijing4	cn-north-4.com.myhuaweicloud.v4.storage.lz13
CN South-Guangzhou (AZ6)	cn-south-1.com.myhuaweicloud.v4.obsv2.storage.lz06
CN East-Shanghai1	cn-east-3.com.myhuaweicloud.v4.storage.lz07
CN-Hong Kong	ap-southeast-1.com.myhuaweicloud.v4.obsv2.storage.lz005

3. If a Linux ECS is available, skip this step. If not, buy a Linux ECS on the Huawei Cloud console. You are advised to buy an ECS with 8 or 16 vCPUs. This ECS must be in the VPC where the SFS Capacity-Oriented file system and general-purpose file system belong and can communicate with the file systems.
4. Install fpart on the Linux ECS.
   ```
   yum -y install fpart rsync
   ```
5. Create two local mount points on the Linux ECS.
   ```
   mkdir /mnt/src
   mkdir /mnt/dst
   ```
6. Obtain the mount points of the SFS Capacity-Oriented file system and general-purpose file system from the console. Then, mount the file systems on the created local mount points (**/mnt/src** and **/mnt/dst**).
   ```
   mount -t nfs -o vers=3,timeo=600,nolock  <mount-point-of-the-SFS-Capacity-Oriented-file-system> /mnt/src
   mount -t nfs -o vers=3,timeo=600,nolock  <mount-point-of-the-general-purpose-file-system> mnt/dst
   ```
7. If the mount fails, install the required tool package by referring to section "Mounting an NFS File System to ECSs (Linux)" in the *Scalable File Service Getting Started*.

Procedure

Step 1 Full migration: Migrate all files from the SFS Capacity-Oriented file system to the general-purpose file system.

1. Log in to the ECS using VNC and run the following command to synchronize data:
   ```
   fpsync -n 500 -f 500 -o "-lptgoDvu --numeric-ids" -v /mnt/src/ /mnt/dst/
   ```

   NOTE

   - **n**: the number of parallel copy threads
   - **f**: the number of files each thread can copy
   - **o**: the rsync configuration option
   - **v**: the output detailed logs
   - If there are a large number of files, you can run the command concurrently for the subdirectories in **/mnt/src**.
   - You can adjust the **n** and **f** options based on the file and directory structure.

2. The default run log files are stored in **/tmp/fpsync**, under which **log** stores task run logs, **parts** stores file partition logs, **queue** stores execution configurations, and **work** stores the copy commands. A run ID is returned after you run the **fpsync** command. You can use this run ID to check the task running status.

Step 2 First incremental migration: If there is data written to the source files from the service application during the full migration, synchronize the incremental data after the full migration is complete.

Run the following command to perform an incremental migration:

```
fpsync -n 500 -f 500 -o "-lptgoDvu --numeric-ids" -v /mnt/src/ /mnt/dst/
```

CAUTION

Take note of the start time and end time of this operation. This period is approximately equal to how long customer services need to be stopped in the following steps.

Step 3 Unmounting the SFS Capacity-Oriented file system: To prevent new data from being written to the file system, before synchronizing the incremental data, stop the service application that uses the source file system on all ECSs and containers.

CAUTION

- Ensure that the SFS Capacity-Oriented file system is unmounted from all service ECSs (or cloud containers). Otherwise, after the migration, data in the general-purpose file system may be inconsistent with that in the SFS Capacity-Oriented file system.
- Do not unmount the SFS Capacity-Oriented file system from the ECS that you use to perform the migration.
- The operations vary in ECS and CCE scenarios. Refer to the operations based on your scenario.
- Ensure that services are stopped when you perform this step.

ECS scenario:

1. Log in to each service ECS that mounts the SFS Capacity-Oriented file system.
2. Unmount the SFS Capacity-Oriented file system.

```
umount -f <local-mount-point>
```

CCE scenario:

1. Log in to the CCE console, find the workload that uses the storage volume, and reduce the number of pods to 0.

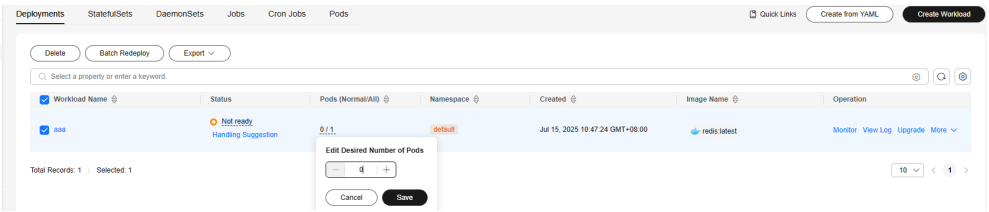

2. Wait until all pods are deleted and then go to the next step.

Step 4 Second incremental migration: Synchronize the incremental data generated after the first incremental migration to the general-purpose file system. After the synchronization, data in the general-purpose file system will be the same as that in the SFS Capacity-Oriented file system.

```
fpsync -n 500 -f 500 -o "-lptgoDvu --numeric-ids" -v /mnt/src/ /mnt/dst/
```

CAUTION

- Check that data in the general-purpose file system is the same as that in the SFS Capacity-Oriented file system.
- Ensure that services are stopped when you perform this step.

Step 5 Mounting the general-purpose file system: Switch the service to the general-purpose file system by mounting all service ECSs (or cloud containers) to the general-purpose file system.

ECS scenario:

1. Log in to the service ECS and run the following command to unmount the SFS Capacity-Oriented file system:
   ```
   umount -f <local-mount-point>
   ```
2. Mount the general-purpose file system.
   ```
   mount -t nfs -o vers=3,timeo=600,noresvport,nolock <mount-point-of-the-general-purpose-file-system>  <local-mount-point>
   ```

CCE scenario: Mount the file system by referring to **Migrating Containerized Application Data from SFS 1.0 to General-Purpose File System or SFS Turbo**.

CAUTION

- Ensure that the mount parameters and directories are correct. Or, services may be affected.
- The operations vary in ECS and CCE scenarios. Refer to the operations based on your scenario.

Step 6 The customer restores services and checks whether services are normal.

Step 7 After observing the services for a period of time, the customer can release the SFS Capacity-Oriented file system resources to prevent unintended charges.

----End

3.3 Migrating Data from SFS Capacity-Oriented to SFS Turbo Using fpart

Solution Overview

This section describes how to use fpart to migrate data from SFS Capacity-Oriented to SFS Turbo. The migration aims to efficiently and securely migrate data

while ensuring data integrity and consistency. The entire process includes a full migration (no impact on services), a first incremental migration (no impact on services), unmounting the SFS Capacity-Oriented file system (services are interrupted), a second incremental migration (services are interrupted), and changing the mount point (services are interrupted).

Constraints

- This solution is only applicable for the data migration using Linux ECSs.
- The Linux ECS, SFS Capacity-Oriented file system, and SFS Turbo file system must be in the same VPC.
- This solution is not applicable for SFS Capacity-Oriented file systems using CIFS or those mounted on Windows clients.
- If the SFS Turbo file system is used by a CCE cluster, the CCE cluster must be upgraded to version 1.15+, and the everest add-on version must be later than 1.1.13.

Solution Architecture

Figure 3-2 Process of migrating data from SFS Capacity-Oriented to SFS Turbo using fpart

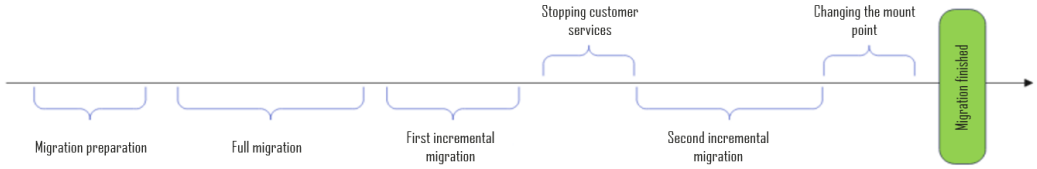

Resource Planning

Table 3-3 describes the resource planning in this solution.

Table 3-3 Resource planning

Resource	Example Configuration	Description
ECS	Specifications: 8 vCPUs | 16 GiB | c7.2xlarge.2 OS: Linux Region: CN-Hong Kong VPC: VPC1	Ensure that the **/mnt/src** and **/mnt/dst** directories have been created.

Preparations

1. Create an SFS Turbo file system on the Huawei Cloud console in the VPC where the SFS Capacity-Oriented file system resides. If no VPC is available, create one in the same region and then create the SFS Turbo file system. For how to create an SFS Turbo file system, see **Creating an SFS Turbo File**

System. For details about how to create a VPC, see **Creating a VPC with a Subnet**.

2. If a Linux ECS is available, skip this step. If not, buy a Linux ECS on the Huawei Cloud console. You are advised to buy an ECS with 8 or 16 vCPUs. This ECS must be in the VPC where the SFS Capacity-Oriented and SFS Turbo file systems belong and can communicate with the file systems.
3. Install fpart on the Linux ECS.
   ```
   yum -y install fpart rsync
   ```
4. Create two local mount points on the Linux ECS.
   ```
   mkdir /mnt/src
   mkdir /mnt/dst
   ```
5. Obtain the mount points of the SFS Capacity-Oriented and SFS Turbo file systems from the console. Then, mount the file systems on the created local mount points (**/mnt/src** and **/mnt/dst**).
   ```
   mount -t nfs -o vers=3,timeo=600,nolock <mount-point-of-the-SFS-Capacity-Oriented-file-system> /mnt/src
   mount -t nfs -o vers=3,timeo=600,nolock <mount-point-of-the-SFS-Turbo-file-system> /mnt/dst
   ```
6. If the mount fails, install the required tool package by referring to section "Mounting an NFS File System to ECSs (Linux)" in the *Scalable File Service Getting Started*.

Procedure

Step 1 Full migration: Migrate all files from the SFS Capacity-Oriented file system to the SFS Turbo file system.

1. Log in to the ECS using VNC and run the following command to synchronize data:
   ```
   fpsync -n 500 -f 500 -o "-lptgoDvu --numeric-ids" -v /mnt/src/ /mnt/dst/
   ```

   NOTE

   - **n**: the number of parallel copy threads
   - **f**: the number of files each thread can copy
   - **o**: the rsync configuration option
   - **v**: the output detailed logs
   - If there are a large number of files, you can run the command concurrently for the subdirectories in **/mnt/src**.
   - You can adjust the **n** and **f** options based on the file and directory structure.
2. The default run log files are stored in **/tmp/fpsync**, under which **log** stores task run logs, **parts** stores file partition logs, **queue** stores execution configurations, and **work** stores the copy commands. You can check the task running status based on these logs.

Step 2 First incremental migration: Synchronize the incremental data generated after the full migration to the SFS Turbo file system.

```
fpsync -n 500 -f 500 -o "-lptgoDvu --numeric-ids" -v /mnt/src/ /mnt/dst/
```

NOTICE

Take note of the start time and end time of this operation. This period is approximately equal to how long customer services need to be stopped in the following steps.

Step 3 Unmounting the SFS Capacity-Oriented file system: To prevent new data from being written to the file system, before synchronizing the incremental data, stop the service application that uses the source file system on all ECSs and containers.

CAUTION

- Ensure that the SFS Capacity-Oriented file system is unmounted from all service ECSs (or cloud containers). Otherwise, after the migration, data in the SFS Turbo file system may be inconsistent with that in the SFS Capacity-Oriented file system.
- Do not unmount the SFS Capacity-Oriented file system from the ECS that you use to perform the migration.
- The operations vary in ECS and CCE scenarios. Refer to the operations based on your scenario.
- Ensure that services are stopped when you perform this step.

ECS scenario:

1. Log in to each service ECS that mounts the SFS Capacity-Oriented file system.
2. Unmount the SFS Capacity-Oriented file system.
   ```
   umount -f <local-mount-point>
   ```

CCE scenario:

1. Log in to the CCE console, find the workload that uses the storage volume, and reduce the number of pods to 0.

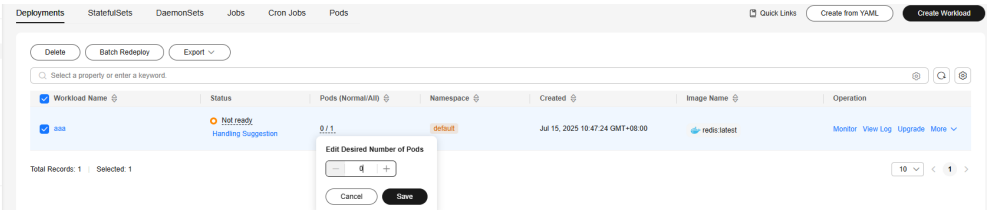

2. Wait until all pods are deleted and then go to the next step.

Step 4 Second incremental migration: Synchronize the incremental data generated after the first incremental migration to the SFS Turbo file system. After the synchronization, data in the SFS Turbo file system will be the same as that in the SFS Capacity-Oriented file system.

```
fpsync -n 500 -f 500 -o "-lptgoDvu --numeric-ids" -v /mnt/src/ /mnt/dst/
```

CAUTION

- Check that data in the SFS Turbo file system is the same as that in the SFS Capacity-Oriented file system.
- Ensure that services are stopped when you perform this step.

Step 5 Mounting the SFS Turbo file system: Switch the service to the SFS Turbo file system by mounting all service ECSs (or cloud containers) to the SFS Turbo file system.

ECS scenario:

1. Log in to the service ECS and run the following command to unmount the SFS Capacity-Oriented file system:
   ```
   umount -f <local-mount-point>
   ```
2. Mount the SFS Turbo file system.
   ```
   mount -t nfs -o vers=3,nolock <mount-point-of-the-SFS-Turbo-file-system>
   ```

CCE scenario: Mount the file system by referring to **Migrating Containerized Application Data from SFS 1.0 to General-Purpose File System or SFS Turbo**.

CAUTION

- Ensure that the mount parameters and directories are correct. Or, services may be affected.
- The operations vary in ECS and CCE scenarios. Refer to the operations based on your scenario.

Step 6 The customer restores services and checks whether services are normal.

Step 7 After observing the services for a period of time, the customer can release the SFS Capacity-Oriented file system resources to prevent unintended charges.

----End

3.4 Migrating Data from SFS Capacity-Oriented to Other File Systems Using rclone

Solution Overview

You can migrate data from SFS Capacity-Oriented to General-Purpose File System or SFS Turbo.

In this solution, a Linux ECS is used to connect the SFS Capacity-Oriented file system and the destination file system.

Constraints

- Only Linux ECSs can be used for data migration.
- The Linux ECS, SFS Capacity-Oriented file system, and destination file system must be in the same VPC. If the destination file system is a general-purpose file system, you need to configure a VPC endpoint.
- Incremental migration is supported, so you can only migrate the changed data.

Prerequisites

- You have created a Linux ECS.
- You have created an SFS Capacity-Oriented file system and a destination file system and have obtained their mount points.

Resource Planning

Table 3-4 describes the resource planning in this solution.

Table 3-4 Resource planning

Resource	Example Configuration	Description
ECS	Specifications: 8 vCPUs | 16 GB | c7.2xlarge.2 OS: Linux Region: CN-Hong Kong VPC: VPC1	Ensure that the **/mnt/src** and **/mnt/dst** directories have been created.

Procedure

Step 1 Log in to the ECS console.

Step 2 Log in to the created Linux ECS, which can access both the SFS Capacity-Oriented and the destination file systems.

Step 3 Mount the SFS Capacity-Oriented file system, which is *file system 1* in this example.

```
mount -t nfs -o vers=3,timeo=600,noresvport,nolock <mount-point-of-file-system-1> /mnt/src
```

Step 4 Mount the general-purpose file system or SFS Turbo file system, which is *file system 2* in this example.

```
mount -t nfs -o vers=3,timeo=600,noresvport,nolock <mount-point-of-file-system-2> /mnt/dst
```

Step 5 Install rclone on the Linux ECS.

```
wget https://downloads.rclone.org/v1.53.4/rclone-v1.53.4-linux-amd64.zip --no-check-certificate
unzip rclone-v1.53.4-linux-amd64.zip
chmod 0755 ./rclone-*/rclone
cp ./rclone-*/rclone /usr/bin/
rm -rf ./rclone-*
```

NOTE

rclone does not retain the file permissions or owner group information on the source. Use rsync if you have such requirements.

Step 6 Synchronize data to the destination file system.

```
rclone copy /mnt/src /mnt/dst -P --transfers 32 --checkers 64 --links --create-empty-src-dirs
```


NOTE

The parameters are described as follows. Set **transfers** and **checkers** based on the system specifications.

- **/mnt/src**: source path
- **/mnt/dst**: destination path
- **--transfers**: number of files that can be transferred concurrently
- **--checkers**: number of local files that can be scanned concurrently
- **-P**: data copy progress
- **--links**: replicates the soft links from the source. They are saved as soft links in the destination.
- **--copy-links**: replicates the content of files to which the soft links point. They are saved as files rather than soft links in the destination.
- **--create-empty-src-dirs**: replicates the empty directories from the source to the destination.

After data synchronization is complete, go to the destination file system to check whether data is migrated.

----End

Verification

Step 1 Log in to the Linux ECS.

Step 2 Check the file synchronization results on the destination server.

```
cd /mnt/dst
ls | wc -l
```

If the data volume is the same as that on the source server, data is migrated successfully.

----End

3.5 Migrating Data from SFS Capacity-Oriented to SFS Turbo (CIFS Protocol)

Solution Overview

This section describes how to migrate data from SFS Capacity-Oriented to SFS Turbo. The migration aims to efficiently and securely migrate data while ensuring data integrity and consistency. The entire process includes a full migration (no impact on services), a first incremental migration (no impact on services), unmounting the SFS Capacity-Oriented file system (services are interrupted), a second incremental migration (services are interrupted), and switching the mount service (services are interrupted).

Figure 3-3 Process of migrating data from SFS Capacity-Oriented to SFS Turbo (CIFS)

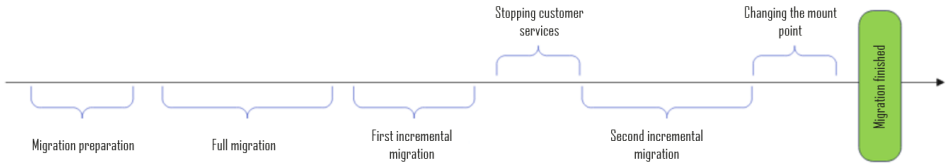

Constraints

- Only Windows ECSs can be used for data migration.
- The Windows ECS, SFS Capacity-Oriented file system, and SFS Turbo file system must be in the same VPC and same security group.
- The SFS Capacity-Oriented file system uses the CIFS protocol, and the SFS Turbo file system uses the SMB protocol.

Preparations

1. Create an SMB SFS Turbo file system in the same VPC and security group as the SFS Capacity-Oriented file system.
   - For how to create an SFS Turbo file system, see **Creating an SFS Turbo File System**.
   - If no VPC is available, create one in the same region and then create the SFS Turbo file system.
   - If your account does not have the security group permissions, when creating the SFS Turbo file system, select **Security Group** under **Advanced Settings** and deselect **Allow traffic on certain ports** on the displayed page (this way, the file system can only be mounted to ECSs in the same security group).
2. Prepare a Windows ECS. The ECS must be in the same VPC and security group as the SFS Capacity-Oriented and SFS Turbo file systems, so that it can communicate with both file systems.

   If no ECS is available, buy one by referring to **Purchasing and Using a Windows ECS**.
3. Log in to the ECS and enable the Workstation and TCP/IP NetBIOS Helper services. For details, see **Mounting an SMB File System in Windows**.
4. Mount both the SFS Capacity-Oriented and SFS Turbo file systems on the Windows ECS.
   a. Obtain the mount information.

      Log in to the **SFS console**. Click the name of the created file system. Obtain the mount command of the SFS Capacity-Oriented file system by following "Mount Information" > "Mount in Windows" on the page. Obtain the mount command of the SFS Turbo file system by following "Basic Info" > "IPv4 Mount Command" > "Command in Windows".

   b. Switch in to the **ECS console** and log in to the Windows ECS.

      i. Open Command Prompt and run the following commands to mount the SFS Capacity-Oriented and SFS Turbo file systems using available drive letters (X and Y are used as example drive letters):

net use X: <mount-address-of-the-SFS Capacity-Oriented-file-system>

net use Y: <mount-address-of-the-SFS-Turbo-file-system>

The format of the file system mount address is ||*<file-system-domain-name>*|*<path>*.

Table 3-5 Variables

Variable	Description
<file-system-domain-name>	Obtain the file system domain name from the file system mount address. For details, see **Viewing a File System**.
<path>	The format is **share-***xxxxxxxx*, where *x* is a digit or letter.

ii. Run **net use**. If information similar to the following is returned, the file systems have been mounted:

```
PS C:\Users\Administrator> net use
New connections will be remembered.

Status       Local       Remote                    Network
-------------------------------------------------------------------------------
OK           X:          \\sfs-nas1.xxx.myhuawei      xxx
OK           Y:          \\f8be395d-xxx-8aa8-6aa460   xxx

The command completed successfully.
```

iii. During mounting, if the system returns the message "You cannot access this shared folder because your organization's security policies block unauthenticated guest access. These policies help protect your PC from unsafe or malicious devices on the network", handle it.

Resource Planning

Table 3-6 describes the resource planning in this solution.

Table 3-6 Resource planning

Resource	Example Configuration
ECS	Specifications: general computing | x1.2u.2g | 8 vCPUs | 16 GiB OS: Windows Region: CN-Hong Kong VPC: VPC1

Procedure

Step 1 Log in to the ECS console.

Step 2 Log in to the Windows ECS **using VNC**.

Step 3 Full migration: Migrate all files from the SFS Capacity-Oriented file system to the SFS Turbo file system.

Run the following command to synchronize data using Robocopy:

```
robocopy X:\ Y:\ /e /w:5 /z /mt:32
```

NOTE

- **/e**: copies all subdirectories (including empty directories).
- **/w**: sets the retry interval (in seconds) of an error.
- **/z**: enables resumable transfer.
- **/mt**: specifies the number of concurrent copy threads. The value ranges from **1** to **128**. The default value is **8**.

Step 4 First incremental migration: Synchronize the incremental data generated after the full migration to the SFS Turbo file system.

```
robocopy X:\ Y:\ /e /w:5 /z /mt:32
```

NOTICE

Take note of the start time and end time of this operation. This period is approximately the time that customer services are interrupted.

Step 5 Unmounting the SFS Capacity-Oriented file system: To prevent new data from being written to the file system, before synchronizing the incremental data, stop the service application that uses the source file system on all ECSs and containers.

CAUTION

- Ensure that the SFS Capacity-Oriented file system is unmounted from all service ECSs (or cloud containers). Otherwise, after the migration, data in the SFS Turbo file system may be inconsistent with that in the SFS Capacity-Oriented file system.
- Do not unmount the SFS Capacity-Oriented file system from the ECS that you use to perform the migration.
- Ensure that services are stopped when you perform this step.

1. Log in to each service ECS that mounts the SFS Capacity-Oriented file system.
2. Unmount the SFS Capacity-Oriented file system.

   ```
   net use <drive-letter>: /del
   ```

Step 6 Second incremental migration: Synchronize the incremental data generated after the first incremental migration to the SFS Turbo file system. After the synchronization, data in the SFS Turbo file system will be the same as that in the SFS Capacity-Oriented file system.

```
robocopy X:\ Y:\ /e /w:5 /z /mt:32
```


CAUTION

- Check that data in the SFS Turbo file system is the same as that in the SFS Capacity-Oriented file system.
- Ensure that services are stopped when you perform this step.

Step 7 Check the migration result.

```
robocopy X:\ Y:\ /e /l /ns /njs /njh /ndl /fp /log:reconcile.txt
```

If the files in the two file systems are the same, the **reconcile.txt** log file is empty. Otherwise, the log file records the file list.

NOTE

- **/e**: Lists directories (including empty directories) only.
- **/l**: Records differences only and does not modify or copy files.
- **/ns**: Does not include file sizes in the log.
- **/njs**: Does not include job summaries.
- **/njh**: Does not include job headers.
- **/ndl**: Does not include folders in the log.
- **/fp**: Includes full paths of files in the log.
- **/log:reconcile.txt**: Writes the migration results to the **reconcile.txt** log. If the log already exists, the existing log will be overwritten.

Step 8 Switching the mount service: Mount all service ECSs to the SFS Turbo file system.

1. Log in to the service ECS and run the following command to unmount the SFS Capacity-Oriented file system:
   ```
   net use <drive-letter>: /del
   ```
2. Mount the SFS Turbo file system.
   ```
   net use <drive-letter>: <mount-address-of-the-SFS-Turbo-file-system>
   ```

CAUTION

- Ensure that the mount parameters and directories are correct. Or, services may be affected.
- Because there are differences between metadata storage and SFS Capacity-oriented, after the mount service is switched to SFS Turbo, it is normal that the data size may increase by 8% to 10%.

Step 9 The customer restores services and checks whether services are normal.

Step 10 After observing the services for a period of time, the customer can release the SFS Capacity-Oriented file system resources to prevent unintended charges.

----End

4 Configuring DR for SFS General-Purpose File Systems

Solution Overview

To quickly recover service data (such as production execution files and team collaboration documents) and reduce downtime, or to build basic disaster recovery (DR) capabilities at a low cost and deploy hierarchical protection through remote backup, you are advised to use the SFS intra-AZ backup solution. This solution requires two SFS general-purpose file systems. One is used for daily services, and the other is for DR. You can manually or periodically back up data of the service file system to the DR file system. The entire process includes a full migration (no impact on services), an incremental migration (no impact on services), scheduled migration task configuration (no impact on services), and migration result check (no impact on services).

Constraints

- This solution is only applicable for the data migration of Linux ECSs.
- The Linux ECS and the two SFS general-purpose file systems must be in the same VPC.
- The general-purpose file systems use the NFS protocol.

Preparations

1. Create a VPC endpoint. For details, see **Configuring a VPC Endpoint**.
2. Create an NFS general-purpose file system and select the same VPC as that of the existing general-purpose file system.
   - For how to create a general-purpose file system, see **Creating a General-Purpose File System**.
   - If no VPC is available, create one in the same region and then create the general-purpose file system.
3. Prepare a Linux ECS. The ECS must be in the same VPC as the two general-purpose file systems, so that it can communicate with both file systems.

   If no ECS is available, buy one by referring to **Purchasing and Using a Linux ECS**.

Resource Planning

Table 4-1 Resource planning

Resource	Example Configuration	Description
ECS	Specifications: 8 vCPUs | 16 GiB | c7.2xlarge.2 OS: Linux Region: CN-Hong Kong VPC: VPC1	Ensure that the **/mnt/src** and **/mnt/dst** directories have been created.

Procedure

Step 1 Mount the two general-purpose file systems by referring to **Mounting an NFS File System to a Linux ECS**. Mount the service file system on **/mnt/src** and mount the DR file system on **/mnt/dst**.

Step 2 Install the migration tool.

```
sudo yum install -y rsync
```

Step 3 Full migration: Migrate all files from the service file system to the DR file system.

1. Create a directory in the DR file system for storing the migrated data, for example, **dst/backup**.
   ```
   mkdir /mnt/dst/backup
   ```
2. Migrate data in the service file system to the backup directory in the DR file system.
   ```
   rsync -avP /mnt/src/ /mnt/dst/backup/
   ```

Step 4 Incremental migration: Synchronize the incremental data generated after the full migration to the DR file system.

```
rsync -avP --delete /mnt/src/ /mnt/dst/backup/
```

Step 5 Use crontab to configure scheduled tasks to migrate incremental data.

1. Access crontab.
   ```
   crontab -e
   ```
2. Configure scheduled incremental migration.

   Add the following information to crontab: (**00** indicates the minute, and **02** indicates the hour. In this example, incremental migration will be performed at 02:00 every day. Change the values as required.)

   ```
   00 02 * * * rsync -avP --delete /mnt/src/ /mnt/dst/backup/ > /tmp/last_rsync_result.log 2>&1 &
   ```
3. Run the following command: (If information configured in **Step 5.2** is returned, the configuration is successful.)
   ```
   crontab -l
   ```

Step 6 Verify the migration results.

Run the following command to check the latest incremental migration data:

```
cat /tmp/last_rsync_result.log
```

----End

