Scalable File Service

FAQs

Issue	04
Date	2024-11-15

Huawei Cloud Computing Technologies Co., Ltd.

Contents

1 Concepts

1.1 What Is SFS?

Scalable File Service (SFS) provides scalable, high-performance file storage. With SFS, you can enjoy shared file access spanning multiple ECSs. SFS supports the Network File System (NFS) protocol. You can seamlessly integrate existing applications and tools with the service.

SFS provides an easy-to-use graphical user interface (GUI). On the GUI, users can create and configure file systems, saving effort in deploying, resizing, and optimizing file systems.

In addition, SFS features high reliability and availability. It can be elastically expanded, and it performs better as its capacity grows. The service is suitable for a wide range of scenarios, including media processing, file sharing, content management and web services, big data, and analytic applications.

1.2 What Is SFS Turbo?

SFS Turbo provides high-performance file storage that can be expanded on demand. With SFS Turbo, you can enjoy shared file access spanning multiple ECSs. SFS Turbo supports the NFS protocol (only NFSv3). You can seamlessly integrate existing applications and tools with SFS Turbo.

SFS Turbo provides an easy-to-use graphical user interface (GUI). On the GUI, users can create and configure file systems, saving effort in deploying, resizing, and optimizing file systems.

In addition, SFS Turbo features high reliability and availability. It can be elastically expanded, and it performs better as its capacity grows. The service is suitable for a wide range of scenarios, including enterprise office, high-performance websites, and software development. For details about the file system types, see **File System Types**.

1.3 What Are the Differences Between SFS, OBS, and EVS?

Table 1-1 shows the comparison between SFS, OBS, and EVS.

Table 1-1 Comparison between SFS, OBS, and EVS

Dimension	SFS	OBS	EVS
Concept	SFS provides on-demand high-performance file storage, which can be shared by multiple ECSs. SFS is similar to a remote directory for Windows or Linux OSs.	OBS provides massive, secure, reliable, and cost-effective data storage for users to store data of any type and size.	EVS provides scalable block storage that features high reliability and high performance to meet various service requirements. An EVS disk is similar to a hard disk on a PC.
Data storage logic	Stores files. Data is sorted and displayed in files and folders.	Stores objects. Files can be stored directly to OBS. The files automatically generate corresponding system metadata. You can also customize the metadata if needed.	Stores binary data and cannot directly store files. To store files, you need to format the file system first.
Access method	SFS file systems can be accessed only after being mounted to ECSs or BMSs through NFS or CIFS. A network address must be specified or mapped to a local directory for access.	OBS buckets can be accessed through the Internet or Direct Connect. The bucket address must be specified for access, and transmission protocols HTTP and HTTPS are used.	EVS disks can be used and accessed from applications only after being attached to ECSs or BMSs and initialized.
Application Scenario	Gene sequencing, image rendering, media processing, file sharing, content management, and web services	Big data analysis, static website hosting, online video on demand (VoD), gene sequencing, and intelligent video surveillance	Industrial design, energy exploration, critical clustered applications, enterprise application systems, and development and testing

Dimension	SFS	OBS	EVS
Capacity	PB-scale	EB-scale	TB-scale
Latency	3–10 ms	10 ms	Sub-millisecond level
IOPS/TPS	10,000 for a single file system	Tens of millions	128,000 for a single disk
Bandwidth	GB/s	TB/s	MB/s
Data sharing	Supported	Supported	Supported
Remote access	Supported	Supported	Not supported
Online editing	Supported	Not supported	Supported
Used independently	Supported	Supported	No (EVS disks can only be used after being attached to ECSs.)

2 Specifications

2.1 What Is the Maximum Size of a Single File Supported for a File System?

For SFS Capacity-Oriented file systems, the maximum supported size of a file is 240 TB.

For SFS Turbo file systems, the maximum supported size of a file is 16 TB.

For general purpose file systems, the maximum supported size of a file is 240 TB.

2.2 What Access Protocols Does SFS Support?

SFS Capacity-Oriented supports standard NFSv3 and CIFS protocols. SFS Turbo and General Purpose File System support the standard NFSv3 protocol.

2.3 How Many File Systems Can I Create with One Account?

Each account can create a maximum of 10 SFS Capacity-Oriented file systems and 10 SFS Turbo file systems.

You can create a maximum of 100 general purpose file systems with one account.

- SFS Capacity-Oriented file systems can be created in batches. To create more than 10 SFS Capacity-Oriented file systems, click **Increase quota** on the page for creating a file system.
- Only one SFS Turbo file system can be created at a time. To create more than 10 SFS Turbo file systems, click **Service Tickets** in the upper right corner of the console to submit a service ticket.

2.4 How Many Cloud Servers Can I Mount a File System To?

You can mount an SFS Capacity-Oriented file system to a maximum of 10,000 cloud servers.

You can mount an SFS Turbo file system to a maximum of 3,000 cloud servers.

You can mount a general purpose file system to a maximum of 10,000 cloud servers.

3 Restrictions

3.1 Can I Expand the File System Capacity If I Start to Run Out of Space?

SFS Capacity-Oriented file systems support online capacity resizing. For details, see **File System Resizing**.

SFS Turbo file systems support online capacity expansion. During an expansion, mounting the file system may fail. This is because the connection used for mounting may experience an I/O delay about 30 seconds (max. 3 minutes).

General purpose file systems have no capacity limit and do not support resizing.

3.2 Can I Migrate My File System Data to Another Region?

Cross-region migration of file system data is currently not supported. It is recommended that you select an appropriate region when purchasing a file system. Alternatively, you can copy the data to a local device and transfer it to another region.

If you are using SFS Turbo file systems, you can back up your file system data and replicate the backups to another region using the CBR service. Then, create new SFS Turbo file systems from the backups. This way, your file system data has been migrated to another region.

3.3 Can a File System Be Mounted Across Accounts?

General purpose file systems cannot be mounted across accounts.

Cross-account mounting for SFS Capacity-Oriented file systems is available.

With VPC peering, an SFS Turbo file system can be accessed across accounts. For details about VPC peering connection and usage instructions, see **VPC Peering Connection**.

4 Networks

4.1 Can a File System Be Accessed Across VPCs?

Yes.

- Multi-VPC access can be configured for an SFS Capacity-Oriented or a general purpose file system so that cloud servers in different VPCs can share the same file system, as long as the VPCs that the servers belong to are added as authorized VPCs of the file system or the server IP addresses are added as authorized IP addresses of the VPCs. For details, see **Configuring Multi-VPC Access**.

- SFS Turbo can work with VPC Peering to allow cloud servers in two or more VPCs of the same region to share the same file system as if they are in the same VPC. For details about VPC peering connection, see **VPC Peering Connection**.

4.2 Does SFS Support Cross-Region Mounting?

SFS Capacity-Oriented file systems cannot be mounted across regions. A file system can only be mounted to ECSs in the same region.

General purpose file systems cannot be mounted across regions.

For an SFS Turbo General file system, cross-region access is supported if the file system is mounted by IP address and communication between VPCs is established using Cloud Connect connections.

4.3 Does the Security Group of a VPC Affect the Use of SFS?

A security group is a collection of access control rules for cloud servers that have the same security protection requirements and are mutually trusted in a VPC. After a security group is created, you can create different access rules for the security group to protect the cloud servers that are added to this security group. The default security group rule allows all outgoing data packets. Cloud servers in

a security group can access each other without the need to add rules. The system creates a security group for each cloud account by default. Users can also create custom security groups by themselves.

For an SFS Turbo file system, the system automatically enables the security group ports required by the NFS protocol after the file system is created. This ensures that the SFS Turbo file system can be successfully mounted to your servers. The inbound ports required by the NFS protocol are ports 111, 2049, 2051, 2052, and 20048. If you need to change the enabled ports, go to the VPC console, choose **Access Control** > **Security Groups**, locate the target security group, and change the ports. You are advised to use an independent security group for an SFS Turbo file system to isolate it from service nodes.

For a general purpose file system, you need to manually add the inbound and outbound rules for the security group. For details, see section "Adding a Security Group Rule" in the *Virtual Private Cloud User Guide*. The inbound ports required by the NFS protocol are ports 111, 2049, and 2050.

For an SFS Capacity-Oriented file system, you need to manually add inbound and outbound rules for the security group. For details, see **Adding a Security Group Rule**. For an SFS Capacity-Oriented file system, the inbound ports required by the NFS protocol are ports 111, 2049, 2050, 2051, and 2052. The inbound port required by the DNS server is port 53 and those required by the CIFS protocol are ports 445 and 135.

Example Configuration

- Inbound rule

Direction	Protocol	Port Range	Source IP Address		Description
Inbound	TCP and UDP	111	IP Address	0.0.0.0/0 (All IP addresses are allowed. It can be modified.)	One port corresponds to one access rule. You need to add information to the ports one by one.

- Outbound rule

Direction	Protocol	Port Range	Source IP Address		Description
Outbound	TCP and UDP	111	IP Address	0.0.0.0/0 (All IP addresses are allowed. It can be modified.)	One port corresponds to one access rule. You need to add information to the ports one by one.

📖 **NOTE**

Enter an IP address range using a mask. For example, enter **192.168.1.0/24**, and do not enter **192.168.1.0-192.168.1.255**. If the source IP address is 0.0.0.0/0, all IP addresses are allowed. For more information, see **Security Groups and Security Group Rules**.

The bidirectional access rule must be configured for port 111. The inbound rule can be set to the front-end service IP range of SFS. You can obtain it by running the following command: **ping** *File system domain name or IP address* or **dig** *File system domain name or IP address*.

For ports 2049, 2050, 2051, and 2052, only the outbound rule needs to be added, which is the same as the outbound rule of port 111.

If the NFS protocol is used, add inbound rules for the following ports: 111 (TCP and UDP), 2049 (TCP), 2051 (TCP), 2052 (TCP), 20048 (UDP and TCP). If UDP is not enabled on port 20048, mounting the file system may take a long time. You can use the **-o tcp** option in the **mount** command to avoid this issue.

4.4 What Can I Do If the Data of the File System That Is Mounted to Two Servers Is Not Synchronized?

Symptom

When file system C is mounted to both server A and server B, there is a delay in synchronizing the file to server B after it is uploaded to server A. However, there is no delay when the file is uploaded to server B separately.

Fault Diagnosis

Add **noac, lookupcache=none** to the mount command.

The **noac** option disables file attribute caching and forces write synchronization. By default, an NFS client's file attribute information is cached using the **ac** option to improve performance, and the client checks file attribute information periodically and updates it if there are any changes. Within the cache validity

period, the client does not check whether file attribute information on the server is changed. By default, the value of this option is **ac**. Set it to **noac**.

The **lookupcache** option is related to directory entry caching, and the value can be **all**, **none**, **pos**, or **positive**. With **lookupcache=none**, the client neither trust the positive nor negative lookup results. In this way, lookup caching is disabled.

Solution

Step 1 Unmount the file system if it has been mounted. For details, see **Unmounting a File System**.

Step 2 Prepare for the mount by referring to **Mounting an NFS File System to ECSs**.

Step 3 Run the following command to mount the file system:

mount -t nfs -o vers=3,timeo=600,noac,lookupcache=none,noresvport,nolock *Shared path Local path*

----End

5 Billing

5.1 How Do I Purchase SFS?

SFS uses pay-per-use billing by default, which means that you are billed by the storage capacity you select during purchase and the duration of use. You can also purchase a yearly or monthly package based on how much space you require and for how long you will use the resources. In case of arrears, you need to renew the service within 15 days, or data in your file systems will be cleared.

📖 **NOTE**

> The size of a resource package is irrelevant to the compute throughput of the file system.

Pay-per-use

Step 1 Register an account.

1. Visit the **Huawei Cloud website**.

2. In the upper right corner of the page, click **Register**.

3. Complete the registration as instructed.

Step 2 Top up your account.

1. Log in to the management console.

2. Click **Top Up** and the top-up page is displayed.

3. Top up the account as prompted.

4. After the top-up is complete, close the dialog box and go back to the management console homepage.

Step 3 Use SFS.

1. Choose **Storage** > **Scalable File Service** to go to the SFS console.

2. Click **Create File System**. You do not need to select a billing mode for an SFS Capacity-Oriented file system. Use the file system after it is created.

📖 **NOTE**

- The SFS Capacity-Oriented file system is billed by the used storage capacity and duration of use.
- The SFS Turbo file system is billed by the storage capacity you select during purchase and the duration of use.

----End

Yearly/Monthly Subscription

You can purchase a resource package of SFS Capacity-Oriented or general purpose file systems by clicking **Buy Storage Package**. Resource packages cannot be used across regions. Check your region before purchasing a resource package.

For an SFS Turbo file system, you can purchase a resource package when creating the file system and change the billing mode from pay-per-use to yearly/monthly.

SFS Capacity-Oriented file systems and general purpose file systems can share a resource package.

Procedure for SFS Capacity-Oriented and General Purpose File System

Step 1 Purchase resource packages.

1. In the upper right corner of the SFS console, click **Buy Storage Package**.
2. On the page that is displayed, set the parameters as instructed.

    You can click **Product Pricing Details** to view the detailed prices.
3. Click **Next**.
4. Confirm the order information and click **Submit**.

    If the order information is incorrect, click **Previous** to modify it and then continue purchasing.
5. Pay for the order.
6. Return to the SFS console to use SFS.

    Resource packages cannot be viewed on the SFS console. To view the resource packages, see **How Do I Check the Usage of an SFS Capacity-Oriented Resource Package?**.

    📖 **NOTE**

    An SFS Capacity-Oriented or a General Purpose File System resource package can be used immediately after payment. If the package capacity is greater than the used capacity of the file system, pay-per-use billing stops immediately. The capacity of the resource packages is preferentially used when you use the file systems.

    SFS Capacity-Oriented and General Purpose File System resource packages cannot be expanded, but can be purchased multiple times. For details, see **Can I Purchase SFS Capacity-Oriented Resource Packages When I Still Have Valid Ones in Use?**

    For example, if a user purchased a 1 TB SFS Capacity-Oriented or General Purpose File System resource package for a year and created a 500 GB SFS Capacity-Oriented file system and used all 500 GB, the used capacity of the resource package is 500 GB. For how to view the resource package usage, see **How Do I Check the Usage of an SFS Capacity-Oriented Resource Package?** One month later, the user creates another 600 GB SFS Capacity-Oriented file system (file system B). Now the purchased 1 TB capacity is used up, and the usage exceeding the package quota (76 GB) is billed on a pay-per-use basis. If you do not want to be billed in pay-per-use mode, you can purchase more resource packages. For details, see **Yearly/Monthly Subscription**.

Step 2 Use SFS.

1. Choose **Storage** > **Scalable File Service** to go to the SFS console.

2. Click **Create File System**. The package capacity is automatically used. You do not need to associate file systems to the resource package. If you have multiple file systems, they can share the same resource package.

----End

Procedure for SFS Turbo

Method 1: Purchase a yearly/monthly file system. Create a file system by following the instructions in **Create a File System** and set the billing mode to **Yearly/Monthly**.

Method 2: In the **Operation** column of the pay-per-use file system, click **Change to Yearly/Monthly** to change the billing mode to yearly/monthly.

5.2 How Do I Renew the Service?

If the arrears are not paid in time, your account may be frozen and your order may be canceled.

Procedure

Step 1 Log in to the management console.

Step 2 In the upper right corner of the page, choose **Billing & Costs** > **Renewal**.

Step 3 On the **Renewals** page, click **Renew** in the **Operation** column for the desired resource.

Step 4 Pay for the order as prompted.

  📖 NOTE

- After the payment, the system automatically pays the arrears.
- For more information about renewal, including auto-renewal, exporting the renewal list, and changing subscriptions, see **Renewal Management**.
- For more information about orders, unsubscription, coupons, and consumption details, see the **Billing Center User Guide**.

----End

5.3 How Do I Check Whether the Subscriber Is in Arrears?

You can view the outstanding amount on the **Billing Center** page. If your account is in arrears, the system processes resources and fees in SFS according to the retention period rules. For more information, see **Service Suspension and Resource Release**. To prevent service suspension and resource release, repay arrears or top up in time.

Procedure

Step 1 Log in to the management console.

Step 2 In the upper right corner of the page, click **Billing & Costs** to go to the Billing Center.

Step 3 On the **Overview** page, you can view the outstanding amount of the current account.

Step 4 Top up your account to pay arrears if any.

----End

📖 NOTE

- For details about how to repay the outstanding amount, see **Making Repayments**.
- For more information about orders, unsubscription, coupons, and consumption details, see the **Billing Center User Guide**.

5.4 Can I Purchase SFS Capacity-Oriented Resource Packages When I Still Have Valid Ones in Use?

You can still purchase resource packages on top of the valid packages. The capacity of the resource package purchased first is used up first.
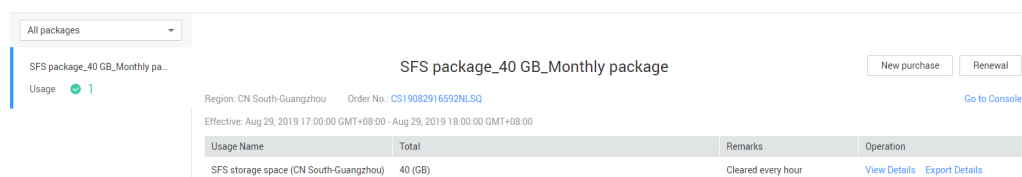
For example: A subscriber has file data of 1.2 TB. The subscriber purchases the resource package A with the capacity of 1 TB on August 15 and the resource package B with the capacity of 1 TB on August 20. Then, from August 20 to September 15, the capacity of the resource package A (1 TB) is used up, and the capacity used in the resource package B is 0.2 TB. The resource package A expires on September 15. If the resource package A is not renewed, the capacity of the resource package B (1 TB) is used up from September 15 to September 20.

5.5 How Do I Check the Usage of an SFS Capacity-Oriented Resource Package?

Procedure

Step 1 Log in to the SFS console.

Step 2 In the upper right corner of the page, choose **Billing & Costs** > **My Packages**.

Step 3 Information of the purchased resource packages is displayed in the list. Click the target resource package. Then click **View Details** in the **Operation** column to view the usage of the resource package. See **Figure 5-1**.

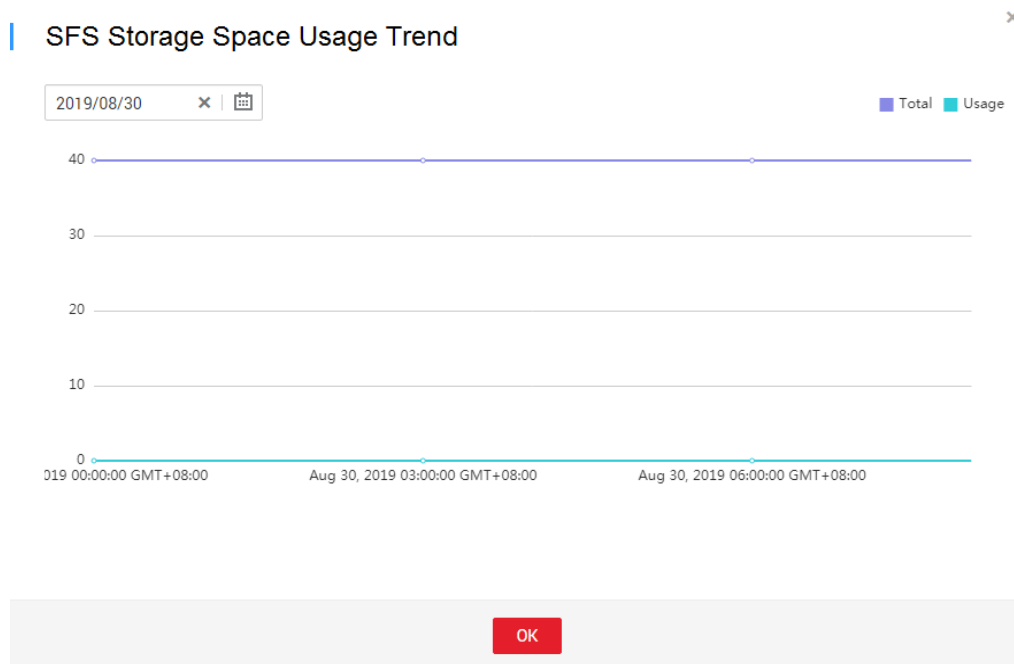**Figure 5-1** Viewing resource package usage

Step 4 View the resource package usage. See **Figure 5-2**.

Figure 5-2 Usage of a resource package

----**End**

5.6 How Do I Adjust the Size of an SFS Capacity-Oriented Resource Package?

If your package capacity is not enough, purchase a new package. The capacity in the old package is preferentially used.

The capacity of a purchased resource package cannot be decreased, so if your package capacity is too much for you, unsubscribe from the package and then purchase a smaller one.

Alternatively, you can use the package until it expires and then purchase a smaller one.

5.7 Do SFS Capacity-Oriented and SFS Turbo Share One Resource Package?

No.

For an SFS Capacity-Oriented file system, you can purchase a resource package separately by clicking **Buy Storage Package**. Resource packages cannot be used across regions. Check your region before purchasing a resource package.

For an SFS Turbo file system, you can purchase a resource package when creating the file system and change the billing mode from pay-per-use to yearly/monthly as needed. Resource packages cannot be purchased separately.

6 Others

6.1 How Do I Access a File System from a Server?

To access a file system, run the **mount** command on a Linux server to mount the file system. Then, you can share the files and directories of the file system.

To access your file system, install the NFS client on a Linux server and run the **mount** command to mount the file system. For a Windows server, install the NFS client, modify the NFS transfer protocol, and run the **mount** command to mount the file system. Alternatively, directly enter the mount point of the CIFS file system as an authorized user to mount the CIFS file system. Then, you can share the files and directories of the file system.

6.2 How Do I Check Whether a File System on a Linux Server Is Available?

- For an SFS Capacity-Oriented file system: Log in to the cloud server as user **root** and run **showmount -e** *File system domain name or IP address*. All available SFS Capacity-Oriented file systems corresponding to the specified domain name or IP address will be returned.
- For a general purpose file system: not supported currently
- For an SFS Turbo file system: Log in to the cloud server as user **root** and run **showmount -e** *File system domain name or IP address*. If the following information is displayed, the SFS Turbo file system corresponding to the file system domain name or IP address in the command is available.
  Export list for *File system domain name or IP address*
  / *

  Or
  Export list for *File system domain name or IP address*
  / IP addresses of all the clients that can access the SFS Turbo file system

6.3 What Resources Does SFS Occupy?

To ensure that file systems can be used properly, the service occupies the following resources:

- For general purpose file systems:

  When a general purpose file system is created, the inbound rules for ports 111, 2049, and 2050 are created in the security group you selected. The source IP address defaults to 0.0.0.0/0 in all rules. You can change the IP address as required.

- For SFS Turbo file systems:

  – When an SFS Turbo file system is created, multiple private IP addresses and virtual IP addresses are created in the subnet entered by the user.

  – When an SFS Turbo file system is created, the inbound rules of ports 111, 445, 2049, 2051, 2052, and 20048 are enabled in the security group entered by the user. The source IP address defaults to 0.0.0.0/0 in all rules. You can change the IP address as required.

When data is written to the folders of a file system, the running memory of the server is occupied, but the storage space of the server disk is not occupied. The file system uses independent space.

6.4 Why Is the Capacity Displayed as 10P After I Mount My SFS Capacity-Oriented File System?

The size of an existing SFS Capacity-Oriented file system with automatic capacity expansion enabled is unlimited. When you run the **df -h** command on the client, the system returns **10P** for display purposes.

6.5 Why the Capacity Is Displayed as 250TB After I Mount My General Purpose File System?

There is no limit on the capacity of a general purpose file system. When you run **df -h** on a client, the system returns **250TB** for display purposes. This value is meaningless, and the available capacity is not limited.

6.6 How Can I Migrate Data Between SFS and OBS?

Background

HUAWEI CLOUD Object Storage Service (OBS) is a stable, secure, efficient, and easy-to-use cloud storage service. With the Representational State Transfer (REST) application programming interfaces (APIs), OBS is able to store any amount and form of unstructured data.

HUAWEI CLOUD Scalable File Service (SFS) is a network attached storage (NAS) service that provides scalable high-performance file storage. With SFS, shared access can be achieved among multiple Elastic Cloud Servers (ECSs), Bare Metal Servers (BMSs), and containers created on Cloud Container Engine (CCE) and Cloud Container Instance (CCI).

How to Migrate

SFS file systems need to be mounted to ECSs, containers, or BMSs. Data migration between SFS and OBS is actually the data migration between servers/containers and OBS. Data in servers or containers is stored in the mounted SFS file systems.

Therefore, after an SFS file system is mounted to a server or container, you can log in to the server or container and use OBS tools, APIs, or SDKs to migrate data. For example, if you want to migrate data from SFS to OBS, use the OBS upload function on the server or container to upload data to OBS. If you want to migrate data from OBS to SFS, use the OBS download function on the server or container to download data to the SFS file system mounted on the server or container. Read and write permissions are required.

As listed in **Table 6-1**, OBS provides various methods for data migration. Select a suitable migration method according to your operating system and actual needs, and migrate data by referring to the upload and download sections in the guide manual.

📖 **NOTE**

> The supported operating systems, data volume, and operation complexity vary according to the migration methods. obsutil is recommended.
>
> To reduce costs, it is recommended that you configure the intranet DNS and migrate data over the HUAWEI CLOUD intranet. For details, see **Accessing OBS over Intranet**.

Table 6-1 Data migration tools provided by OBS

Tool	Supported OS (Refer to each tool guide for supported versions.)	Document
OBS Console	Windows	**Console Operation Guide**
OBS Browser+	Windows	**OBS Browser+ Tool Guide**
obsutil (Recommended)	Windows/Linux	**obsutil Tool Guide**
SDK	All	**SDK Reference**
API	All	**API Reference**

6.7 Can a File System Be Accessed Across Multiple AZs?

1. A single file system can be created only in one AZ, for example, **AZ 1**, but can be mounted to and accessed from any AZ.

2. A file system does not support data redundancy across AZs. If the AZ where a file system resides is unavailable, the file system is unavailable.

6.8 Can I Upgrade an SFS Capacity-Oriented File System to an SFS Turbo File System?

No. If data in the SFS Capacity-Oriented file system is no longer required, delete or unsubscribe from it and then purchase an SFS Turbo file system. If you still need the data in the SFS Capacity-Oriented file system, purchase an SFS Turbo file system and **migrate data to the SFS Turbo file system**. After the data is migrated, delete or unsubscribe from the SFS Capacity-Oriented file system.

6.9 Can I Upgrade an SFS Turbo File System from Standard to Standard-Enhanced?

No. The type of an existing SFS Turbo file system cannot be changed.

If you want an SFS Turbo file system of another type, delete or unsubscribe from the current file system and purchase a new one with your desired type. If you still need the data in your old file system, purchase a new file system and **migrate data to the new file system**. After the data is migrated, delete or unsubscribe from the old file system.

6.10 How Can I Migrate Data Between SFS and EVS?

Mount a file system and an EVS disk to the same ECS, and then manually replicate data between the file system and EVS disk.

6.11 Can I Directly Access SFS from On-premises Devices?

SFS Turbo supports on-premises access through IDC private lines or other methods. After network communication is established, you can access an SFS Turbo file system from your on-premises devices.

SFS Capacity-Oriented file systems currently do not support on-premises access, but you can **migrate your data to SFS** and then use the file system.

6.12 How Do I Delete .nfs Files?

NFS .nfs Files

The .nfs files are temporary files in NFS. If you try to delete a file, and the file is still open, a .nfs file will appear. The .nfs files are used by NFS clients to manage the deletion of open files in the file system. If one process deletes a file while another process still has it open, the client will rename the file to .nfsxxx. If the last open to this file is closed, the client will automatically delete the file. If the client crashes before the file is cleared, the file will be left in the file system.

Clearing .nfs Files

The .nfs files need to be cleared. You can run the **rm -f** command to delete them. The file system will not be affected by the deletion. If an error is reported when you delete a .nfs file, do as follows:

Figure 6-1 Deletion error

```
$ rm -f .nfs0000000001f0df8c0000XXXX

rm: cannot remove `smkit/SM_DOMAIN/.nfs0000000001f0df8c0000XXXX': Device or resource
busy
```

Run the **lsof** command to obtain the ID of the process that has the file open.

Figure 6-2 Viewing the process ID

```
$ lsof .nfs0000000001f0df8c0000XXXX

COMMAND   PID    USER   FD    TYPE DEVICE SIZE/OFF     NODE NAME

java    25887 <UID> mem    REG   0,22   98117 32545366 .nfs0000000001f09a560000XXXX
```

If the process can be stopped, run the **kill -9** *Process ID* command to stop the process and then delete the file.

6.13 Why My File System Used Space Increases After I Migrate from SFS Capacity-Oriented to SFS Turbo?

SFS Turbo file systems contain metadata, which occupies about 8% to 10% file system space. That is why the used space of your file system increases after a data migration from SFS Capacity-Oriented file systems to SFS Turbo file systems. The metadata consists of the file system management data, such as the file size, file system owner, and file modification time.

6.14 How Can I Improve the Copy and Delete Efficiency with an SFS Turbo File System?

Common Linux commands, such as **cp**, **rm**, and **tar**, are executed sequentially. To take the concurrency advantage of cloud file systems, run commands concurrently to improve efficiency.

6.15 How Do Second- and Third-level Directory Permissions of an SFS Turbo File System Be Inherited?

Subdirectories in SFS Turbo file systems cannot inherit permissions of their parent directories.

6.16 How Do I Deploy SFS Turbo on CCE?

Complete the deployment on the CCE console based on your services by referring to **Storage** or **Storage (FlexVolume)**.