Scalable File Service

User Guide

Issue 22

Date 2020-04-01



HUAWEI TECHNOLOGIES CO., LTD.

Copyright © Huawei Technologies Co., Ltd. 2021. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions


and other Huawei trademarks are trademarks 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 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.

Contents

1 Permissions Management 1
1.1 Creating a User and Granting SFS Permissions 1
1.2 Creating a Custom Policy 2
2 File System Management 5
3 Network Configuration 8
3.1 Configuring VPCs 8
3.2 Configuring Cross-Account Access 13
3.3 Configuring DNS 17
4 File System Resizing 21
5 Quotas 24
6 Encryption 26
7 Backup 27
8 Monitoring 29
8.1 SFS Metrics 29
8.2 SFS Turbo Metrics 31
8.3 Creating Alarm Rules 33
9 Auditing 37
10 Typical Applications 39
10.1 HPC 39
10.2 Media Processing 41
10.3 Enterprise Website/App Background 42
10.4 Log Printing 43
11 Other Operations 45
11.1 Testing SFS Turbo Performance 45
11.2 Mounting a File System to an ECS Running Linux as a Non-root User 50
11.3 Configuring DNS for Cache-accelerated SFS Turbo 52
11.4 Data Migration 55
11.4.1 Migrating Data Using Direct Connect 55
11.4.2 Migrating Data Using the Internet 56

 Copyright © Huawei Technologies Co., Ltd.

11.4.3 Migrating Data Between an SFS Capacity-Oriented File System and an SFS Turbo File System.....59

A Change History.....61

 Copyright © Huawei Technologies Co., Ltd.

1 Permissions Management

1.1 Creating a User and Granting SFS Permissions

This chapter describes how to use IAM to implement fine-grained permissions control for your SFS resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing SFS resources.
- Grant only the permissions required for users to perform a specific task.

If your HUAWEI CLOUD account does not require individual IAM users, skip this section.

This section describes the procedure for granting permissions (see **Figure 1-1**).

Prerequisites

Learn about the permissions (see **System-defined roles and policies**) supported by SFS and choose policies or roles according to your requirements. For the permissions of other services, see **System Permissions**.

Restrictions

- All system-defined policies and custom policies are supported in SFS Capacity-Oriented file systems.
- Only system-defined policies are supported in SFS Turbo file systems and custom policies are not supported.

 Copyright © Huawei Technologies Co., Ltd.

Process Flow

Figure 1-1 Process for granting SFS permissions

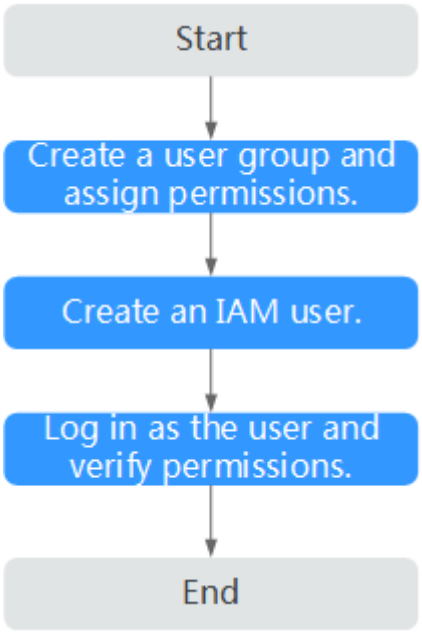

1. **Create a user group and assign permissions** to it.

   Create a user group on the IAM console, and attach the **SFS ReadOnlyAccess** or **SFS Turbo ReadOnlyAccess** policy to the group.

2. **Create a user and add it to a user group**.

   Create a user on the IAM console and add the user to the group created in **1**.

3. **Log in** and verify permissions.

   Log in to SFS Console using the created user, and verify that the user only has read permissions for SFS.

   - Choose **Scalable File Service**. Click **Create File System** on SFS Console. If a message appears indicating that you have insufficient permissions to perform the operation, the **SFS ReadOnlyAccess** or **SFS Turbo ReadOnlyAccess** policy has already taken effect.
   - Choose any other service. If a message appears indicating that you have insufficient permissions to access the service, the **SFS ReadOnlyAccess** or **SFS Turbo ReadOnlyAccess** policy has already taken effect.

1.2 Creating a Custom Policy

Custom policies can be created to supplement the system-defined policies of SFS. For the actions supported for custom policies, see **Permissions Policies and Supported Actions**.

You can create custom policies in either of the following two ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.

Copyright © Huawei Technologies Co., Ltd.

- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see **Creating a Custom Policy**. This section provides examples of common custom SFS policies.

Restrictions

A custom policy applies only to SFS Capacity-Oriented file systems, not SFS Turbo file systems.

Example Custom Policies

- Example 1: Allowing users to create file systems

```
{
    "Version": "1.1",
    "Statement": [
        {
            "Action": [
                "sfs:shares:createShare"
            ],
            "Effect": "Allow"
        }
    ]
}
```

- Example 2: Denying file system deletion

  A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

  The following method can be used if you need to assign permissions of the **SFS FullAccess** policy to a user but also forbid the user from deleting file systems. Create a custom policy for denying file system deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on SFS except deleting file systems. The following is an example of a deny policy:

```
{
    "Version": "1.1",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": [
                "sfs:shares:deleteShare"
            ]
        }
    ]
}
```

- Example 3: Defining permissions for multiple services in a policy

  A custom policy can contain actions of multiple services that are all of the global or project-level type. The following is an example policy containing actions of multiple services:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "sfs:shares:createShare",
        "sfs:shares:deleteShare",
        "sfs:shares:updateShare"
```

 Copyright © Huawei Technologies Co., Ltd.

```
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "ecs:servers:delete"
            ]
        }
    ]
}
```

 Copyright © Huawei Technologies Co., Ltd.

2 File System Management

Viewing a File System

You can search for file systems by file system name keyword or file system status, and view their basic information.

Procedure

Step 1 Log in to SFS Console.

Step 2 In the file system list, view the file systems you have created. **Table 2-1** describes the parameters of each file system.

Table 2-1 Parameter description

Parameter	Description
Name	Name of the file system, for example, **sfs-name-001**
AZ	Availability zone where the file system is located
Status	Possible values are **Available**, **Unavailable**, **Frozen**, **Creating**, **Deleting**, **Deletion error**, **Expanding**, **Expansion error**, **Capacity reducing**, **Capacity reduction error**, and **Capacity reduction failed**.
Type	File system type.
Protocol Type	NFS and CIFS protocols are supported.
Used Capacity (GB)	Used space of the file system for storing data **NOTE** This information is refreshed every 15 minutes.
Maximum Capacity (GB)	Maximum capacity of the file system
Encrypted	Encryption status of the created file system. The value can be **Yes** or **No**.
Enterprise Project	Enterprise project to which the file system belongs

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Mount Address	Mount point of the file system. The format of an NFS file system is *File system domain name*:/*Path* or *File system IP address*:/. The format of a CIFS file system is *File system domain name**Path*. **NOTE** If the mount point is too long to display completely, you can adjust the column width.
Operation	For an SFS Capacity-Oriented file system, operations include resizing, deletion, and monitoring indicator viewing. For an SFS Turbo file system, valid operations include deletion, monitoring indicator viewing, subscription renewal, and unsubscription. **NOTE** After a yearly/monthly SFS Turbo file system is created, you can perform operations such as subscription renewal and unsubscription only one to two minutes later.

Step 3 Click the name of a file system to view detailed information about the file system. See **Figure 2-1**.

Figure 2-1 File system information

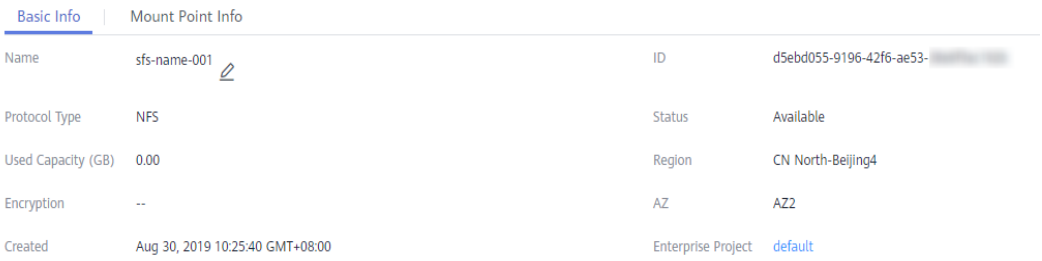

Step 4 (Optional) Search for file systems by file system name keyword, key ID, or file system status.

You can search for file systems in different enterprise projects using the **All Projects** search box in the middle of the page.

----End

Deleting a File System

After a file system is deleted, data in it cannot be restored. To prevent data loss, before deleting a file system, ensure that files in it have been backed up.

Prerequisites

You have unmounted the file system to be deleted. For details about how to unmount the file system, see **Unmounting a File System**.

 Copyright © Huawei Technologies Co., Ltd.

Procedure

Step 1 Log in to SFS Console.

Step 2 In the file system list, click **Delete** or **Unsubscribe** in the row of the file system you want to delete. If you want to delete more than one file system at a time, select the file systems, and then click **Delete** in the upper left part of the file system list. In the dialog box that is displayed, confirm the information, enter **Delete** in the text box, and then click **Yes**. The batch deletion function can be used to delete SFS file systems only.

Step 3 In the displayed dialog box for an SFS Capacity-Oriented file system, as shown in Figure 2-2, confirm the information, enter **Delete** in the text box, and then click **Yes**. After clicking **Unsubscribe** for an SFS Turbo file system, complete the unsubscription as prompted.

NOTE

Only **Available** and **Unavailable** file systems can be deleted or unsubscribed.

Figure 2-2 Deleting a file system

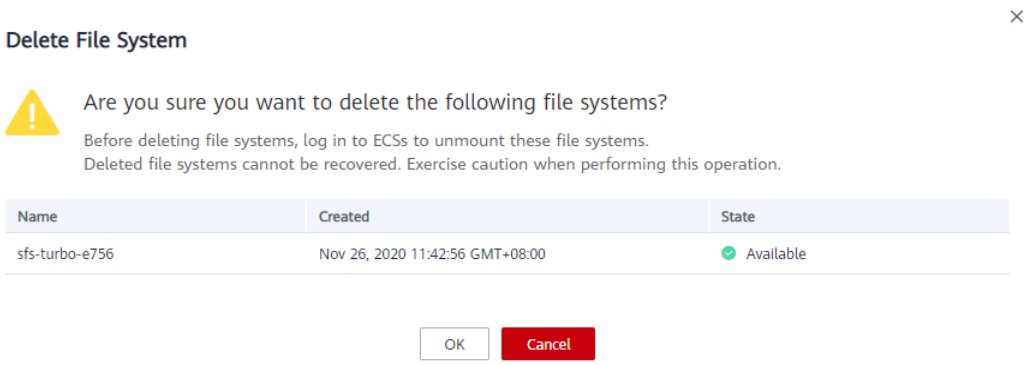

Step 4 Check the file system list to confirm that the file system is deleted successfully.

----End

Copyright © Huawei Technologies Co., Ltd.

3 Network Configuration

3.1 Configuring VPCs

VPC provisions an isolated virtual network environment defined and managed by yourself, improving the security of cloud resources and simplifying network deployment. When using SFS, a file system and the associated ECSs need to belong to the same VPC for file sharing.

In addition, VPC can use network access control lists (ACLs) to implement access control. A network ACL is an access control policy system for one or more subnets. Based on inbound and outbound rules, it determines whether data packets are allowed in or out of any associated subnet. In the VPC list of a file system, each time an authorization address is added and the corresponding permissions are set, a network ACL is created.

For more information about VPC, see the **Virtual Private Cloud**.

Scenarios

Multiple VPCs can be configured for an SFS Capacity-Oriented file system so that ECSs belonging to different VPCs can share the same file system, as long as the VPCs that the ECSs belong to are added to the VPC list of the file system or the ECSs are added to the authorized addresses of the VPCs.

An SFS Turbo file system allows two or more VPCs in the same region to interconnect with each other through VPC peering connection. In this case, different VPCs are in the same network, and ECSs under these VPCs can share the same file system. For details about VPC peering connection, see **VPC Peering Connection**.

This section describes how to access an SFS Capacity-Oriented file system from different VPCs.

Restrictions

- You can add a maximum of 20 VPCs for each file system. A maximum of 400 ACL rules for added VPCs can be created. When adding a VPC, the default IP address 0.0.0.0/0 is automatically added.

 Copyright © Huawei Technologies Co., Ltd.

- If a VPC bound to the file system has been deleted from the VPC console, the IP address/address segment of this VPC in the VPC list of the file system can still be seen as activated. However, this VPC cannot be used any longer and you are advised to delete the VPC from the list.

Procedure

Step 1 Log in to SFS Console.

Step 2 In the file system list, click the name of the target file system. On the page that is displayed, locate the **Authorizations** area.

Step 3 If no VPCs are available, apply for one. You can add multiple VPCs for a file system. Click **Add Authorized VPC** and the **Add Authorized VPC** dialog box is displayed. See **Figure 3-1**.

You can select multiple VPCs from the drop-down list.

Figure 3-1 Adding VPCs

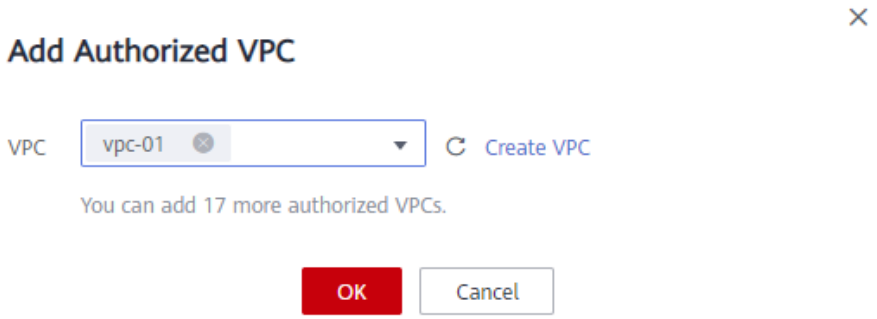

Step 4 Click **OK**. A successfully added VPC is displayed in the list. When adding a VPC, the default IP address **0.0.0.0/0** is automatically added. The default read/write permission is **Read-write**, the default user permission is **no_all_squash**, and the default root permission is **no_root_squash**.

Step 5 View the VPC information in the VPC list. For details about the parameters, see **Table 3-1**.

Table 3-1 Parameter description

Parameter	Description
Name	Name of the added VPC, for example, **vpc-01**
Authorized Addresses/Segments	Number of added IP addresses or IP address segments

Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Operation	The value can be **Add** or **Delete**. **Add**: Adds an authorized VPC. This operation configures the IP address, read/write permission, user permission, user root permission, and priority. For details, see **Table 3-2**. **Delete**: Deletes this VPC.

Step 6 Click ⌄ on the left of the VPC name to view details about the IP addresses/segments added to this VPC. You can add, edit, or delete IP addresses/segments. In the **Operation** column of the target VPC, click **Add**. The **Add Authorized Address/Segment** dialog box is displayed. See **Figure 3-2**. **Table 3-2** describes the parameters to be configured.

Figure 3-2 Adding an authorized address or segment

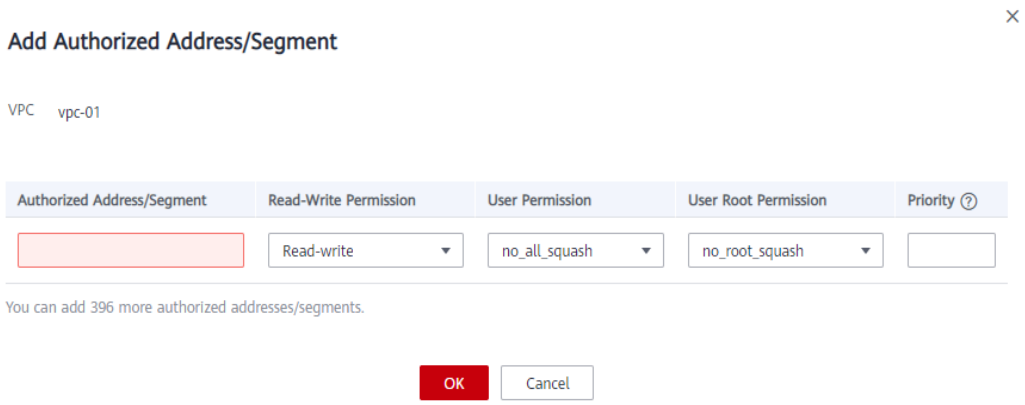

 Copyright © Huawei Technologies Co., Ltd.

Table 3-2 Parameter description

Parameter	Description
Authorized Address/Segment	• Only one IPv4 address or address segment can be entered. • The entered IPv4 address or address segment must be valid and cannot be an IP address or address segment starting with 0 except 0.0.0.0/0. The value **0.0.0.0/0** indicates any IP address in the VPC. In addition, the IP address or address segment cannot start with 127 or any number from 224 to 255, such as 127.0.0.1, 224.0.0.1, or 255.255.255.255. This is because IP addresses or address segments starting with any number from 224 to 239 are class D addresses and they are reserved for multicast. IP addresses or address segments starting with any number from 240 to 255 are class E addresses and they are reserved for research purposes. If an invalid IP address or address segment is used, the access rule may fail to be added or the added access rule cannot take effect. • Multiple addresses separated by commas (,), such as **10.0.1.32,10.5.5.10** are not allowed. • An address segment, for example, 192.168.1.0 to 192.168.1.255, needs to be in the mask format like 192.168.1.0/24. Other formats such as 192.168.1.0-255 are not allowed. The number of bits in a subnet mask must be an integer ranging from 0 to 31. The number of bits **0** is valid only in 0.0.0.0/0.
Read-Write Permission	The value can be **Read-write** or **Read-only**. The default value is **Read-write**.
User Permission	Specifies whether to retain the user identifier (UID) and group identifier (GID) of the shared directory. The default value is **no_all_squash**. • **all_squash**: The UID and GID of a shared directory are mapped to user **nobody**, which is applicable to public directories. • **no_all_squash**: The UID and GID of a shared directory are retained. This parameter is not involved when an authorized address is added for a CIFS file system.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
User Root Permission	Specifies whether to allow the root permission of the client. The default value is **no_root_squash**. ● **root_squash**: Clients cannot access as the **root** user. When a client accesses as the **root** user, the user is mapped to the **nobody** user. ● **no_root_squash**: Clients are allowed to access as the **root** user who has full control and access permissions of the root directories. This parameter is not involved when an authorized address is added for a CIFS file system.
Priority	The value must be an integer ranging from **0** to **100**. **0** indicates the highest priority, and **100** indicates the lowest priority. In the same VPC, the permission of the IP address or address segment with the highest priority is preferentially used. If some IP addresses or address segments are of the same priority, the permission of the most recently added or modified one is used. For example, if the IP address for mounting is 10.1.1.32 and both 10.1.1.32 (read/write) with priority 100 and 10.1.1.0/24 (read-only) with priority **50** meet the requirements, the permission of 10.1.1.0/24 (read-only) with priority **50** is used. That is, if there is no other authorized priority, the permission of all IP addresses in the 10.1.1.0/24 segment, including 10.1.1.32, is read-only.

NOTE

For an ECS in VPC A, its IP address can be added to the authorized IP address list of VPC B, but the file system of VPC B cannot be mounted to this ECS. The VPC used by the ECS and the file system must be the same one.

----End

Verification

After another VPC is configured for the file system, if the file system can be mounted to ECSs in the VPC and the ECSs can access the file system, the configuration is successful.

Example

A user creates an SFS Capacity-Oriented file system A in VPC-B. The network segment is **10.0.0.0/16**. The user has an ECS D in VPC-C, using the private IP address **192.168.10.11** in network segment **192.168.10.0/24**. If the user wants to mount file system A to ECS D and allow the file system to be read and written, the user needs to add VPC-C to file system A's VPC list, add ECS D's private IP address or address segment to the authorized addresses of VPC-C, and then set **Read-Write Permission** to **Read-write**.

 Copyright © Huawei Technologies Co., Ltd.

The user purchases an ECS F that uses the private IP address **192.168.10.22** in the VPC-C network segment **192.168.10.0/24**. If the user wants ECS F to have only the read permission for file system A and its read priority to be lower than that of ECS D, the user needs to add ECS F's private IP address to VPC-C's authorized addresses, set **Read-Write Permission** to **Read-only**, and set **Priority** to an integer between 0 and 100 and greater than the priority set for ECS D.

3.2 Configuring Cross-Account Access

Scenarios

In addition to multi-VPC access, SFS Capacity-Oriented also supports cross-VPC access across accounts.

If VPC IDs used by other accounts are added to the VPC list of a file system and IP addresses or address segments of ECSs are added to the authorized address list, ECSs of different VPCs and of different accounts can share the same file system.

For more information about VPC, see the **Virtual Private Cloud**.

An SFS Turbo file system allows two or more VPCs in the same region to interconnect with each other through VPC peering connection. In this case, different VPCs are in the same network, and ECSs in these VPCs can share the same file system whatever these VPCs belong to the same account or not. For details about VPC peering connection, see **VPC Peering Connection**.

This section describes how to use different accounts to access an SFS Capacity-Oriented file system from different VPCs. Currently, only **CN North-Beijing4** supports cross-account access.

Restrictions

- You can add a maximum of 20 VPCs for each file system. A maximum of 400 ACL rules for added VPCs can be created.
- If a VPC bound to the file system has been deleted from the VPC console, the IP address/address segment of this VPC in the VPC list of the file system can still be seen as activated. However, this VPC cannot be used any longer and you are advised to delete the VPC from the list.

Procedure

Step 1 Log in to SFS Console.

Step 2 In the file system list, click the name of the target file system. On the page that is displayed, locate the **Authorizations** area.

Step 3 Click **Tenant authorized to add VPC** to add VPCs used by other accounts for the file system. A dialog box is displayed. See **Figure 3-3**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-3 Adding a VPC of an authorized tenant

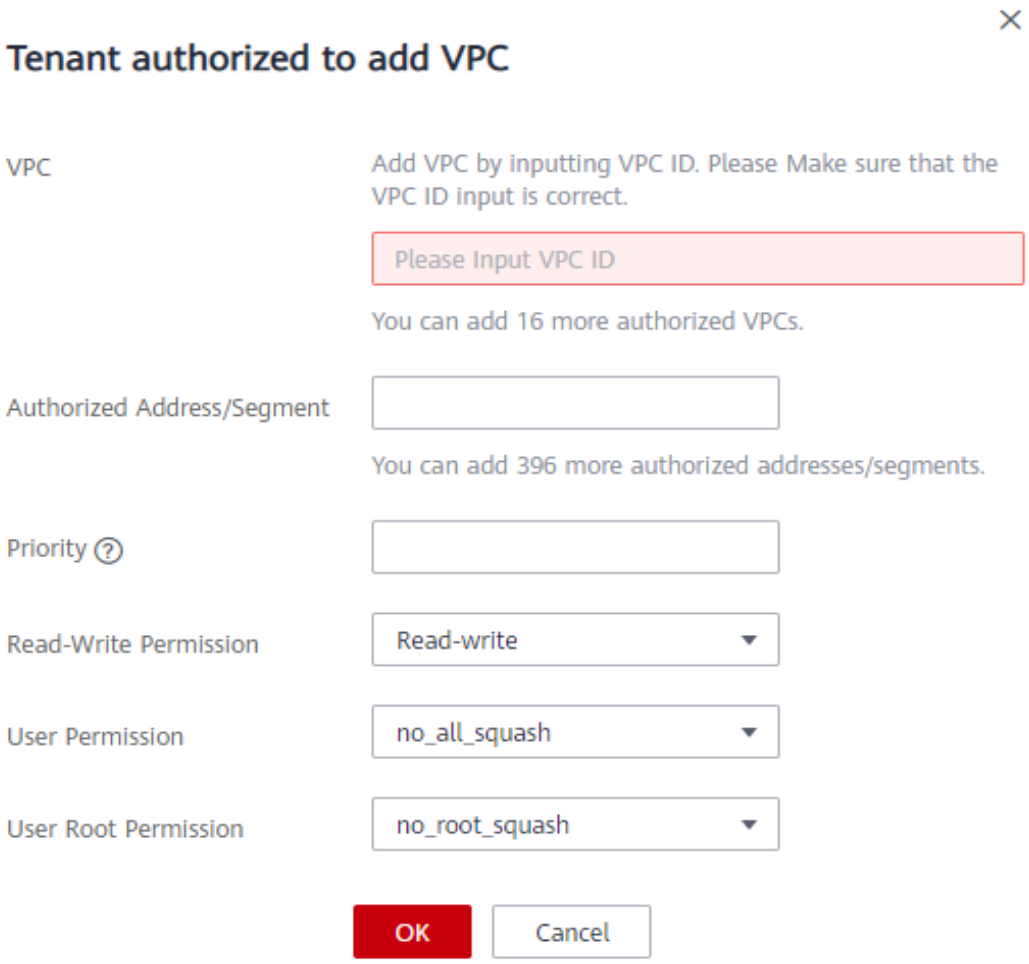

Table 3-3 describes the parameters to be configured.

Table 3-3 Description

Parameter	Description
VPC	Enter the VPC ID of the VPC to be added. You can obtain the VPC ID on the details page of the target VPC on the VPC console.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Authorized Address/Segment	• Only one IPv4 address or address segment can be entered. • The entered IPv4 address or address segment must be valid and cannot be an IP address or address segment starting with 0 except 0.0.0.0/0. The value **0.0.0.0/0** indicates any IP address in the VPC. In addition, the IP address or address segment cannot start with 127 or any number from 224 to 255, such as 127.0.0.1, 224.0.0.1, or 255.255.255.255. This is because IP addresses or address segments starting with any number from 224 to 239 are class D addresses and they are reserved for multicast. IP addresses or address segments starting with any number from 240 to 255 are class E addresses and they are reserved for research purposes. If an invalid IP address or address segment is used, the access rule may fail to be added or the added access rule cannot take effect. • Multiple addresses separated by commas (,), such as **10.0.1.32,10.5.5.10** are not allowed. • An address segment, for example, 192.168.1.0 to 192.168.1.255, needs to be in the mask format like 192.168.1.0/24. Other formats such as 192.168.1.0-255 are not allowed. The number of bits in a subnet mask must be an integer ranging from 0 to 31. The number of bits **0** is valid only in 0.0.0.0/0.
Priority	The value must be an integer ranging from **0** to **100**. **0** indicates the highest priority, and **100** indicates the lowest priority. In the same VPC, the permission of the IP address or address segment with the highest priority is preferentially used. If some IP addresses or address segments are of the same priority, the permission of the most recently added or modified one is used. For example, if the IP address for mounting is 10.1.1.32 and both 10.1.1.32 (read/write) with priority **100** and 10.1.1.0/24 (read-only) with priority **50** meet the requirements, the permission of 10.1.1.0/24 (read-only) with priority **50** is used. That is, if there is no other authorized priority, the permission of all IP addresses in the 10.1.1.0/24 segment, including 10.1.1.32, is read-only.
Read&Write Permissions	The value can be **Read&Write** or **Read-only**. The default value is **Read&Write**.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
User Permission	Specifies whether to retain the user identifier (UID) and group identifier (GID) of the shared directory. The default value is **no_all_squash**. • **all_squash**: The UID and GID of a shared directory are mapped to user **nobody**, which is applicable to public directories. • **no_all_squash**: The UID and GID of a shared directory are retained. This parameter is not involved when an authorized address is added for a CIFS file system.
User Root Permission	Specifies whether to allow the root permission of the client. The default value is **no_root_squash**. • **root_squash**: Clients cannot access as the **root** user. When a client accesses as the **root** user, the user is mapped to the **nobody** user. • **no_root_squash**: Clients are allowed to access as the **root** user who has full control and access permissions of the root directories. This parameter is not involved when an authorized address is added for a CIFS file system.

Step 4 Click **OK**. The added VPC is displayed in the list.

Step 5 Click ⌄ on the left of the VPC name to view details about the IP addresses/segments added to this VPC. You can add, edit, or delete IP addresses/segments. In the **Operation** column of the target VPC, click **Add**. The **Add Authorized Address/Segment** dialog box is displayed. See **Figure 3-4**. **Table 3-3** describes the parameters to be added.

Figure 3-4 Adding an authorized address or segment

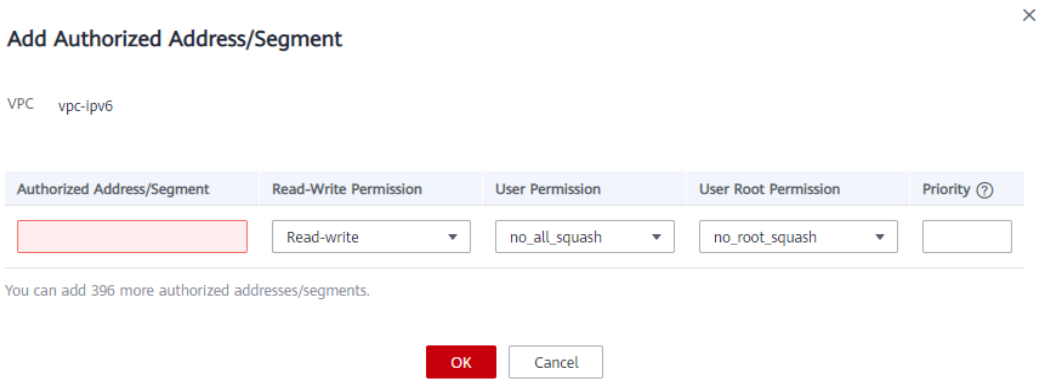

----End

 Copyright © Huawei Technologies Co., Ltd.

Verification

After another user's VPC is configured for the file system, if the file system can be mounted to ECSs in the VPC and the ECSs can access the file system, the configuration is successful.

3.3 Configuring DNS

A DNS server is used to resolve domain names of file systems. The IP address of the DNS server in the CN Northeast region is 100.125.6.250. For IP addresses in other regions, see **What Are the Private DNS Server Addresses Provided by HUAWEI CLOUD DNS?**

Scenarios

By default, the IP address of the DNS server used to resolve domain names of file systems is automatically configured on ECSs when creating ECSs. No manual configuration is needed except when the resolution fails due to a change in the DNS server IP address.

Windows Server 2012 is used as an example in the operation procedures for Windows.

Procedure (Linux)

Step 1 Log in to the ECS as user **root**.

Step 2 Run the **vi /etc/resolv.conf** command to edit the **/etc/resolv.conf** file. Add the DNS server IP address above the existing nameserver information. See **Figure 3-5**.

Figure 3-5 Configuring DNS



The format is as follows:

```
nameserver 100.125.1.250
nameserver 100.125.17.29
```

Step 3 Press **Esc**, input **:wq**, and press **Enter** to save the changes and exit the vi editor.

Step 4 Run the following command to check whether the IP address is successfully added:

cat /etc/resolv.conf

Step 5 Run the following command to check whether an IP address can be resolved from the file system domain name:

nslookup *File system domain name*

NOTE

Obtain the file system domain name from the file system mount point.

Copyright © Huawei Technologies Co., Ltd.

Step 6 (Optional) In a network environment of the DHCP server, edit the **/etc/resolv.conf** file to prevent the file from being automatically modified upon an ECS startup, and prevent the DNS server IP address added in **Step 2** from being reset.

1. Run the following command to lock the file:

   chattr +i /etc/resolv.conf

   NOTE

   Run the **chattr -i /etc/resolv.conf** command to unlock the file if needed.

2. Run the following command to check whether the editing is successful:

   lsattr /etc/resolv.conf

   If the information shown in **Figure 3-6** is displayed, the file is locked.

Figure 3-6 A locked file

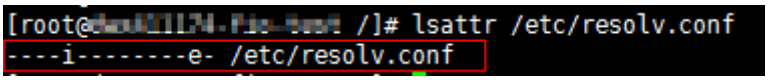

----End

Procedure (Windows)

Step 1 Go to the ECS console and log in to the ECS running Windows Server 2012.

Step 2 Click **This PC** in the lower left corner.

Step 3 On the page that is displayed, right-click **Network** and choose **Properties** from the drop-down list. The **Network and Sharing Center** page is displayed, as shown in **Figure 3-7**. Click **Local Area Connection**.

Figure 3-7 Page for network and sharing center

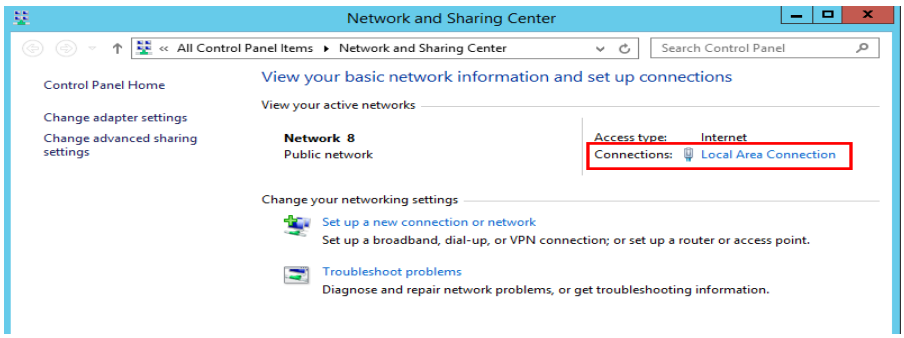

Step 4 In the **Activity** area, select **Properties**. See **Figure 3-8**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-8 Local area connection

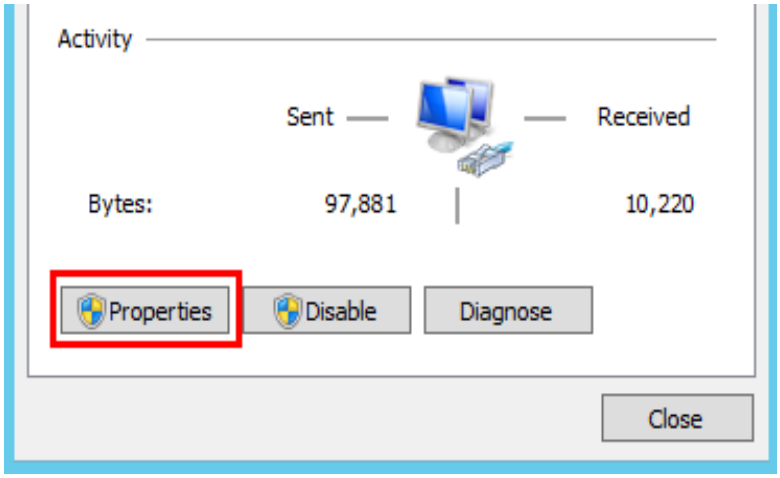

Step 5 In the **Local Area Connection Properties** dialog box that is displayed, select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**. See **Figure 3-9**.

Figure 3-9 Local area connection properties

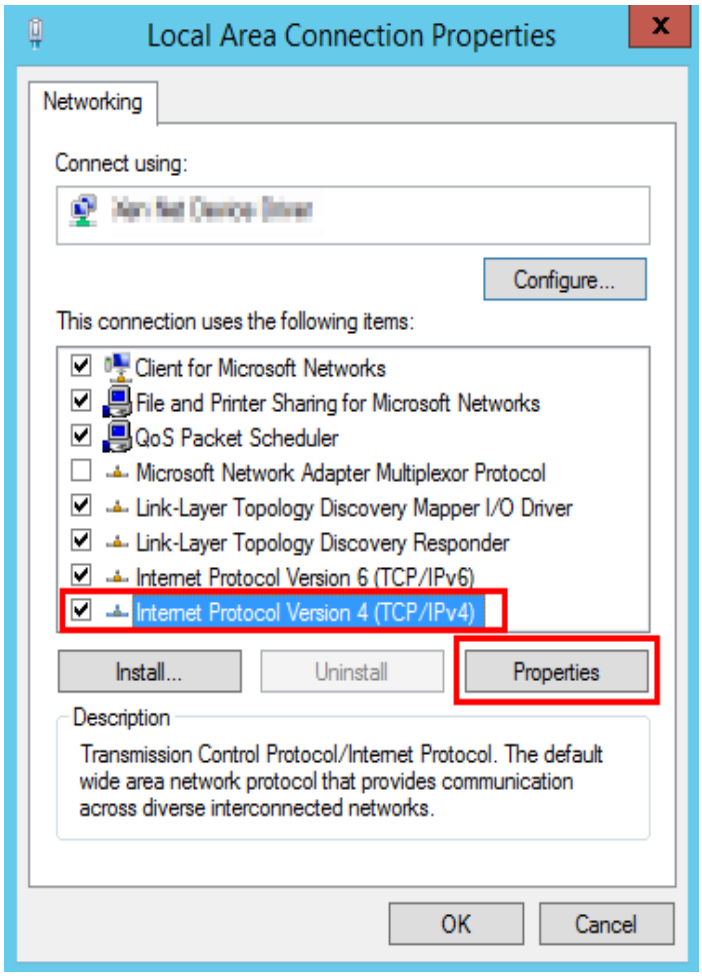

Step 6 In the dialog box that is displayed, select **Use the following DNS server addresses:** and configure DNS, as shown in **Figure 3-10**. The DNS server IP address in Northeast China is 100.125.6.250. For details about DNS server IP

 Copyright © Huawei Technologies Co., Ltd.

addresses in other regions, see **What Are the Private DNS Server Addresses Provided by HUAWEI CLOUD DNS?** After completing the configuration, click **OK**.

Figure 3-10 Configuring DNS on Windows

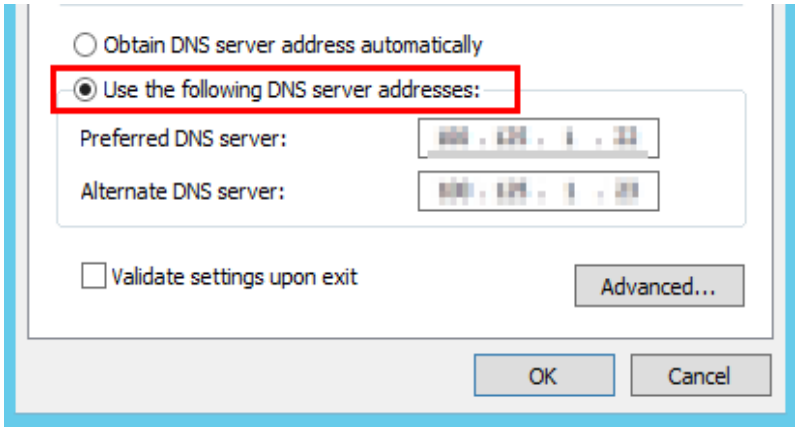

----End

Copyright © Huawei Technologies Co., Ltd.

4 File System Resizing

Scenarios

You can expand or shrink the capacity of a file system when needed.

Limitations and Constraints

SFS Capacity-Oriented file systems support resizing if the auto capacity expansion function is disabled. You can only enable the auto capacity expansion function when creating a file system. After the auto capacity expansion function is enabled, you cannot reset the maximum capacity. In addition, after the auto capacity expansion function is enabled, you cannot disable it.

SFS Capacity-Oriented file systems support online resizing, which does not affect services.

SFS Turbo file systems support online capacity expansion. During capacity expansion, I/Os are suspended for 10 to 15 seconds. You are advised to expand the capacity during off-peak hours. The online capacity expansion button may not be available for some instances of earlier versions. If capacity expansion is required, click **here** to submit a service ticket.

Rules for Resizing

The rules for resizing an SFS Capacity-Oriented file system are as follows:

- Expanding a file system

  Total capacity of a file system after expansion ≤ (Capacity quota of the cloud account - Total capacity of all the other file systems owned by the cloud account)

  For example, cloud account A has a quota of 500 TB. This account has already created three file systems: SFS1 (350 TB), SFS2 (50 TB), and SFS3 (70 TB). If this account needs to expand SFS2, the new capacity of SFS2 cannot be greater than 80 TB. Otherwise, the system will display a message indicating an insufficient quota and the expansion operation will fail.
- Shrinking a file system
  - When a shrink error or failure occurs on a file system, it takes approximately five minutes for the file system to restore to the available state.

 Copyright © Huawei Technologies Co., Ltd.

- After a shrink operation fails, you can only reattempt to shrink the file system storage capacity but cannot expand it directly.
- Total capacity of a file system after shrinking ≥ Used capacity of the file system

  For example, cloud account B has created a file system, SFS1. The total capacity and used capacity of SFS1 are 50 TB and 10 TB respectively. When shrinking SFS1, the user cannot set the new capacity to be smaller than 10 TB.

Procedure

Step 1 Log in to SFS Console.

Step 2 In the file system list, click **Resize** or **Expand Capacity** in the row of the desired file system. The following dialog box is displayed. See **Figure 4-1**.

Figure 4-1 Resizing a file system

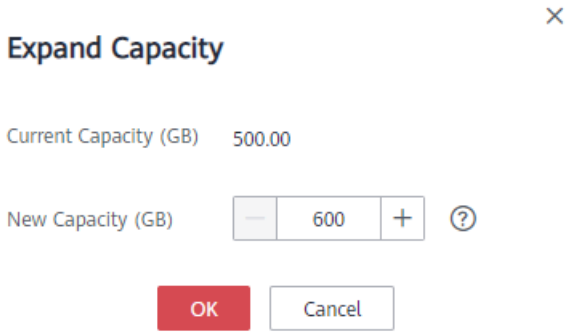

Step 3 Enter a new maximum capacity of the file system based on service requirements, and click **OK**. **Table 4-1** describes the parameters.

Table 4-1 Parameter description

Parameter	Description
Used Capacity (GB)	Used capacity of the current file system
Maximum Capacity (GB)	Maximum capacity of the current file system
New Maximum Capacity (GB)	Target maximum capacity of the file system after expanding or shrinking. The value ranges from **1 GB** to **512,000 GB**. **NOTE** The new maximum capacity cannot be smaller than the used capacity.

Step 4 In the dialog box that is displayed, confirm the information and click **OK**.

 Copyright © Huawei Technologies Co., Ltd.

Step 5 In the file system list, check the capacity information after resizing.

----End

Capacity Expansion on Yearly/Monthly SFS Turbo Subscription

Step 1 Log in to the management console and select **Storage** > **Scalable File Service**.

Step 2 In the file system list, click **Expand** on the row of the SFS Turbo file system to be expanded. The **Expand Capacity** page is displayed. See **Figure 4-2**.

Figure 4-2 Expanding the capacity of an SFS Turbo file system

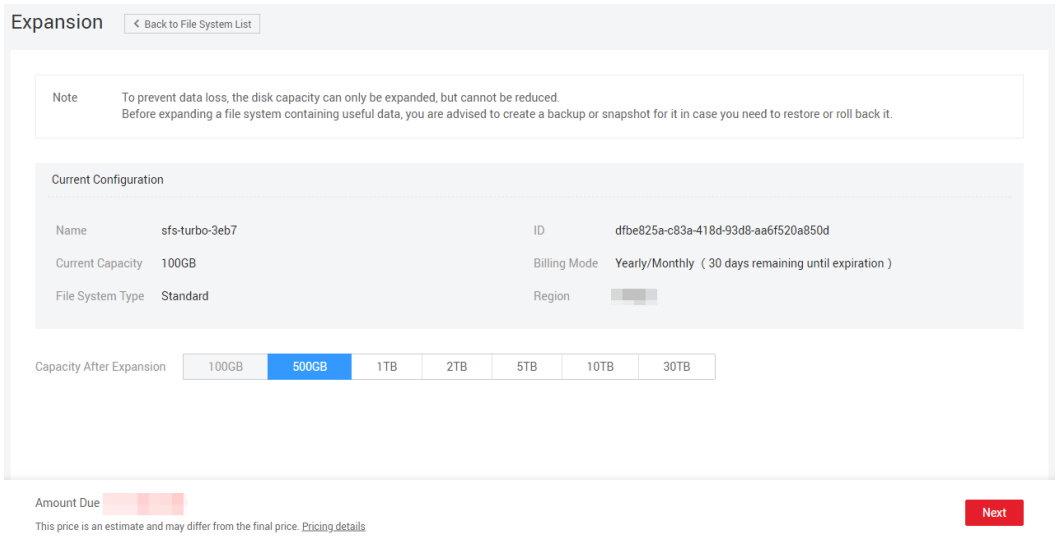

Step 3 Select the capacity after expansion based on service requirements and then click **Next**.

Step 4 Confirm the resource information and click **Submit**.

Step 5 After the payment is successful, you can return to the file system list and click the file system name to check the total capacity after expansion.

----End

Copyright © Huawei Technologies Co., Ltd.

5 Quotas

What Is Quota?

Quotas are enforced for service resources on the platform to prevent unforeseen spikes in resource usage. Quotas can limit the number or amount of resources available to users, such as the maximum number of ECSs or EVS disks that can be created.

If the existing resource quota cannot meet your service requirements, you can apply for a higher quota.

How Do I View My Quotas?

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, choose **Resources** > **My Quotas**.

   The **Service Quota** page is displayed.

   Figure 5-1 My Quotas

   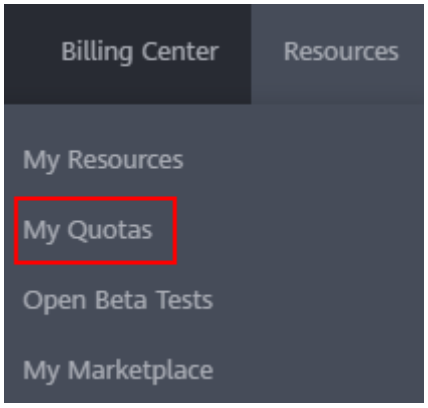

4. View the used and total quota of each type of resources on the displayed page.

   If a quota cannot meet service requirements, apply for a higher quota.

 Copyright © Huawei Technologies Co., Ltd.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources** > **My Quotas**.
   The **Service Quota** page is displayed.

   Figure 5-2 My Quotas

   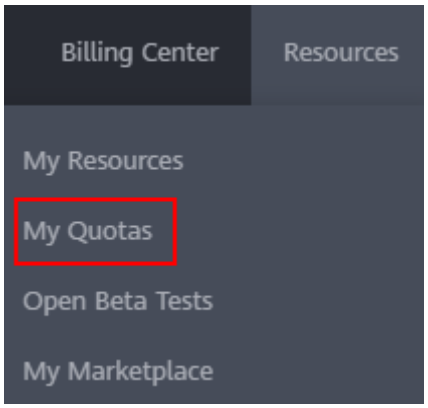

3. Click **Increase Quota**.
4. On the **Create Service Ticket** page, configure parameters as required.
   In **Problem Description** area, fill in the content and reason for adjustment.
5. After all necessary parameters are configured, select **I have read and agree to the Tenant Authorization Letter and Privacy Statement** and click **Submit**.

 Copyright © Huawei Technologies Co., Ltd.

6 Encryption

Creating an Encrypted File System

To use the file system encryption function, you need to authorize SFS Capacity-Oriented to access KMS when creating an SFS Capacity-Oriented file system. If you have the Security Administrator rights, grant SFS the permissions to access KMS directly. Otherwise, you need to contact the system administrator to obtain the "Security Administrator" rights first. For details, see **File System Encryption**.

For SFS Turbo file systems, no authorization is required.

You can create a file system that is encrypted or not, but you cannot change the encryption settings of an existing file system.

For details about how to create an encrypted file system, see **Creating a File System**.

Unmounting an Encrypted File System

If the CMK used by the encrypted file system is disabled or planned to be deleted, the file system can only be used within a certain period of time (30s by default). Exercise caution in this case.

For details about how to unmount the file system, see **Unmounting a File System**.

 Copyright © Huawei Technologies Co., Ltd.

7 Backup

Only SFS Turbo file systems can be backed up using CBR while SFS Capacity-Oriented file systems cannot.

If you need to back up SFS Capacity-Oriented file systems, copy the data to OBS buckets for backup or back up the data locally.

The new backup function is now available. The original backup data will be automatically cleared. Go to the CBR page and back up your data there in a timely manner to avoid data loss.

Scenarios

A backup is a complete copy of an SFS Turbo file system at a specific time and it records all configuration data and service data at that time.

For example, if a file system is faulty or encounters a logical error (for example, mis-deletion, hacker attacks, and virus infection), you can use data backups to restore data quickly.

Procedure

Ensure that the target file system is available. Otherwise, the backup task cannot start. This topic describes how to manually create a file system backup.

Step 1 Log in to SFS Console.

Step 2 In the file system list, locate the SFS Turbo file system to be backed up and click its name. In the **Operation** column, click **Create Backup**.

Step 3 Create a backup vault by following the instructions in **Purchasing a File System Backup Vault** in the *Cloud Backup and Recovery User Guide*. Then, create a backup by following the instructions in **Creating a File System Backup**.

Step 4 The system automatically backs up the file system.

You can view the backup creation status on the **Backup** tab page. When the **Status** of the backup changes to **Available**, the backup has been successfully created.

 Copyright © Huawei Technologies Co., Ltd.

Step 5 If the file system is faulty, use a backup to create a new file system. For details, see **Using a Backup to Create a File System** in the *Cloud Backup and Recovery User Guide*.

----End

 Copyright © Huawei Technologies Co., Ltd.

8 Monitoring

8.1 SFS Metrics

Function

This topic describes metrics reported by Scalable File Service (SFS) as well as their namespaces and dimensions. You can use the console or **APIs** provided by Cloud Eye to query the metrics generated for SFS.

Namespace

SYS.SFS

Metrics

Metric	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Original Metric)
read_bandwidth	Read Bandwidth	Read bandwidth of a file system within a monitoring period Unit: byte/s	≥ 0 bytes/s	SFS file system	4 minutes
write_bandwidth	Write Bandwidth	Write bandwidth of a file system within a monitoring period Unit: byte/s	≥ 0 bytes/s	SFS file system	4 minutes
rw_bandwidth	Read and Write Bandwidth	Read and write bandwidth of a file system within a monitoring period Unit: byte/s	≥ 0 bytes/s	SFS file system	4 minutes

 Copyright © Huawei Technologies Co., Ltd.

Metric	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Original Metric)
read_ops	Read OPS	Read operations of a file system within a monitoring period Unit: count/s	≥ 0 counts/s	SFS file system	4 minutes
write_ops	Write OPS	Write operations of a file system within a monitoring period Unit: count/s	≥ 0 counts/s	SFS file system	4 minutes
rw_ops	Read Write OPS	Read and write operations of a file system within a monitoring period Unit: count/s	≥ 0 counts/s	SFS file system	4 minutes
used_capacity	Used Capacity	Used capacity of a file system within a monitoring period Unit: byte	≥ 0 bytes	SFS file system	4 minutes

Dimension

Key	Value
share_id	SFS file system

Viewing Monitoring Statistics

Step 1 Log in to the management console.

Step 2 View the monitoring graphs using either of the following methods.

- Method 1: Choose **Service List** > **Storage** > **Scalable File Service**. In the file system list, click **View Metric** in the **Operation** column of the target file system.
- Method 2: Choose **Management & Deployment** > **Cloud Eye** > **Cloud Service Monitoring** > **Scalable File Service**. In the file system list, click **View Metric** in the **Operation** column of the target file system.

Step 3 View the SFS file system monitoring data by metric or monitored duration.

Figure 8-1 shows the monitoring graphs. For more information about Cloud Eye, see the *Cloud Eye User Guide*.

 Copyright © Huawei Technologies Co., Ltd.

Figure 8-1 SFS monitoring graphs



----End

8.2 SFS Turbo Metrics

Function

This topic describes metrics reported by Elastic File Service (EFS) to Cloud Eye as well as their namespaces and dimensions. You can use the console or **APIs** provided by Cloud Eye to query the metrics generated for EFS.

Namespace

SYS.EFS

Metrics

Metric	Metric Name	Description	Value Range	Monitored Object
client_connections	Client Connections	Number of client connections	≥ 0	SFS Turbo file system
data_read_io_bytes	Read Bandwidth	Data read I/O load Unit: byte/s	≥ 0 bytes/s	SFS Turbo file system
data_write_io_bytes	Write Bandwidth	Data write I/O load Unit: byte/s	≥ 0 bytes/s	SFS Turbo file system
metadata_io_bytes	Metadata Read and Write Bandwidth	Metadata read and write I/O load Unit: byte/s	≥ 0 bytes/s	SFS Turbo file system
total_io_bytes	Total Bandwidth	Total I/O load Unit: byte/s	≥ 0 bytes/s	SFS Turbo file system
iops	IOPS	I/O operations per unit time	≥ 0	SFS Turbo file system

 Copyright © Huawei Technologies Co., Ltd.

Metric	Metric Name	Description	Value Range	Monitored Object
used_capacity	Used Capacity	Used capacity of a file system Unit: byte	≥ 0 bytes	SFS Turbo file system
used_capacity_percent	Capacity Usage	Percentage of used capacity in the total capacity Unit: Percent	0% to 100%	SFS Turbo file system

Dimension

Key	Value
efs_instance_id	Instance

Viewing Monitoring Statistics

Step 1 Log in to the management console.

Step 2 View the monitoring graphs using either of the following methods.

- Method 1: Choose **Service List** > **Storage** > **Scalable File Service**. In the file system list, click **View Metric** in the **Operation** column of the target file system.
- Method 2: Choose **Management & Deployment** > **Cloud Eye** > **Cloud Service Monitoring** > **Elastic File Service**. In the file system list, click **View Metric** in the **Operation** column of the target file system.

Step 3 View the SFS Turbo file system monitoring data by metric or monitored duration.

Figure 8-2 shows the monitoring graphs. For more information about Cloud Eye, see the *Cloud Eye User Guide*.

Figure 8-2 SFS Turbo monitoring graphs

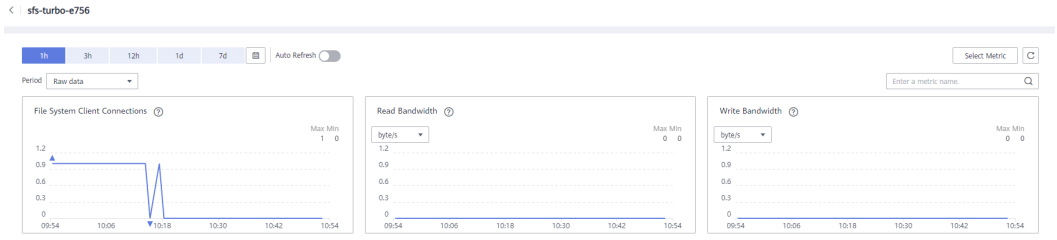

----End

 Copyright © Huawei Technologies Co., Ltd.

8.3 Creating Alarm Rules

The alarm function is based on collected metrics. You can set alarm rules for key metrics of SFS. When the metric data triggers the conditions set in the alarm rule, Cloud Eye sends emails, or text messages, to you, or sends HTTP/HTTPS requests to the servers. In this way, you are immediately informed of cloud service exceptions and can quickly handle the faults to avoid service losses.

Cloud Eye uses Simple Message Notification (SMN) to notify users. This requires you to create a topic and add relevant subscribers for this topic on the SMN console first. Then when you create alarm rules, you can enable the **Alarm Notification** function and select the created topic. When an error occurs, Cloud Eye can broadcast alarm information to those subscribers in real time.

Creating an Alarm Rule

1. Log in to the management console.
2. Choose **Management & Deployment** > **Cloud Eye** > **Cloud Service Monitoring** > **Scalable File Service**, or choose **Management & Deployment** > **Cloud Eye** > **Cloud Service Monitoring** > **Elastic File Service**.
3. Click **Create Alarm Rule** in the **Operation** column of the target file system.
4. On the **Create Alarm Rule** page, set parameters as prompted.
   a. Set **Name** and **Description**.

      Figure 8-3 Specifying the rule name

      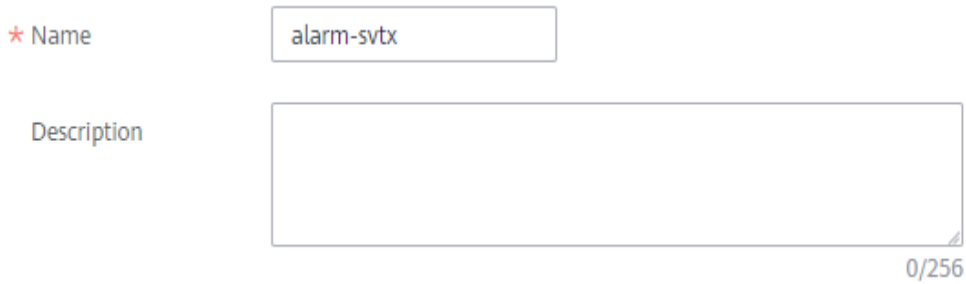

      Table 8-1 Parameters for specifying the rule name

Parameter	Description	Example Value
Name	Name of the alarm rule. The system generates a name randomly but you can change it.	alarm-b6al
Description	Alarm rule description. This parameter is optional.	-

   b. Set **Alarm Content** parameters.

 Copyright © Huawei Technologies Co., Ltd.

Figure 8-4 Configuring the alarm content using a template

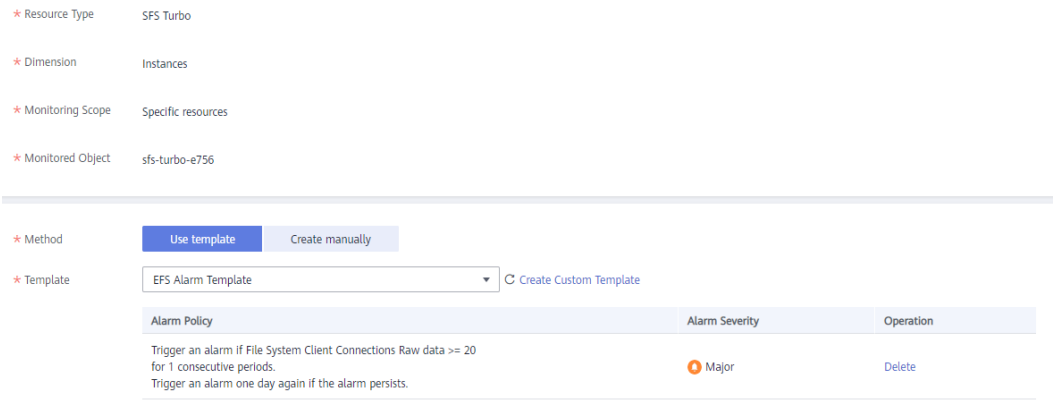

Figure 8-5 Configuring the alarm content manually

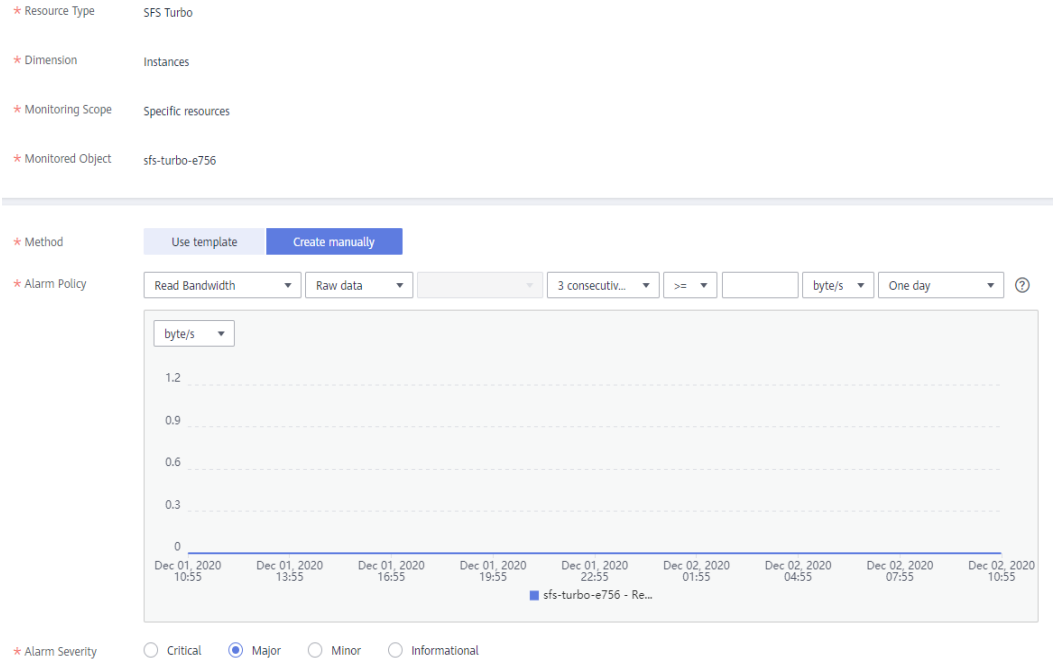

Table 8-2 Parameters for configuring the alarm content

Parameter	Description	Example Value
Method	Select **Use template** or **Create manually** as required. **NOTE** If you set **Monitoring Scope** to **Specific resources**, you can set **Method** to **Use template**.	Use template
Template	Select the template to be imported.	-

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description	Example Value
Alarm Policy	Specifies the policy that triggers an alarm. When **Resource Type** is set to **Website Monitoring**, **Log Monitoring**, **Custom Monitoring**, or a specific cloud service, the alarm policy takes effect periodically. When **Resource Type** is set to **Event Monitoring**, the specific event is an instant operation instead of a periodic action.	-
Alarm Severity	Severity of an alarm. Valid values are **Critical**, **Major**, **Minor**, and **Informational**.	Major

c. Set **Alarm Notification** parameters.

Figure 8-6 Configuring alarm notification

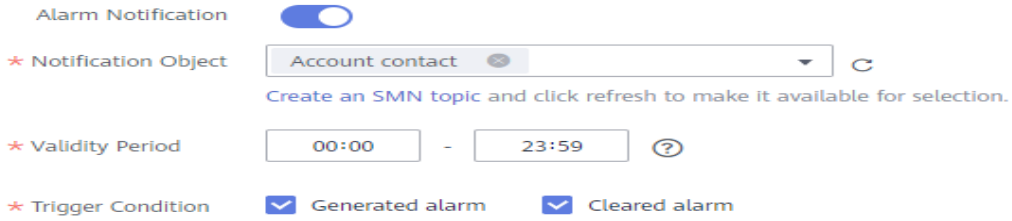

Table 8-3 Parameters for configuring alarm notification

Parameter	Description	Example Value
Alarm Notification	Specifies whether to notify users when alarms are triggered. Notifications can be sent by emails or text messages, or HTTP/HTTPS requests sent to the servers. You can enable (recommended) or disable **Alarm Notification**.	Enable
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule. If **Validity Period** is set to **00:00–8:00**, Cloud Eye sends notifications only within 00:00–8:00.	-

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description	Example Value
Notification Object	Name of the topic to which the alarm notification is sent. If you enable **Alarm Notification**, you need to select a topic. If no desired topics are available, create one first, whereupon the SMN service is invoked. For details about how to create a topic, see the *Simple Message Notification User Guide*.	-
Trigger Condition	Condition for triggering the alarm. You can select **Generated alarm**, **Cleared alarm**, or both.	-

d. Click **Create**.

After the alarm rule is created, if the metric data reaches the specified threshold, Cloud Eye immediately informs you that an exception has occurred. For details about other operations, see the *Cloud Eye User Guide*.

 Copyright © Huawei Technologies Co., Ltd.

9 Auditing

Scenarios

Cloud Trace Service (CTS) records operations of SFS resources, facilitating query, audit, and backtracking.

Currently, only SFS Turbo file systems support recording of resource operations using CTS. SFS Capacity-Oriented file systems do not support this function.

Prerequisites

You have enabled CTS and the tracker is normal. For details about how to enable CTS, see section **Enabling CTS** in the *Cloud Trace Service Getting Started.*

Operations

Table 9-1 SFS operations traced by CTS

Operation	Resource Type	Trace
Creating a shared file system	sfs	createShare
Modifying a shared file system	sfs	updateShareInfo
Deleting a shared file system	sfs	deleteShare
Adding a share access rule	sfs	addShareACL
Deleting a share access rule	sfs	deleteShareACL
Expanding a shared file system	sfs	extendShare

 Copyright © Huawei Technologies Co., Ltd.

Operation	Resource Type	Trace
Shrinking a shared file system	sfs	shrinkShare

Querying Traces

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and project.

Step 3 Choose **Management & Deployment** > **Cloud Trace Service**. The **Cloud Trace Service** page is displayed.

Step 4 In the navigation pane on the left, choose **Trace List**.

Step 5 On the trace list page, set **Trace Source**, **Resource Type**, and **Search By**, and click **Query** to query the specified traces.

For details about other operations, see section "Querying Real-Time Traces" in the *Cloud Trace Service User Guide*.

----End

Disabling or Enabling a Tracker

This section describes how to disable an existing tracker on the CTS console. After the tracker is disabled, the system will stop recording operations, but you can still view existing operation records.

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and project.

Step 3 Choose **Management & Deployment** > **Cloud Trace Service**. The **Cloud Trace Service** page is displayed.

Step 4 Click **Trackers** in the left pane.

Step 5 Click **Disable** on the right of the tracker information.

Step 6 Click **Yes**.

Step 7 After the tracker is disabled, the available operation changes from **Disable** to **Enable**. To enable the tracker again, click **Enable** and then click **Yes**. The system will start recording operations again.

----End

 Copyright © Huawei Technologies Co., Ltd.

10 Typical Applications

10.1 HPC

Context

HPC is short for high-performance computing. An HPC system or environment is made up of a single computer system with many CPUs, or a cluster of multiple computer clusters. It can handle a large amount of data and perform high-performance computing that would be rather difficult for PCs. HPC has ultra-high capability in floating-point computation and can be used for compute-intensive and data-intensive fields, such as industrial design, bioscience, energy exploration, image rendering, and heterogeneous computing. Different scenarios put different requirements on the file system:

- Industrial design: In automobile manufacturing, CAE and CAD simulation software are widely used. When the software is operating, compute nodes need to communicate with each other closely, which requires high bandwidth and low latency of the file system.
- Bioscience: The file system should have high bandwidth and large storage, and be easy to expand.
  - Bioinformatics: To sequence, stitch, and compare genes.
  - Molecular dynamics: To simulate the changes of proteins at molecular and atomic levels.
  - New drug R&D: To complete high-throughput screening (HTS) to shorten the R&D cycle and reduce the investment.
- Energy exploration: Field operations, geologic prospecting, geological data processing and interpretation, and identification of oil and gas reservoirs all require large memory and high bandwidth of the file system.
- Image rendering: Image processing, 3D rendering, and frequent processing of small files require high read/write performance, large capacity, and high bandwidth of file systems.
- Heterogeneous computing: Compute elements may have different instruction set architectures, requiring the file system provide high bandwidth and low latency.

 Copyright © Huawei Technologies Co., Ltd.

SFS is a shared storage service based on file systems. It features high-speed data sharing, dynamic storage tiering, as well as on-demand, smooth, and online resizing. These outstanding features empower SFS to meet the demanding requirements of HPC on storage capacity, throughput, IOPS, and latency.

A biological company needs to perform plenty of gene sequencing using software. However, due to the trivial steps, slow deployment, complex process, and low efficiency, self-built clusters are reluctant to keep abreast of business development. However, things are getting better since the company resorted to professional HPC service process management software. With massive compute and storage resource of the cloud platform, the initial investment and cost during O&M are greatly reduced, the service rollout time is shortened, and efficiency is boosted.

Configuration Process

1. Organize the files of DNA sequencing to be uploaded.
2. Log in to SFS Console. Create a file system to store the files of DNA sequencing.
3. Log in to the ECSs that function as the head node and compute node, and mount the file system.
4. On the head node, upload the files to the file system.
5. On the compute node, edit the files.

Prerequisites

- A VPC has been created.
- ECSs that function as head nodes and compute nodes have been created, and have been assigned to the VPC. For details about how to upload on-premises gene sequencing files to SFS Capacity-Oriented, see **Migrating Data Using Direct Connect**.
- SFS has been enabled.

Example Configuration

Step 1 Log in to SFS Console.

Step 2 In the navigation pane, choose **SFS Capacity-Oriented**. In the upper right corner of the page, click **Create File System**.

Step 3 On the **Create File System** page, set parameters as instructed.

Step 4 After the configuration is complete, click **Create Now**.

For details about how to mount a file system to an ECS running Linux, see **Mounting an NFS File System to ECSs (Linux)**. For details about how to mount a file system running Windows, see **Mounting an NFS File System to ECSs (Windows)** and **Mounting a CIFS File System to ECSs (Windows)**.

Step 5 Log in to the head node, and upload the files to the file system.

Step 6 Start gene sequencing, and the compute node obtains the gene sequencing file from the mounted file system for calculation.

----End

 Copyright © Huawei Technologies Co., Ltd.

10.2 Media Processing

Context

Media processing involves uploading, downloading, cataloging, transcoding, and archiving media materials, as well as storing, invoking, and managing audio and video data. Media processing has the following requirements on shared file systems:

- Media materials feature a high video bit rate and a large scale. The capacity of file systems must be large and easy to be expanded.
- Acquisition, editing, and synthesis of audio and video data require stable and low-latency file systems.
- Concurrent editing requires file systems to deliver reliable and easy-to-use data sharing.
- Video rendering and special effects need processing small files frequently. The file systems must offer high I/O performance.

SFS is a shared storage service based on file systems. It features high-speed data sharing, dynamic storage tiering, as well as on-demand, smooth, and online resizing. These outstanding features empower SFS to meet the demanding requirements of media processing on storage capacity, throughput, IOPS, and latency.

A TV channel has a large volume of audio and video materials to process. The work will be done on multiple editing workstations. The TV channel uses SFS to enable file sharing among the editing workstations. First, a file system is mounted to ECSs that function as upload workstations and editing workstations. Then raw materials are uploaded to the shared file system through the upload workstations. Then, the editing workstations concurrently edit the materials in the shared file system.

Configuration Process

1. Organize the material files that are to be uploaded.
2. Log in to SFS Console. Create a file system to store the material files.
3. Log in to the ECSs that function as upload workstations and editing workstations, and mount the file system.
4. On the upload workstations, upload the material files to the file system.
5. On the editing stations, edit the material files.

Prerequisites

- A VPC has been created.
- ECSs that function as upload workstations and editing workstations have been created, and have been assigned to the VPC. For details about how to upload on-premises material files to SFS Capacity-Oriented, see **Migrating Data Using Direct Connect**.
- SFS has been enabled.

 Copyright © Huawei Technologies Co., Ltd.

Example Configuration

Step 1 Log in to SFS Console.

Step 2 In the navigation pane, choose **SFS Capacity-Oriented**. In the upper right corner of the page, click **Create File System**.

Step 3 On the **Create File System** page, set parameters as instructed.

Step 4 After the configuration is complete, click **Create Now**.

For details about how to mount a file system to an ECS running Linux, see **Mounting an NFS File System to ECSs (Linux)**. For details about how to mount a file system running Windows, see **Mounting an NFS File System to ECSs (Windows)** and **Mounting a CIFS File System to ECSs (Windows)**.

Step 5 Log in to the upload workstations, and upload the material files to the file system.

Step 6 Log in to the editing workstations, and edit the material files.

----End

10.3 Enterprise Website/App Background

Context

For I/O-intensive website services, SFS Turbo can provide shared website source code directories and storage for multiple web servers, enabling low-latency and high-IOPS concurrent share access. Features of such services are as follows:

- A large number of small files: Static website files need to be stored, including HTML files, JSON files, and static images.
- Read I/O intensive: Scope of data reading is large, and data writing is relatively small.
- Multiple web servers access an SFS Turbo background to achieve high availability of website services.

Configuration Process

1. Sort out the website files.
2. Log in to SFS Console. Create an SFS Turbo file system to store the website files.
3. Log in to the ECS that functions as the compute node and mount the file system.
4. On the head node, upload the files to the file system.
5. Start the web server.

Prerequisites

- A VPC has been created.
- ECSs that function as head nodes and compute nodes have been created, and have been assigned to the VPC. For details about how to upload on-premises website files to SFS Turbo, see **Migrating Data Using Direct Connect**.

 Copyright © Huawei Technologies Co., Ltd.

- SFS has been enabled.

Example Configuration

Step 1 Log in to SFS Console.

Step 2 In the navigation pane, choose **SFS Capacity-Oriented**. In the upper right corner of the page, click **Create File System**.

Step 3 On the **Create File System** page, set parameters as instructed.

Step 4 After the configuration is complete, click **Create Now**.

For details about how to mount a file system to an ECS running Linux, see **Mounting an NFS File System to ECSs (Linux)**. For details about how to mount a file system running Windows, see **Mounting an NFS File System to ECSs (Windows)** and **Mounting a CIFS File System to ECSs (Windows)**.

Step 5 Log in to the head node, and upload the files to the file system.

Step 6 Start the web server.

----End

10.4 Log Printing

Context

SFS Turbo can provide multiple service nodes for shared log output directories, facilitating log collection and management of distributed applications. Features of such services are as follows:

- A shared file system is mounted to multiple service hosts and logs are printed concurrently.
- Large file size and small I/O: The size of a single log file is large, but the I/O of each log writing is small.
- Write I/O intensive: Write I/O of small blocks is the major service.

Configuration Process

1. Log in to SFS Console. Create an SFS Turbo file system to store the log files.
2. Log in to the ECS that functions as the compute node and mount the file system.
3. Configure the log directory to the shared file system. It is recommended that each host use different log files.
4. Start applications.

Prerequisites

- A VPC has been created.
- ECSs that function as head nodes and compute nodes have been created, and have been assigned to the VPC. For details about how to upload on-premises log files to SFS Turbo, see **Migrating Data Using Direct Connect**.

 Copyright © Huawei Technologies Co., Ltd.

- SFS has been enabled.

Example Configuration

Step 1 Log in to SFS Console.

Step 2 In the navigation pane, choose **SFS Capacity-Oriented**. In the upper right corner of the page, click **Create File System**.

Step 3 On the **Create File System** page, set parameters as instructed.

Step 4 After the configuration is complete, click **Create Now**.

For details about how to mount a file system to an ECS running Linux, see **Mounting an NFS File System to ECSs (Linux)**. For details about how to mount a file system running Windows, see **Mounting an NFS File System to ECSs (Windows)** and **Mounting a CIFS File System to ECSs (Windows)**.

Step 5 Configure the log directory to the shared file system. It is recommended that each host use different log files.

Step 6 Start applications.

----End

 Copyright © Huawei Technologies Co., Ltd.

11 Other Operations

11.1 Testing SFS Turbo Performance

fio is an open-source I/O pressure testing tool. You can use fio to test the throughput and IOPS of SFS.

Prerequisites

fio has been installed on the ECS. It can be downloaded from **the official website** or from **GitHub**.

Note and Description

The test performance depends on the network bandwidth between the client and server, as well as the capacity of the file system.

Installing fio

The following uses a Linux CentOS system as an example:

1. Download fio.

   yum install fio
2. Install the libaio engine.

   yum install libaio-devel
3. Check the fio version.

   fio --version

File System Performance Data

The performance indicators of SFS Turbo file systems include IOPS and throughput. For details, see **Table 11-1**.

 Copyright © Huawei Technologies Co., Ltd.

Table 11-1 Sample performance data

Parameter	SFS Turbo Standard	SFS Turbo Performance
Maximum capacity	32 TB	32 TB
Maximum IOPS	5,000	20000
Maximum throughput	150 MB/s	350 MB/s
Formula used to calculate the IOPS	IOPS = Min. (5,000, (1,200 + 6 x Capacity))	IOPS = Min. (20,000, (1,500 + 20 x Capacity))

IOPS Calculation Formula

- IOPS of a single file system = Min. (Maximum IOPS, (Baseline IOPS + IOPS per GB x Capacity))

  For an SFS Turbo Performance file system:

  - If the file system capacity is 500 GB: IOPS = Min. (20,000, (1,500 + 20 x 500)) = 11,500
  - If the file system capacity is 1,000 GB: IOPS = Min. (20,000, (1,500 + 20 x 1,000)) = 20,000
- No performance calculation formula is available for the SFS Turbo Standard - Enhanced and SFS Turbo Performance - Enhanced file systems. The IOPS of an SFS Turbo Standard - Enhanced file system is 15,000, and that of an SFS Turbo Performance - Enhanced file system is 100,000.

Common Test Configuration Example

NOTE

The following estimated values are obtained from the test on a single ECS. You are advised to use multiple ECSs to test the performance of **SFS**.

This section uses SFS Turbo Performance as an example to describe the specifications of the ECS.

Specifications: General computing-plus | c3.xlarge.4 | 4 vCPUs | 16 GB

Image: CentOS 7.5 64-bit

Mixed read/write with a read/write ratio of 7:3

- fio command:

  fio --randrepeat=1 --ioengine=libaio --name=test -output=output.log --direct=1 --filename=*/mnt/nfs/test_fio* --bs=4k --iodepth=128 --size=10240M --readwrite=rw --rwmixwrite=30 --fallocate=none

  NOTE

  /mnt/nfs/test_fio indicates the location of the file to be tested. The location must be specific to the file name, which is the **test_fio** file in the **/mnt/nfs** directory in this example. Set it based on the site requirements.

 Copyright © Huawei Technologies Co., Ltd.

- fio result:

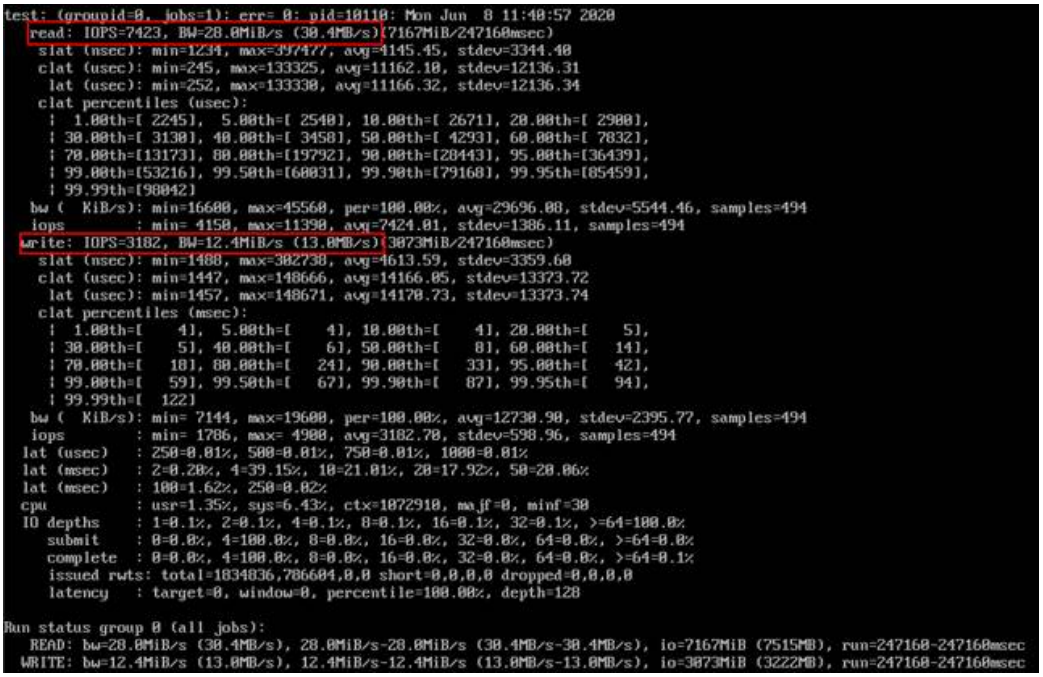

Mixed read/write with a read/write ratio of 3:7

- fio command:

  fio --randrepeat=1 --ioengine=libaio --name=test -output=output.log --direct=1 --filename=*/mnt/nfs/test_fio* --bs=4k --iodepth=128 --size=10240M --readwrite=rw --rwmixwrite=70 --fallocate=none

  NOTE

  /mnt/nfs/test_fio indicates the location of the file to be tested. The location must be specific to the file name, which is the **test_fio** file in the **/mnt/nfs** directory in this example. Set it based on the site requirements.

- Results:

 Copyright © Huawei Technologies Co., Ltd.

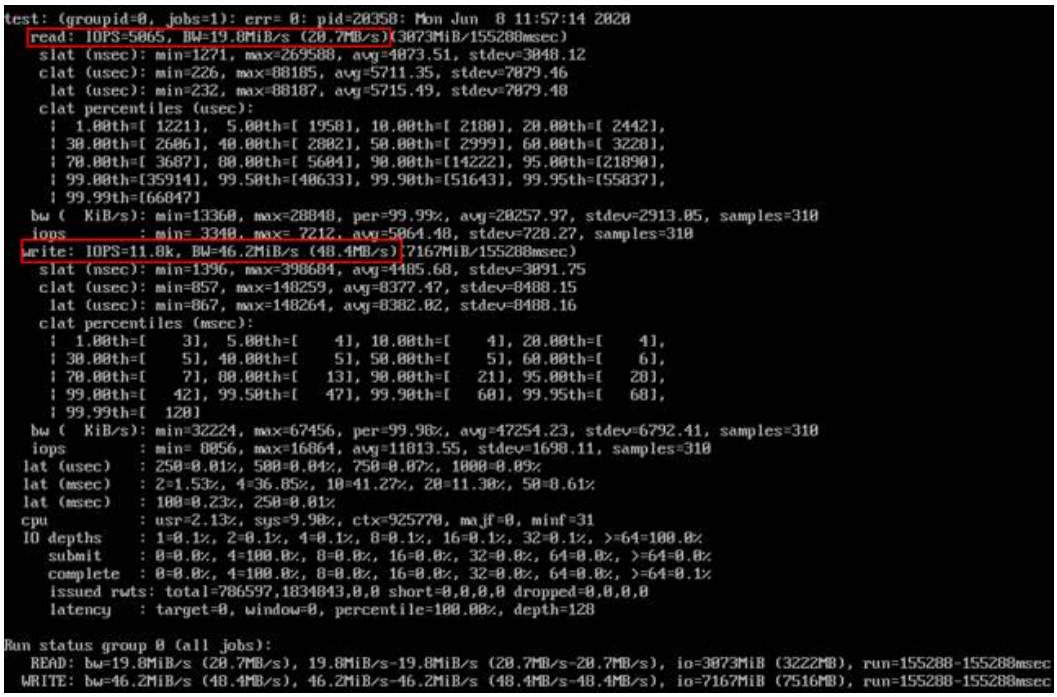

Read IOPS

- fio command:

  fio --randrepeat=1 --ioengine=libaio --name=test -output=output.log --direct=1 --filename=*/mnt/sfs-turbo/test_fio* --bs=4k --iodepth=128 --size=10240M --readwrite=read --fallocate=none

  NOTE

  /mnt/sfs-turbo/test_fio indicates the location of the file to be tested. The location must be specific to the file name, which is the **test_fio** file in the **/mnt/sfs-turbo** directory in this example. Set it based on the site requirements.

- Results:

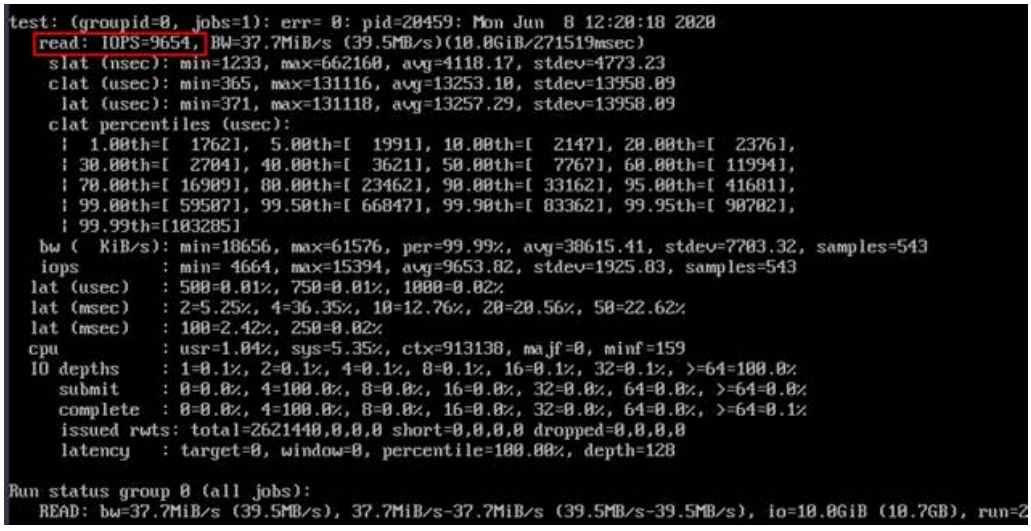

Write IOPS

- fio command:

fio --randrepeat=1 --ioengine=libaio --name=test -output=output.log --direct=1 --filename=*/mnt/sfs-turbo/test_fio* --bs=4k --iodepth=128 --size=10240M --readwrite=write --fallocate=none

NOTE

/mnt/sfs-turbo/test_fio indicates the location of the file to be tested. The location must be specific to the file name, which is the **test_fio** file in the **/mnt/sfs-turbo** directory in this example. Set it based on the site requirements.

- Results:

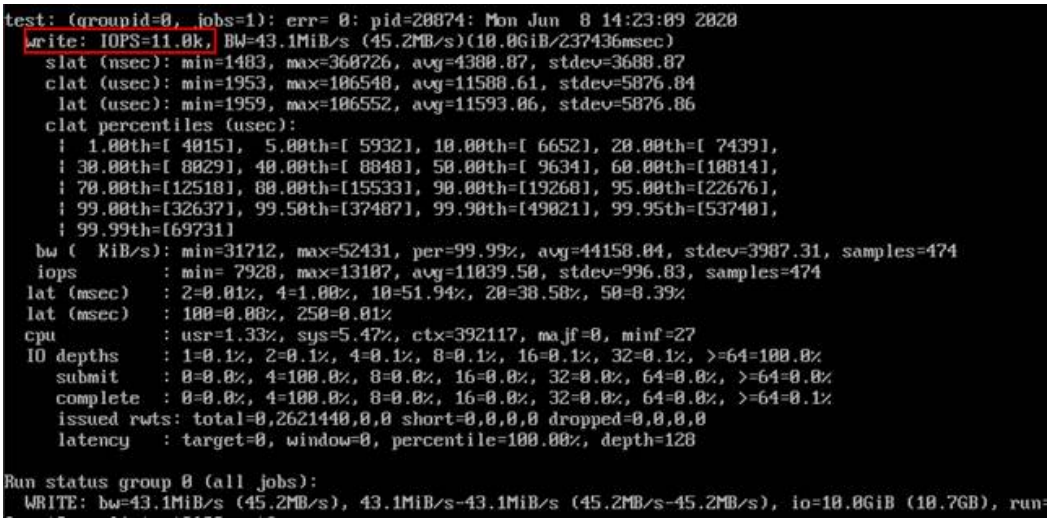

Read bandwidth

- fio command:

  fio --randrepeat=1 --ioengine=libaio --name=test -output=output.log --direct=1 --filename=*/mnt/sfs-turbo/test_fio* --bs=1M --iodepth=128 --size=10240M --readwrite=read --fallocate=none

  NOTE

  /mnt/sfs-turbo/test_fio indicates the location of the file to be tested. The location must be specific to the file name, which is the **test_fio** file in the **/mnt/sfs-turbo** directory in this example. Set it based on the site requirements.

- Results:

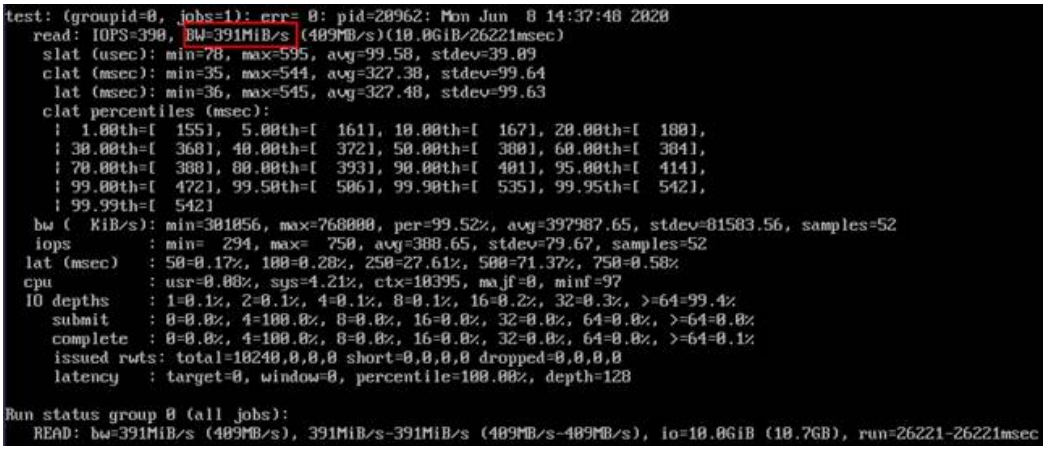

Write bandwidth

 Copyright © Huawei Technologies Co., Ltd.

- fio command:

  fio --randrepeat=1 --ioengine=libaio --name=test -output=output.log --direct=1 --filename=*/mnt/sfs-turbo/test_fio* --bs=1M --iodepth=128 --size=10240M --readwrite=write --fallocate=none

  NOTE

  /mnt/sfs-turbo/test_fio indicates the location of the file to be tested. The location must be specific to the file name, which is the **test_fio** file in the **/mnt/sfs-turbo** directory in this example. Set it based on the site requirements.

- Results:

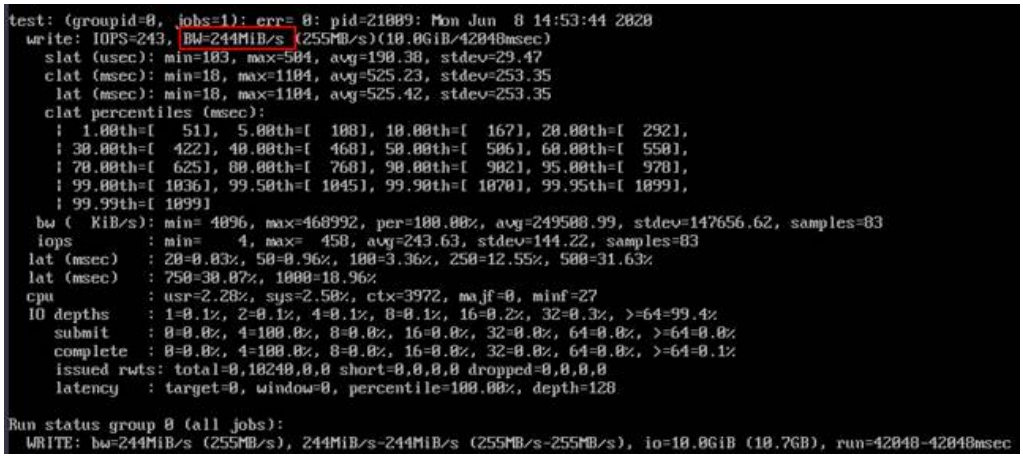

11.2 Mounting a File System to an ECS Running Linux as a Non-root User

Description

By default, a Linux ECS allows only the **root** user to run the **mount** command for mounting a file system. However, if the permissions of user **root** are assigned to other common users, such users can also run the **mount** command for file system mounting. The following describes how to mount a file system to a Linux ECS as a common user. The EulerOS is used as an example.

Prerequisites

- A non-**root** user has been created on the ECS.
- A file system has been created and can be mounted to the ECS by the **root** user.
- You have obtained the mount point of the file system.

Procedure

Step 1 Log in to the ECS as user **root**.

Step 2 Assign the permissions of user **root** to the non-**root** user.

1. Run the **chmod 777 /etc/sudoers** command to change the **sudoers** file to be editable.

 Copyright © Huawei Technologies Co., Ltd.

2. Run the **vi /etc/resolv.conf** command to edit the **sudoers** file.
3. Add a common user under the **root** account. In the following figure, user **Mike** is added.

Figure 11-1 Adding a user

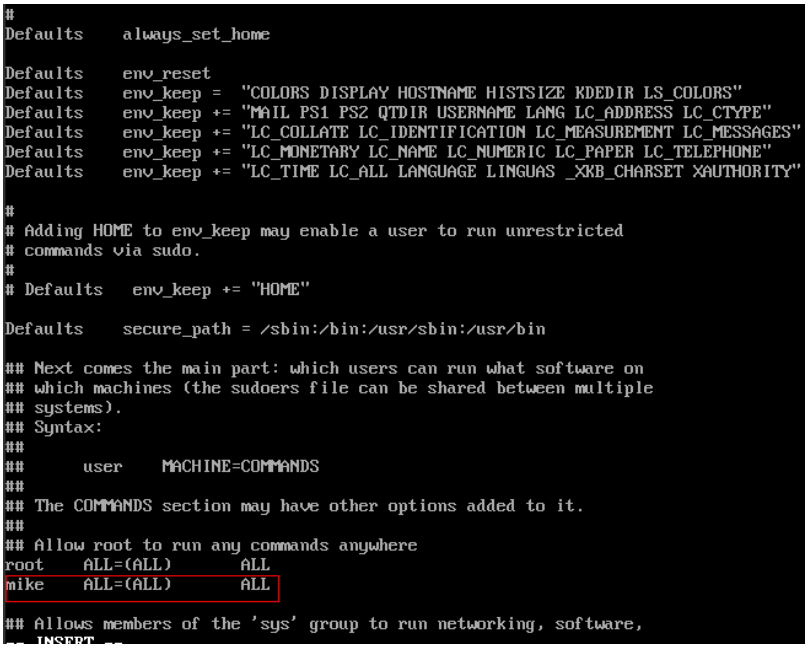

4. Press **Esc**, input **:wq**, and press **Enter** to save and exit.
5. Run the **chmod 440 /etc/sudoers** command to change the **sudoers** file to be read-only.

Step 3 Log in to the ECS as user **Mike**.

Step 4 Run the following command to mount the file system. For details about the mounting parameters, see **Table 11-2**.

mount -t nfs -o vers=3,timeo=600,noresvport,nolock *Mount point Local path*

Table 11-2 Parameter description

Parameter	Description
Mount point	The format for an SFS Capacity-Oriented file system is *File system domain name*:/*Path*, for example, **example.com:/share-***xxx*. The format for an SFS Turbo file system is *File system IP address*:/, for example, **192.168.0.0:/**. **NOTE** *x* is a digit or letter. If the mount point is too long to display completely, expand the column to view the full mount point.
Local path	Local path on the ECS, used to mount the file system, for example, **/local_path**.

 Copyright © Huawei Technologies Co., Ltd.

Step 5 Run the following command to view the mounted file system.

mount -l

If the command output contains the following information, the file system is mounted successfully.

```
example.com:/share-xxx on /local_path type nfs (rw,vers=3,timeo=600,nolock,addr=)
```

----End

11.3 Configuring DNS for Cache-accelerated SFS Turbo

Configure DNS for the cache-accelerated SFS Turbo servers so that you can use the cache-accelerated SFS Turbo file systems.

Procedure

Step 1 Log in to the HUAWEI CLOUD management console and choose **Network** > **Domain Name Service** > **DNS Resolution** > **Private Zones**.

Step 2 Click **Create Private Zone** in the upper right corner.

 Copyright © Huawei Technologies Co., Ltd.

Figure 11-2 Create Private Zone

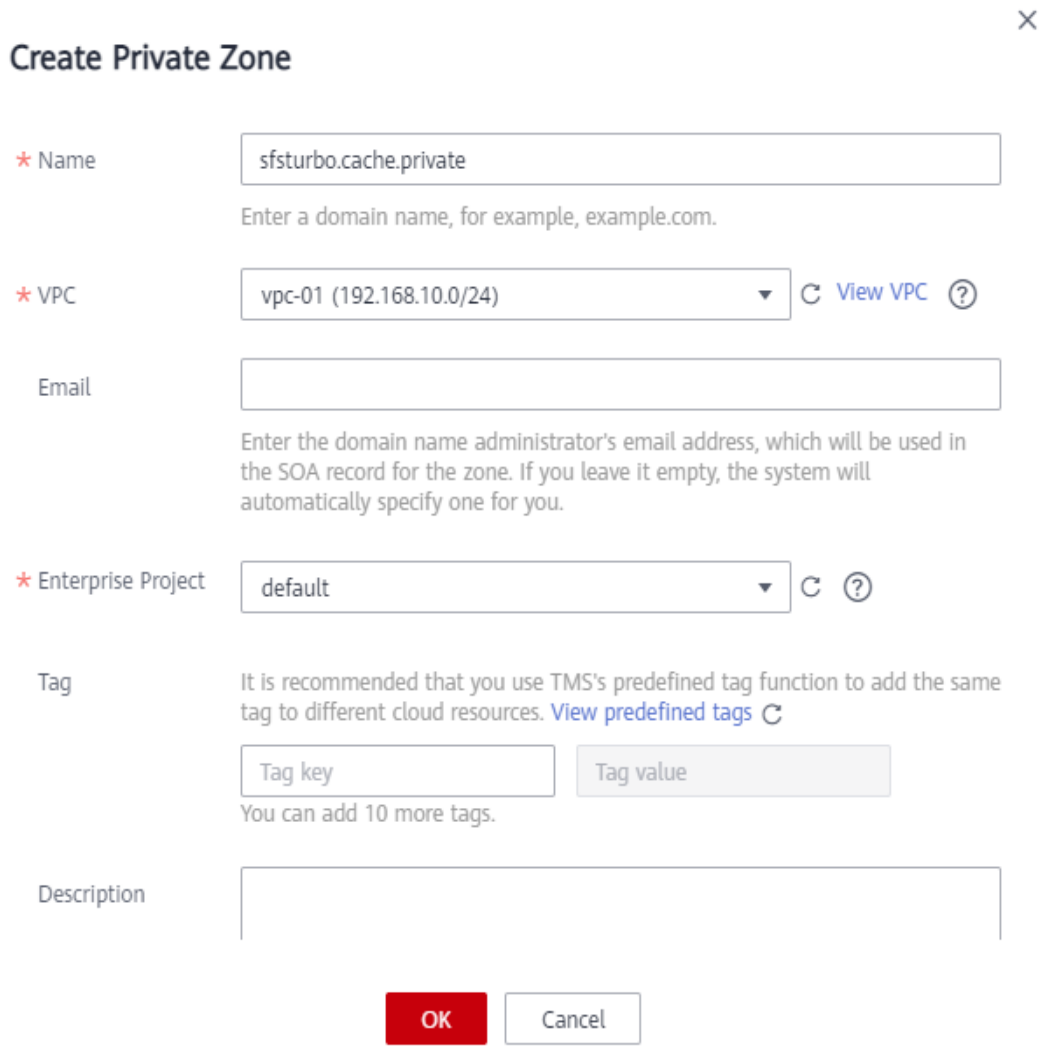

NOTE

Name: Enter a domain name that has not been used. **sfsturbo.cache.private** is an example domain name.

Enterprise Project: Select the VPC and enterprise project where the cache-accelerated SFS Turbo file system belongs.

For detailed parameter settings, see **Creating a Private Zone**.

Step 3 Confirm the settings and click **OK**.

Step 4 Click the name of the created public zone to go to the record sets page.

Figure 11-3 Viewing a private zone

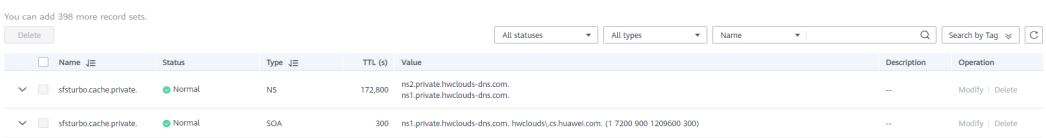

 Copyright © Huawei Technologies Co., Ltd.

Step 5 Click **Add Record Set** in the upper right corner. Set required parameters based on **Figure 11-4**.

Figure 11-4 Add Record Set

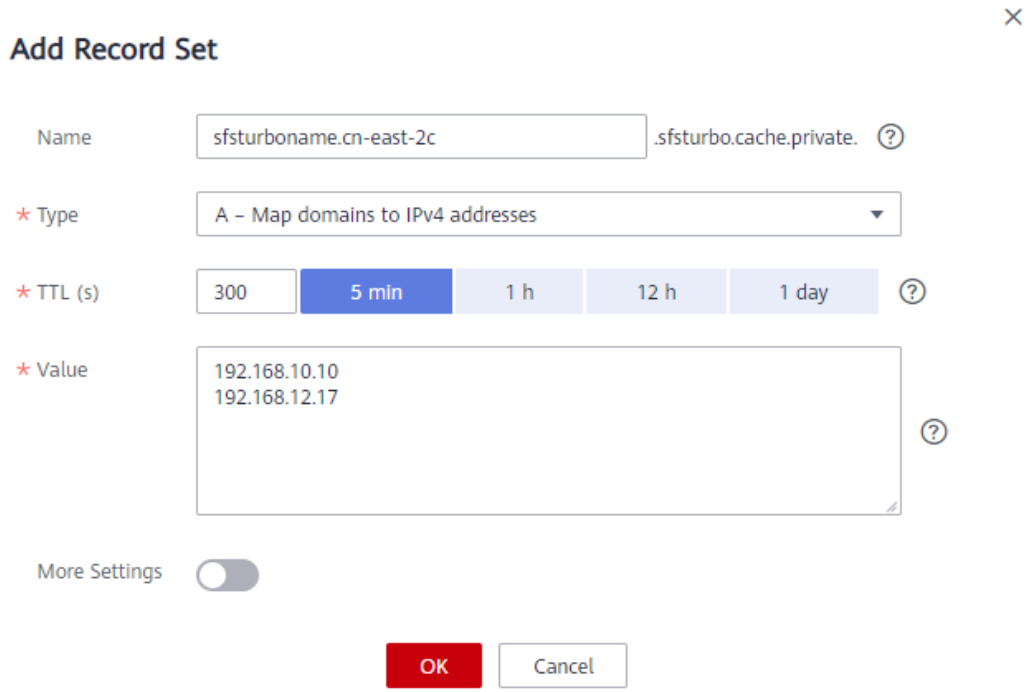

NOTE

Name: It is recommended that you set the name in the format of [*cache-accelerated SFS Turbo file system name*].[*region*]. In this example, **sfsturboname.cn-east-2c** indicates that the name of the cache-accelerated SFS Turbo file system is **sfsturboname**, and the region is **cn-east-2**, in which **c** indicates CN East-Shanghai2 AZ3.

Type: Select the A type.

TTL: Select **5 min**.

Value: indicates the IP addresses of destination servers to which requests are forwarded after the domain name is accessed. Set this parameter to the IP addresses of the servers housing cache-accelerated SFS Turbo file systems.

Step 6 Confirm the settings and click **OK**. The record set is successfully added.

Figure 11-5 Record set added successfully

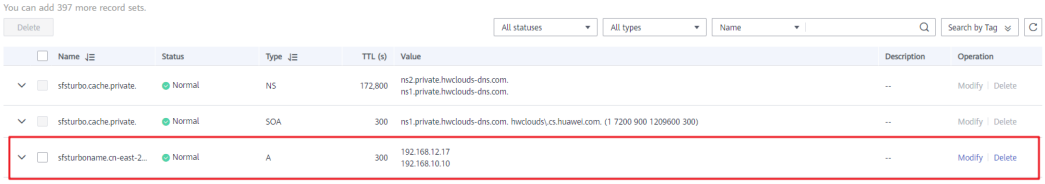

Step 7 After the record set is successfully added, you can mount the cache-accelerated SFS Turbo file system using domain name **sfsturboname.cn-east-2c.sfsturbo.cache.private**.

----End

 Copyright © Huawei Technologies Co., Ltd.

11.4 Data Migration

11.4.1 Migrating Data Using Direct Connect

Context

Users can migrate data from local Network Attached Storage (NAS) devices to SFS Turbo using Direct Connect.

In this solution, a Linux ECS is created to enable the connection between a local NAS device and SFS Turbo and to migrate data from the NAS device to the cloud.

You can also refer to this solution to migrate data from the NAS in the cloud to SFS Turbo. For details, see **Migrating Data from the NAS in the Cloud to SFS**.

Limitations and Constraints

- Only ECSs running Linux can be used for data migration.
- The UID and GID of your file are no longer consistent after synchronization.
- The file access modes are no longer consistent after synchronization.

Prerequisites

- You have enabled and configured Direct Connect. For details, see ***Direct Connect User Guide***.
- You have obtained a Linux ECS.
- You have created an SFS Turbo file system and have obtained the mount point of the file system.
- You have obtained the mount point of the local NAS device.

Procedure

Step 1 Log in to the ECS management console.

Step 2 Log in to the created Linux ECS to access the local NAS device and the SFS Turbo file system.

Step 3 Run the following mount command to access the local NAS storage device:

```
mount -t nfs -o vers=3,timeo=600,noresvport,nolock Mount point of the local NAS device /mnt/src
```

Step 4 Run the following mount command to access the file system:

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

Step 5 Run the following command on the Linux ECS to install the rclone tool:

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

Step 6 Run the following command to synchronize data:

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

 Copyright © Huawei Technologies Co., Ltd.

NOTE

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

- **transfers**: number of files that can be transferred concurrently
- **checkers**: number of local files that can be scanned concurrently
- **P**: data copy progress
- **copy-links**: copying soft links

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

----End

Migrating Data from the NAS in the Cloud to SFS

To migrate data from a NAS device in the cloud to your SFS Turbo file system, ensure that the NAS and file system are in the same VPC or you can use Cloud Connect.

For details about how to configure Cloud Connect, see ***Direct Connect User Guide***.

11.4.2 Migrating Data Using the Internet

Context

Users can migrate data from local NAS devices to SFS Turbo using the Internet for cloud service expansion.

In this solution, a Linux server is created in the cloud and on-premises respectively for data migration from the local NAS devices to the cloud. The inbound and outbound directions of port 22 on these two servers are accessible. The on-premises server is used to access the local NAS devices, and the ECS is used to access SFS Turbo.

You can also refer to this solution to migrate data from the NAS in the cloud to SFS Turbo. For details, see **Migrating Data from Cloud NAS to SFS Turbo**.

Limitations and Constraints

- Data cannot be migrated from local NAS devices to SFS Capacity-Oriented file systems using the Internet.
- Only ECSs running Linux can be used for data migration.
- The UID and GID of a file are no longer consistent after synchronization.
- The file access modes are no longer consistent after synchronization.
- The inbound and outbound directions of port 22 are accessible.

Prerequisites

- A Linux server has been created on the cloud and on-premises respectively.
- The elastic IP addresses have been configured for the servers to ensure that the two servers can communicate with each other.

 Copyright © Huawei Technologies Co., Ltd.

- You have created an SFS Turbo file system and have obtained the mount point of the file system.
- You have obtained the mount point of the local NAS device.

Procedure

Step 1 Log in to the ECS management console.

Step 2 Log in to the created on-premises server **client1** and run the following command to access the local NAS device:

```
mount -t nfs -o vers=3,timeo=600,noresvport,nolock Mount point of the local NAS device /mnt/src
```

Step 3 Log in to the created Linux ECS **client2** and run the following command to access the SFS Turbo file system:

```
mount -t nfs -o vers=3,timeo=600,noresvport,nolock Mount point of the SFS Turbo file system  /mnt/dst
```

Step 4 Run the following command on **client1** to install the rclone tool:

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

Step 5 Run the following command on **client1** to configure the environment:

```
rclone config
No remotes found - make a new one
n) New remote
s) Set configuration password
q) Quit config
n/s/q> n
name> remote name (New name)
Type of storage to configure.
Enter a string value. Press Enter for the default ("").
Choose a number from below, or type in your own value
24 / SSH/SFTP Connection
   \ "sftp"
Storage> 24 (Select the SSH/SFTP number)
SSH host to connect to
Enter a string value. Press Enter for the default ("").
Choose a number from below, or type in your own value
 1 / Connect to example.com
   \ "example.com"
host> ip address (IP address of client2)
SSH username, leave blank for current username, root
Enter a string value. Press Enter for the default ("").
user> user name (Username of client2)
SSH port, leave blank to use default (22)
Enter a string value. Press Enter for the default ("").
port> 22
SSH password, leave blank to use ssh-agent.
y) Yes type in my own password
g) Generate random password
n) No leave this optional password blank
y/g/n> y
Enter the password:
password: (Password for logging in to client2)
Confirm the password:
password: (Confirm the password for logging in to client2)
Path to PEM-encoded private key file, leave blank or set key-use-agent to use ssh-agent.
Enter a string value. Press Enter for the default ("").
key_file> (Press Enter)
The passphrase to decrypt the PEM-encoded private key file.

Only PEM encrypted key files (old OpenSSH format) are supported. Encrypted keys
in the new OpenSSH format can't be used.
```

 Copyright © Huawei Technologies Co., Ltd.

```
y) Yes type in my own password
g) Generate random password
n) No leave this optional password blank
y/g/n> n
When set forces the usage of the ssh-agent.
When key-file is also set, the ".pub" file of the specified key-file is read and only the associated key is requested from the ssh-agent. This allows to avoid `Too many authentication failures for *username*` errors when the ssh-agent contains many keys.
Enter a boolean value (true or false). Press Enter for the default ("false").
key_use_agent> (Press Enter)
Enable the use of the aes128-cbc cipher. This cipher is insecure and may allow plaintext data to be recovered by an attacker.
Enter a boolean value (true or false). Press Enter for the default ("false").
Choose a number from below, or type in your own value
 1 / Use default Cipher list.
   \ "false"
 2 / Enables the use of the aes128-cbc cipher.
   \ "true"
use_insecure_cipher> (Press Enter)
Disable the execution of SSH commands to determine if remote file hashing is available.
Leave blank or set to false to enable hashing (recommended), set to true to disable hashing.
Enter a boolean value (true or false). Press Enter for the default ("false").
disable_hashcheck>
Edit advanced config? (y/n)
y) Yes
n) No
y/n> n
Remote config
--------------------
[remote_name]
type = sftp
host=(client2 ip)
user=(client2 user name)
port = 22
pass = *** ENCRYPTED ***
key_file_pass = *** ENCRYPTED ***
--------------------
y) Yes this is OK
e) Edit this remote
d) Delete this remote
y/e/d> y
Current remotes:

Name                 Type
====                 ====
remote_name          sftp

e) Edit existing remote
n) New remote
d) Delete remote
r) Rename remote
c) Copy remote
s) Set configuration password
q) Quit config
e/n/d/r/c/s/q> q
```

Step 6 Run the following command to view the **rclone.conf** file in **/root/.config/rclone/rclone.conf**:

```
cat /root/.config/rclone/rclone.conf
[remote_name]
type = sftp
host=(client2 ip)
user=(client2 user name)
port = 22
pass = ***
key_file_pass = ***
```

Step 7 Run the following command on **client1** to synchronize data:

```
rclone copy /mnt/src remote_name:/mnt/dst -P --transfers 32 --checkers 64
```

 Copyright © Huawei Technologies Co., Ltd.

NOTE

- Replace *remote_name* in the command with the actual remote name.
- The parameters are described as follows. Set **transfers** and **checkers** based on the system specifications.
  - **transfers**: number of files that can be transferred concurrently
  - **checkers**: number of local files that can be scanned concurrently
  - **P**: data copy progress

After data synchronization is complete, go to the target SFS Turbo file system to check whether the migration is successful.

----End

Migrating Data from Cloud NAS to SFS Turbo

To migrate data from the cloud NAS to SFS Turbo using the Internet, perform the preceding steps.

11.4.3 Migrating Data Between an SFS Capacity-Oriented File System and an SFS Turbo File System

Context

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

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

Limitations and Constraints

- Only ECSs running Linux can be used for data migration.
- The Linux ECS, SFS Capacity-Oriented file system, and SFS Turbo file system must be in the same VPC.

Prerequisites

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

Procedure

Step 1 Log in to the ECS management console.

Step 2 Log in to the created Linux ECS that can access SFS Capacity-Oriented and SFS Turbo file systems.

Step 3 Run the following command to mount file system 1 (either the SFS Capacity-Oriented 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 [Mount point of file system 1] /mnt/src
```

 Copyright © Huawei Technologies Co., Ltd.

Step 4 Run the following command to mount file system 2 (the other system that you have 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 [Mount point of file system 2] /mnt/dst
```

Step 5 Run the following command on the Linux ECS to install the rclone tool:

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

Step 6 Run the following command to synchronize data:

```
rclone copy /mnt/src /mnt/dst -P --transfers 32 --checkers 64
```

NOTE

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

- **transfers**: number of files that can be transferred concurrently
- **checkers**: number of files that can be scanned concurrently
- **P**: data copy progress

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

----End

 Copyright © Huawei Technologies Co., Ltd.

A Change History

Release Date	What's New
2020-09-14	This issue is the twenty-fourth official release. Updated the following content: Changed "SFS file system" to "SFS Capacity-Oriented file system."
2020-06-08	This issue is the twenty-third official release. Updated the following content: Updated the screenshots in section "SFS Turbo Performance Test."
2020-04-01	This issue is the twenty-second official release. Updated the following content: Deleted the content related to the open beta test of policy-based access control.
2019-11-11	This issue is the twenty-first official release. Updated the following content: Added section "Configuring Cross-Account Access."
2019-08-23	This issue is the twentieth official release. Updated the following content: Added session "Migrating Data to SFS Turbo."
2019-07-22	This issue is the nineteenth official release. Updated the following content: Added section "Creating Alarm Rules."
2019-05-30	This issue is the eighteenth official release. Updated the following content: • Added section "Permissions Management."

 Copyright © Huawei Technologies Co., Ltd.

Release Date	What's New
2019-04-30	This issue is the seventeenth official release. Updated the following content: ● The capacity expansion feature of the SFS Turbo file system is brought offline, and the related description is deleted from section "File System Resizing." ● Added monitoring indicators to section "Monitoring." ● Added descriptions about the performance calculation formula in section "SFS Turbo Performance Test."
2019-03-30	This issue is the sixteenth official release. Updated the following content: ● Modified the outline structure and optimized section names. ● Added section "SFS Turbo Performance Test."
2019-03-01	This issue is the fifteenth official release. Updated the following content: ● Added the contents related to enterprise projects in section "Viewing File Systems." ● Optimized the description of user rights and user root permission parameters in section "Managing VPCs."
2019-02-15	This issue is the fourteenth official release. Updated the following content: ● Added section "Quotas."
2019-01-30	This issue is the thirteen official release. Updated the following content: ● Added the viewing monitoring statistics operation in section "Viewing a File System." ● Optimized section "Viewing Monitoring Statistics." ● Changed the "Troubleshooting" section to "Scalable File Service Troubleshooting" for separate publishing.
2018-12-07	This issue is the twelfth official release. Updated the following content: ● Added the support for the CIFS protocol and related descriptions. ● Added description about expanding the capacity for an SFS Turbo file system in section "File System Resizing."

Copyright © Huawei Technologies Co., Ltd.

Release Date	What's New
2018-11-02	This issue is the eleventh official release. Updated the following content: ● Added the description of the SFS Turbo file system. ● Added the description of file systems that support encryption. ● Added the description of managing file system's VPCs. ● Added the description of mounting a file system to a Windows OS.
2018-06-15	This issue is the tenth official release. Updated the following content: ● Added section "Viewing Monitoring Statistics." ● Added section "Querying Traces."
2018-05-25	This issue is the ninth official release. Updated the following content: Split the document into several parts for release.
2018-02-07	This issue is the eighth official release. Updated the following content: ● Updated section "Configuring DNS", and changed the IP addresses of the DNS server for resolving domain names of file systems. ● Updated section "Limitations and Constraints", and added the sentence "You can mount file systems to all Elastic Cloud Servers (ECSs) that support the NFSv3 protocol." ● Deleted "Can a File System Be Mounted to a Windows-based ECS" from "FAQs." ● Added "Does the Security Group of VPC Affect SFS" in "FAQs." ● Added section "How Do I Purchase SFS?"
2018-01-12	This issue is the seventh official release. Updated the following content: ● Updated the screenshots based on changes in the GUI. ● Added the secondary confirmation and images for deleting a file system in section "Deleting a File System." ● Optimized section "SFS." ● Updated section "Application Scenarios" and section "Limitations and Constraints."
2017-09-23	This issue is the sixth official release. Updated the following content: ● Adjusted the document content structure. ● Added chapter "Typical Applications."

 Copyright © Huawei Technologies Co., Ltd.

Release Date	What's New
2017-04-10	This issue is the fifth official release. Updated the following content: ● Updated the prerequisites and steps, and added a description of mounting a file system with a domain name in section "Mounting a File System to ECSs." ● Added section "Configuring DNS."
2017-03-22	This issue is the fourth official release. Updated the following content: ● Added the versions of operating systems to which SFS file systems can be mounted. ● Updated the description in section "How Is File System Data Stored."
2016-12-30	This issue is the third official release. Updated the following content: ● Updated the steps in section "Mounting a File System to ECSs." Specifically, deleted firewall configurations and added the operation of copying the shared path. ● Updated the screenshots in section "Viewing a File System."
2016-09-30	This issue is the second official release. Updated the following content: ● Added section "Mounting a File System Fails" and section "An ECS Fails to Access a File System."
2016-08-25	This issue is the first official release.

Copyright © Huawei Technologies Co., Ltd.