Scalable File Service

User Guide

Issue	01
Date	2025-12-09

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. 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 Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions


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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

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

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

Contents

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

1 Permissions Management

1.1 Creating an IAM User and Granting SFS Permissions

This section 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.

Figure 1-1 shows the process flow for granting permissions.

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**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Process Flow

Figure 1-1 Process for granting SFS permissions

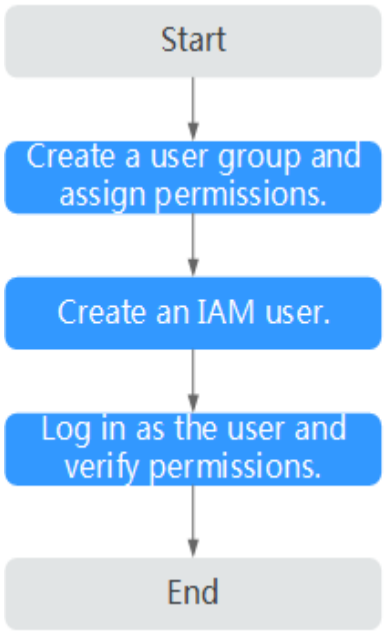

1. **Create a user group and assign permissions** to it.

   Create a user group on the IAM console and assign the **SFS3 ReadOnlyAccess** permissions 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 the console as the created user and switch to the authorized region. Choose **Scalable File Service** > **General Purpose File System** > **File Systems**. Click **Create File System** in the upper right corner. If a message appears indicating that you have insufficient permissions, the **SFS3 ReadOnlyAccess** policy is in effect.

1.2 Creating Custom Policies for General-Purpose File Systems

You can create custom policies to supplement the system-defined policies of SFS. For details about actions supported in custom policies, see **Permissions and Supported Actions**.

To create a custom policy, choose either visual editor or JSON.

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy grammar.
- JSON: Create a JSON policy from scratch or based on an existing policy.

For details, see **Creating a Custom Policy**.The following lists examples of common SFS custom policies.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Example Custom Policies

- Example 1: Grant permission to create general-purpose file systems.

```
{
    "Version": "1.1",
    "Statement": [
        {
            "Action": [
                "sfs3:fileSystem:createFileSystem"
            ],
            "Effect": "Allow"
        }
    ]
}
```

- Example 2: Grant permission to deny general-purpose file system deletion.

  A policy with only "Deny" permissions must be used together with other policies. If the permissions granted to an IAM user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

  Assume that you want to grant the permissions of the **SFS3 FullAccess** policy to a user but want to prevent them from deleting general-purpose file systems. You can create a custom policy for denying file system deletion, and attach this policy together with the **SFS3 FullAccess** policy to the user. As an explicit deny in any policy overrides any allows, the user can perform all operations on general-purpose file systems excepting deleting them. Example policy denying file system deletion:

```
{
    "Version": "1.1",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": [
                "sfs3:fileSystem:deleteFileSystem"
            ]
        }
    ]
}
```

1.3 Managing General-Purpose File System Permissions

Overview

You can add authorization rules for a general-purpose file system to grant different permissions to different clients.

On the console, each new general-purpose file system has a default rule (Read/Write, no_root_squash, All IP addresses). This rule grants all client users read/write permissions to access the file system and does not map the root user to an unprivileged account. You can edit or delete this rule if needed.

Constraints

An IP address cannot be added to two authorization rules at the same time.

Authorized IP Addresses

You can configure authorized IP addresses using CIDR blocks.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

A CIDR block uses a variable-length subnet mask to show the ratio of the network bits to host address bits within a range of IP addresses.

A suffix value is added at the end of an IP address to form a CIDR block. This suffix shows the bits of the network address. For example, 192.1.1.0/24 is an IPv4 CIDR block, in which the first 24 bits (192.1.1) are the network address.

Any IP address whose first 24 bits are the same as those of 192.1.1.0 will be applied with this authorization rule. In other words, 192.1.1.1 and 192.1.1.1/32 have the same effect.

Types of Permissions

There are read/write permissions and user permissions.

Table 1-1 Read/Write permissions

Permission	Description
Read/Write	Client users have the read/write permissions.
Read-only	Client users have the read-only permissions.

Table 1-2 User permissions

Permission	Description
no_root_squash	All client users (including **root**) access the file system as who they are, instead of being mapped to the **nobody** user.
all_squash	All client users access the file system as the **nobody** user.
root_squash	The **root** user accesses the file system as the **nobody** user.

Adding Authorization Rules

You can add authorization rules on the console for permissions management.

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

Step 2 In the file system list, find the general-purpose file system you want to add authorization and click its name to go to its details page.

Step 3 On the **Permissions Management** tab, click **Add Authorization Rule**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 1-2 Add Authorization Rule

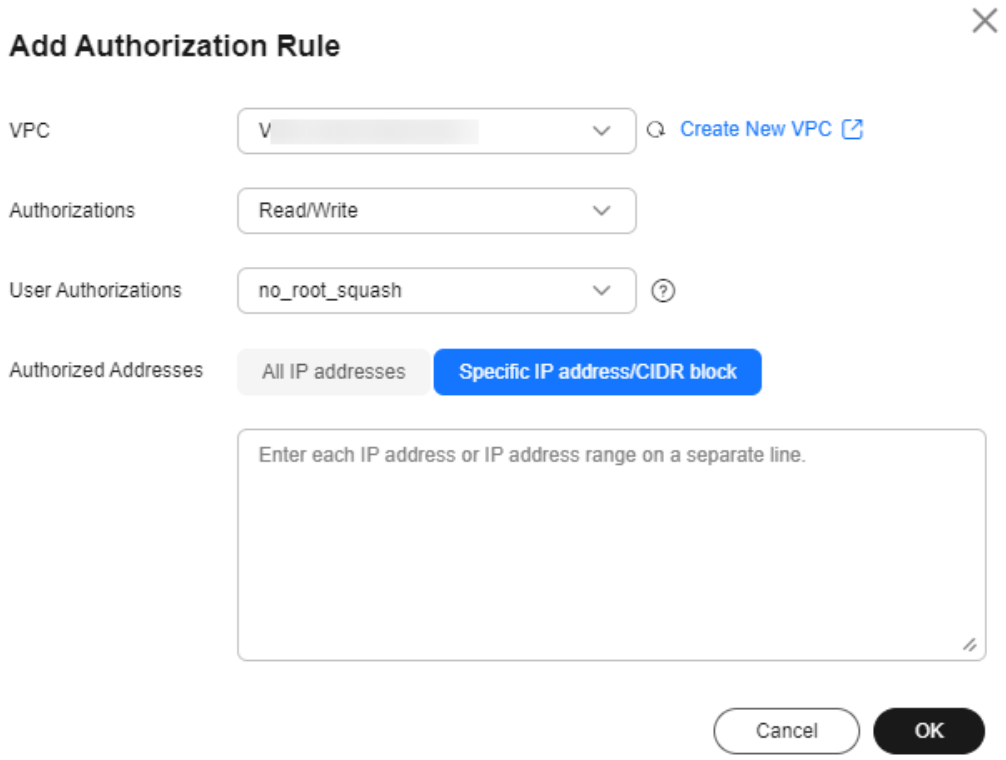

Step 4 On the displayed page, add authorization based on **Table 1-3**.

Table 1-3 Parameter description

Parameter	Description
VPC	Select the VPC you want to add, for example, **vpc-30e0**. If no VPC is available, create one.
Authorizations	You can select **Read/Write** or **Read-only**. **Read/Write** is preselected.
User Authorizations	You can select **no_root_squash**, **root_squash**, or **all_squash**. • **no_root_squash** allows the **root** user on the client to access the general-purpose file system as **root**. • **root_squash** allows the **root** user on the client to access the general-purpose file system as the **nobody** user. • **all_squash** allows any user on the client to access the general-purpose file system as the **nobody** user, and the user can modify and delete the file system.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
Authorized Addresses	You can select **All IP addresses** or **Specific IP address/CIDR block**. **All IP addresses** is preselected. • Enter a valid IPv4 address or range that is not starting with 0 except 0.0.0.0/0. If you add **0.0.0.0/0**, any IP address within this VPC will be authorized to access the file system. Do not enter an IP address or IP address range starting with any number ranging from 224 to 255, for example 224.0.0.1 or 255.255.255.255, because class D and class E IP addresses are not supported. IP addresses starting with 127 are also not supported. If you enter an invalid IP address or IP address range, the rule may fail to be added, or the authorization will not work. • If you enter an IP address range, enter it in the format of *IP address/mask*. For example, enter 192.168.1.0/24. Do not use the following format:192.168.1.0-255 or 192.168.1.0-192.168.1.255. The number of bits in a subnet mask must be an integer ranging from 0 to 31, and mask value **0** is valid only in 0.0.0.0/0. • For details about IP address ranges, see **Authorized IP Addresses**. **NOTE** If you select **Specific IP address/CIDR block**, you can add multiple IP addresses or CIDR blocks. Enter each one on a separate line. After the authorized addresses are added, you can click the number shown under **Authorized Addresses** in the permissions management list to check their information.

Step 5 Confirm the information and click **OK**.

----End

Related Operations

You can click **Edit** in the **Operation** column of a rule to modify the read/write permission and user permission, or click **Delete** to delete a rule.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

2 General-Purpose File System Management

2.1 Configuring a VPC Endpoint

Background

VPC Endpoint provides reliable channels to connect VPCs to general-purpose file systems. By configuring VPC endpoints, compute resources in VPCs can access general-purpose file systems. To find more about VPC Endpoint, see **What Is VPC Endpoint?** Before mounting a general-purpose file system to compute resources, you need to create a VPC endpoint in the region where the compute resources reside. You can check if VPC Endpoint is supported in a region on the console.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 2-1 Configuration process

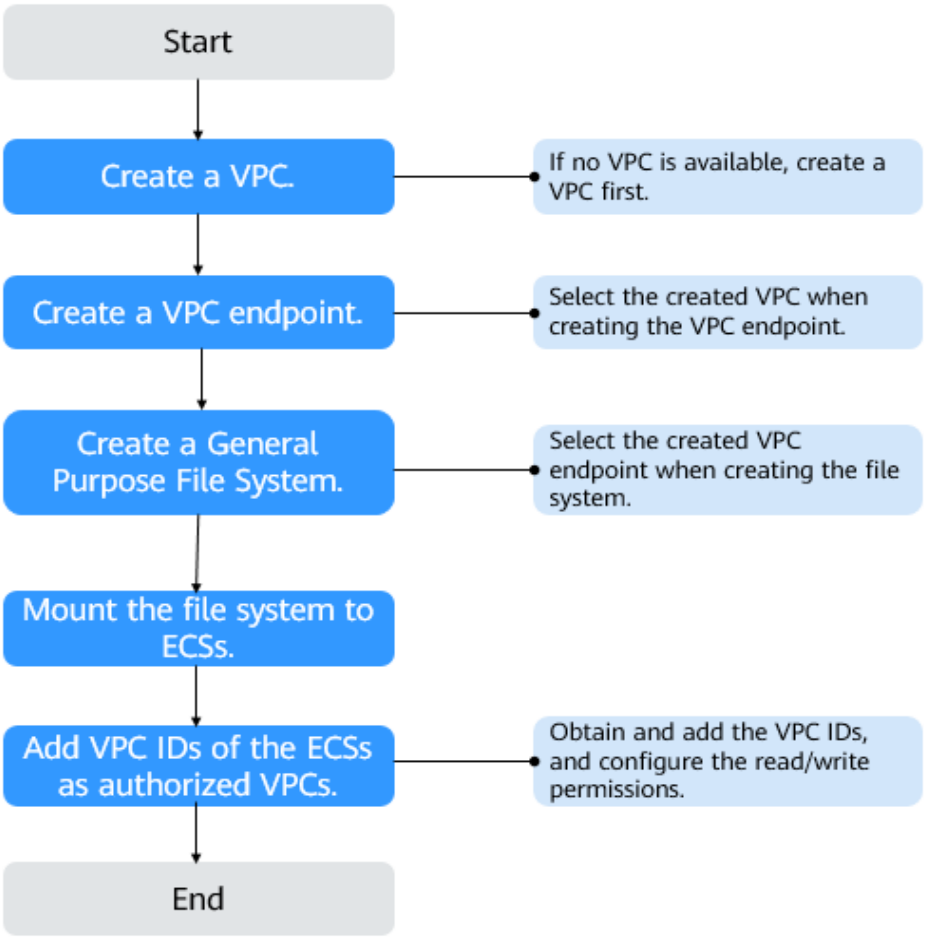

Prerequisites

1. A VPC is available.

   If no VPC is available, create one by referring to **Creating a VPC with a Subnet** in the *Virtual Private Cloud User Guide*.

2. ECSs are available and they are in the created VPC.

   If no ECSs are available, buy ECSs by referring to **Purchasing and Using an ECS** and **Logging in to an ECS**.

Procedure

Step 1 **Log in to the VPC Endpoint console**.

Step 2 On the **VPC Endpoints** page, click **Buy VPC Endpoint**.

The **Buy VPC Endpoint** page is displayed.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 2-2 Buy VPC Endpoint

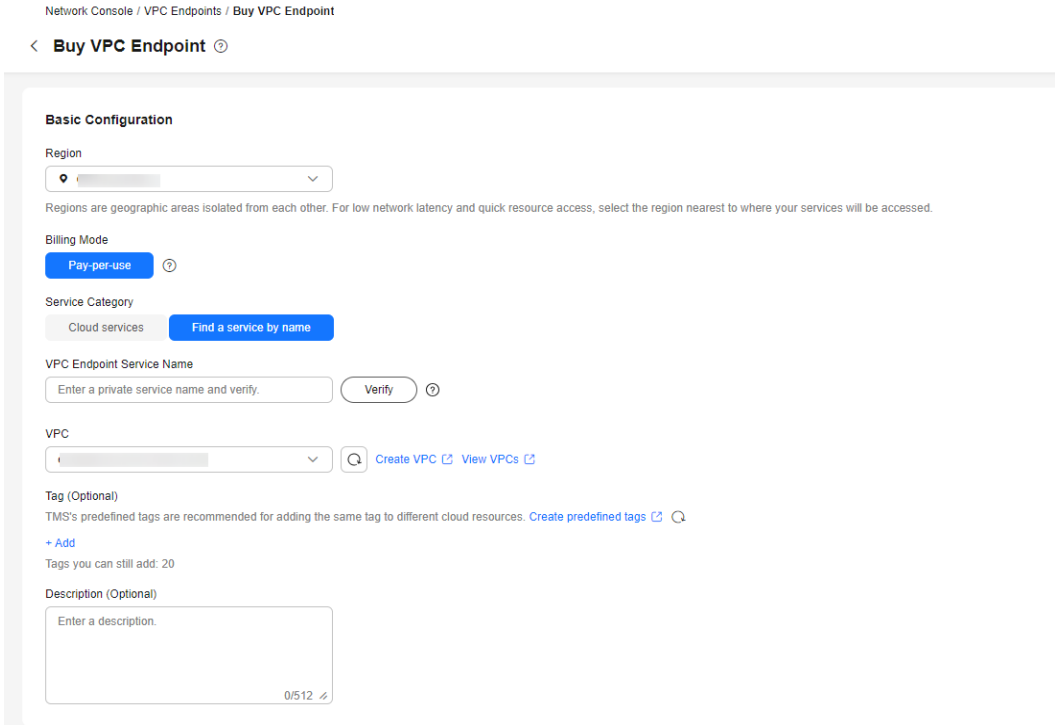

Step 3 Configure VPC endpoint parameters.

Table 2-1 VPC endpoint parameters

Parameter	Description
Region	Region where the VPC endpoint is located. Ensure that this region is the same as the one where the planned general-purpose file system resides.
Billing Mode	**Pay-per-use** is preselected by default, but you will not be billed for the endpoint purchased for general-purpose file systems.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
Service Category	Select **Find a service by name**. Enter a VPC endpoint service name based on the region selected. • If the CN North-Beijing4 region is selected, enter **cn-north-4.com.myhuaweicloud.v4.storage.lz13**. • If the CN South-Guangzhou region (AZ1) is selected, enter **cn-south-1.com.myhuaweicloud.v4.obsv2**. **NOTE** General-purpose file systems created in AZ1 of the CN South-Guangzhou region cannot be mounted to containers. • If the CN South-Guangzhou region (AZ6) is selected, enter **cn-south-1.com.myhuaweicloud.v4.obsv2.storage.lz06**. • If the CN East-Shanghai1 region is selected, enter **cn-east-3.com.myhuaweicloud.v4.storage.lz07**. • If the CN Southwest-Guiyang1 region is selected, enter **com.myhuaweicloud.cn-southwest-2.ipv4.sfs**. • If the CN-Hong Kong region is selected, enter **ap-southeast-1.com.myhuaweicloud.v4.obsv2.storage.lz005**. After entering the service name, click **Verify**. If **Service name found** is displayed, proceed with subsequent steps. If **Service name not found** is displayed, check whether the entered service name is correct. If the problem persists, **submit a service ticket**.
VPC	VPC where the planned general-purpose file system and ECSs reside.
Tag	Optional VPC endpoint tags. Each tag consists of a key and a value. You can add a maximum of 10 tags to a VPC endpoint. Tag keys and values must meet the requirements listed in **Table 2-2**. **NOTE** If you have created a predefined tag in TMS, you can select the corresponding tag key and value. For details about predefined tags, see **Predefined Tags**.

Table 2-2 describes the tag parameters.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Table 2-2 Tag parameters

Parameter	Description	Example Value
Tag key	Each tag has a unique key. You can customize a key or select the key of an existing tag created in TMS. A tag key: • Cannot be an empty string. • Must be unique for each resource. • Can contain a maximum of 128 characters. • Can contain letters, digits, spaces, and the following characters: _.:=+-@; cannot start or end with a space, or start with _sys_.	Key_0001
Tag value	A tag value can be repetitive or left blank. A tag value: • Can be an empty string. • Can contain a maximum of 255 characters. • Can contain letters, digits, spaces, and the following characters _.:/=+-@; cannot start or end with a space.	Value_0001

Step 4 Click **Next**.

- If you do not need to modify the configuration, click **Submit**.
- If you need to modify the configuration, click **Previous**, modify the configuration as needed, and then click **Submit**.

Step 5 Go back to the VPC endpoint list and check whether the status of the VPC endpoint changes to **Accepted**. If so, the VPC endpoint has been connected to the VPC endpoint service.

----End

2.2 Creating a General-Purpose File System

You can create a general-purpose file system and mount it to multiple cloud servers. The servers can then share this file system.

Prerequisites

1. A VPC is available.

   If no VPC is available, create one by referring to **Creating a VPC with a Subnet** in the *Virtual Private Cloud User Guide*.

2. ECSs are available and they are in the created VPC.

   If no ECSs are available, buy ECSs by referring to **Purchasing and Using an ECS** and **Logging in to an ECS**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Creating a General-Purpose File System

Step 1 Log in to the SFS console.

Step 2 In the upper right corner of the page, click **Create File System**.

Step 3 Set the parameters shown in **Figure 2-3**. **Table 2-3** describes the parameters.

Figure 2-3 Creating a general-purpose file system

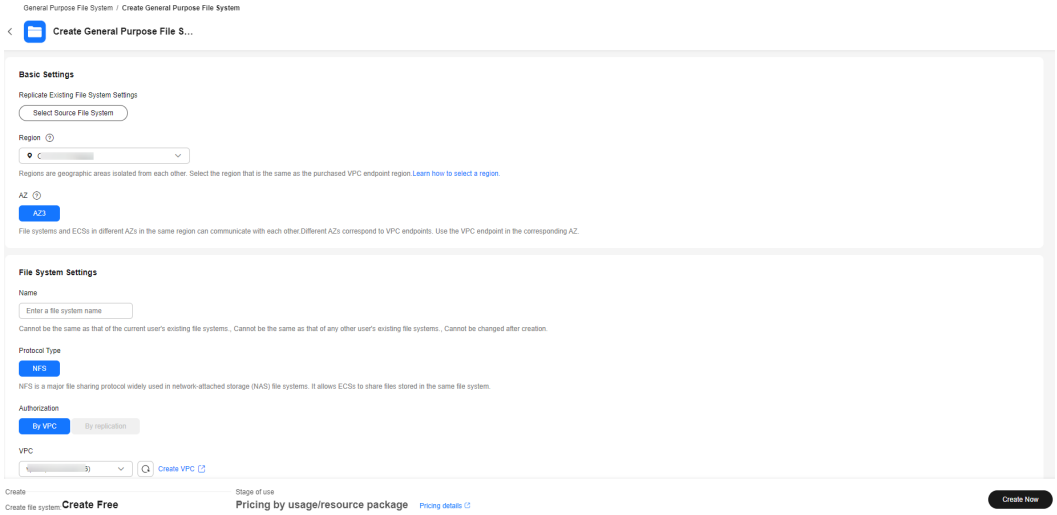

Table 2-3 General-purpose file system parameters

Parameter	Description
Replicate Existing File System Settings	Optional. Click **Select Source File System**. On the displayed page, select a source general-purpose file system from the list. After you click **OK**, the system automatically copies the region, AZ, protocol, authorization, and tags of the source file system. You can change some or all of the replicated settings, if needed.
Region	Mandatory The region of the tenant. Select a region from the drop-down list in the upper left corner of the page. Select the region where the cloud servers and VPC endpoint reside.
AZ	A geographical area with an independent network and an independent power supply. You are advised to select the AZ where the cloud servers reside.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
Name	The user-defined name of the general-purpose file system. It cannot be changed after the general-purpose file system is created. Set an appropriate name when creating the file system. In SFS, general-purpose file systems are named according to the globally applied DNS naming rules: ● The name of a general-purpose file system must be globally unique. It cannot be the same as the name of any existing general-purpose file system, including one created by the current user or any other user. You must wait at least 30 minutes before you can reuse the name of a deleted general-purpose file system. ● The name must be 3 to 63 characters long. Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. ● The name cannot start or end with a period (.) or hyphen (-), and cannot contain two consecutive periods (..) or contain a period (.) and a hyphen (-) adjacent to each other. ● The name cannot be an IP address. **NOTE** When you use virtual-hosted-style URLs to access SFS over HTTPS, if the general-purpose file system name contains periods (.), the certificate verification will fail. In this case, you are advised not to use periods (.) in file system names.
Protocol Type	General-purpose file systems support file access using NFS (NFSv3 only).
Authorization	You can select **By VPC** or **By replication**. **By VPC**: Select the VPC where the ECSs and VPC endpoint reside. An ECS cannot access a file system in a different VPC. Select the VPC where your ECSs reside. You can add VPCs later on the file system details page. **By replication**: This option is available after you select to replicate settings from a source file system. If you select authorization by replication, you can go to the source file system's details page to view the replicated permissions.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
VPC	If you select authorization by VPC, you need to manually select a VPC. Select the IAM project to which the target VPC belongs and then select the target VPC. For details about IAM projects, see **Project Management**. Select the VPC where the ECSs and VPC endpoint reside. An ECS cannot access a general-purpose file system in a different VPC. Select the VPC where your ECSs reside. You can add VPCs later on the file system details page.
Tag	Optional When creating a general-purpose file system, you can add tags to it. Tags help you identify file systems, and you can classify and search for file systems by tag. A tag is composed of a key-value pair. • Key: mandatory. A tag key can contain a maximum of 128 characters. It can contain letters, digits, and spaces representable in UTF-8 and special characters (_.:=+-@). It cannot start or end with a space and cannot be left empty. Tag keys starting with _sys_ are system tags, and you cannot start a tag key with _sys_. • Value: optional. A tag value can contain a maximum of 255 characters. It can contain letters, digits, and spaces representable in UTF-8 and special characters (_.:=+-@) and can be left empty. It cannot start or end with a space. **NOTICE** • You can add a maximum of 20 tags to a general-purpose file system. • Tag keys of the same general-purpose file system must be unique. • Except for tagging the general-purpose file system during file system creation, you can also add, modify, or delete tags for existing general-purpose file systems.
Purchase Resource Packages	Optional You can select this parameter and buy desired resource packages. The packages take effect immediately after payment. For more information, see **Resource Packages** in the *Scalable File Service Billing*. **NOTICE** A resource package can only be used to pay for the storage used by general-purpose file systems in the same region. Any usage in excess of the package quota will be billed in the pay-per-use mode. Only available in the CN-Hong Kong region.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Step 4 Click **Create Now**.

Step 5 Confirm the file system information and click **Submit**.

Step 6 Go back to the file system list.

If you can see the general-purpose file system in the list, it is created successfully. If not, the creation fails. **Submit a service ticket** for technical consultation.

Figure 2-4 General-purpose file system created

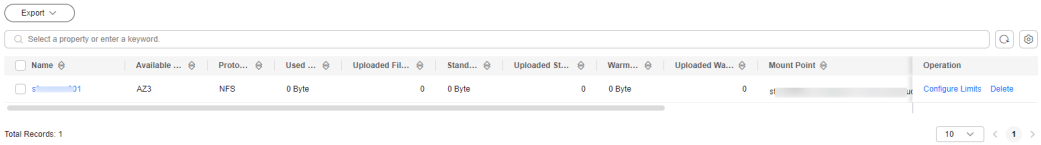

----End

2.3 Mounting a General-Purpose File System

2.3.1 Mounting a General-Purpose File System to Linux ECSs

After creating a general-purpose file system, you need to mount it to cloud servers so that they can share the file system. This section describes how to mount a file system to ECSs.

NOTE

- The operations of mounting a file system to BMSs and containers (created on CCE) are the same as those of ECSs. To use file systems for CCE, see **Storage > Storage Overview** or **Storage > SFS** and then complete the deployment on the CCE console.
- The mount operation may vary depending on the server OS. Perform operations based on your server OS.
- General-purpose file systems cannot be mounted to 32-bit Linux servers.

Constraints

NOTE

This constraint only applies to local paths (mount points) and does not affect other files or directories.

Metadata of the local paths (mount points) cannot be modified. Specifically, the following operations cannot be performed on the local paths' metadata:

- **touch**: Update file access time and modified time.

- **rm**: Delete files or directories.

- **cp**: Replicate files or directories.

- **mv**: Move files or directories.

- **rename**: Rename files or directories.

- **chmod**: Modify permissions on files or directories.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

- **chown**: Change the owners of files or directories.
- **chgrp**: Change the group of a file or directory.
- **ln**: Create hard links.
- **link**: Create hard links.
- **unlink**: Delete hard links.

The **atime**, **ctime**, and **mtime** attributes of a local path (root directory of the mount point) are the current time of the file system server. Each time a root directory attribute is queried, the current time of the server is returned.

Prerequisites

- You have checked the type of the OS on each ECS. Different OSs use different commands to install the NFS client.
- You have created a general-purpose file system and obtained its mount point from the file system list.
- At least one ECS that is in the same VPC as the general-purpose file system is available.
- You have configured the IP address of the DNS server on the ECSs. The DNS server is used to resolve the domain name of the general-purpose file system. For details, see **Binding an EIP**.
- Before mounting a general-purpose file system to a LinuxECS, you need to configure a VPC endpoint. For details, see **Configuring a VPC Endpoint**.

Mounting a General-Purpose File System

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

Step 2 **Log in the ECS** as user **root**. You can log in using the console or a remote access tool (such as PuTTY).

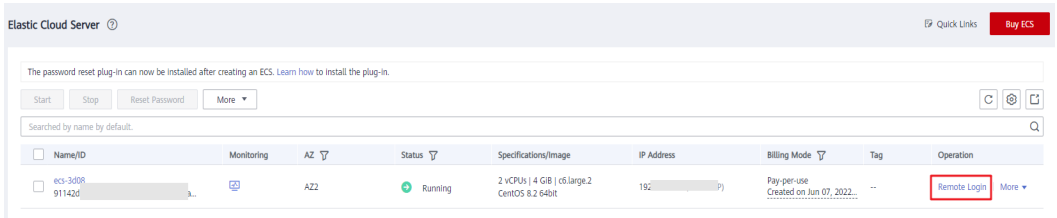

Step 3 Install the NFS client.

1. Check whether the NFS software package is installed.
   - On CentOS, Red Hat, Oracle Enterprise Linux, SUSE, EulerOS, Fedora, or openSUSE, run the following command:

     rpm -qa|grep nfs
   - On Debian or Ubuntu, run the following command:

     dpkg -l nfs-common

   If a command output similar to the following is displayed, the NFS software package has been installed and you can go to **Step 4**. If information similar to the following is not displayed, go to **Step 3.2**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

- On CentOS, Red Hat, EulerOS, Fedora, or Oracle Enterprise Linux, the command output is as follows:
  ```
  libnfsidmap
  nfs-utils
  ```
- On SUSE or openSUSE, the command output is as follows:
  ```
  nfsidmap
  nfs-client
  ```
- On Debian or Ubuntu, the command output is as follows:
  ```
  nfs-common
  ```

2. Install the NFS software package.

   NOTE

   The following commands require that ECSs be connected to the Internet. Or, the installation will fail.

   - On CentOS, Red Hat, EulerOS, Fedora, or Oracle Enterprise Linux, run the following command:
     ```
     sudo yum -y install nfs-utils
     ```
   - On Debian or Ubuntu, run the following command:
     ```
     sudo apt-get install nfs-common
     ```
   - On SUSE or openSUSE, run the following command:
     ```
     zypper install nfs-client
     ```

Step 4 Check whether the domain name in the mount point of the general-purpose file system can be resolved.

```
nslookup <domain-name-of-the-general-purpose-file-system>
```

NOTE

- Obtain the domain name of a general-purpose file system from its mount point. For example, if the mount point of a general-purpose file system is *xxx*:**/sfs-name-001**, *xxx* is the file system domain name, and **sfs-name-001** is the file system name.
- If the **nslookup** command cannot be used, you can run **yum install bind-utils** to install the **bind-utils** software package.

- If the resolution succeeds, proceed with the following steps.
- If the domain name cannot be resolved, configure the DNS server IP address before mounting the general-purpose file system. For details, see **Configuring a DNS Server for Domain Name Resolution**.

Step 5 Mount the NFS file system.

- To mount the file system root directory, run the following commands:

```
mkdir <local-path>
mount -t nfs -o vers=3,timeo=600,noresvport,nolock,tcp <mount-point> <local-path>
```

- To mount a file system subdirectory, run the following commands:

```
mkdir <local-path>|<subdirectory>
mount -t nfs -o vers=3,timeo=600,noresvport,nolock,tcp <file-system-domain-name>:/<subdirectory>
<local-path>
```

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

NOTE

- You can obtain the file system domain name from its mount point in the file system list. For example, if the mount point of a general-purpose file system is *xxx*:**/sfs-name-001**, *xxx* is the file system domain name, and **sfs-name-001** is the file system name.
- If any other resources, such as a disk, have been mounted on the desired local path, create a new path. (NFS clients do not refuse repeated mounts. If there are repeated mounts, information of the last successful mount is displayed.)
- A file system can only be mounted to the ECSs that are in the same VPC as the file system.

Table 2-4 describes the variables in the mount command.

Table 2-4 Parameter description

Parameter	Description
<local-path>	A local directory on the ECS used to mount the file system, for example, **/local_path**.
<mount-point>	The format of a general-purpose file system is *<file-system-domain-name>*:/*<file-system-name>*, for example, **example.com:/***xxx*. **NOTE** • Variable *x* is a digit or letter. • If the mount point is too long to display completely, you can adjust the column width. • Hover over the mount point to view the full mount command.
vers	The file system version. Only NFSv3 is supported currently, so the value is fixed at **3**.
timeo	The waiting time before the NFS client retransmits a request. The unit is 0.1 second. The recommended value is **600**.
noresvport	Whether the NFS client uses a new TCP port when it re-establishes a network connection to the NFS server. It is strongly recommended that you specify **noresvport**, which ensures that your file system remains uninterrupted after a network reconnection or recovery.
lock/nolock	Whether to use the NLM protocol to lock files on the client. If **nolock** is specified, the lock is valid only for applications on the same client. It is invalid for applications on any other clients. **nolock** is recommended. If this parameter is not specified, **lock** is used. Then, other clients cannot write data to the file system. General-purpose file systems do not support operations of non-local locks. If a client uses a non-local lock, it will experience slow writes due to the failure to obtain the lock. In this case, specify **nolock** to avoid such issues.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
tcp/udp	The protocol used by NFS clients to send requests to the server. You can use either UDP or TCP. General-Purpose File System does not support UDP. Therefore, you need to set **proto** to **tcp** for general-purpose file systems.

Figure 2-5 Mount Point

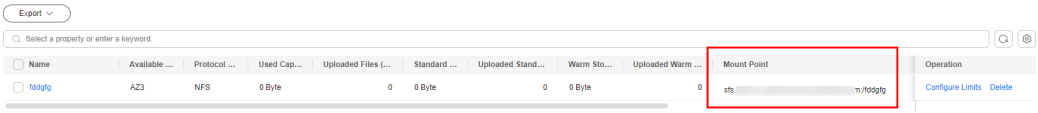

When mounting the file system, you can add performance optimization options described in **Table 2-5**. Use commas (,) to separate them. For example:

```
mount -t nfs -o vers=3,timeo=600,nolock,rsize=1048576,wsize=1048576,hard,retrans=3,tcp,noresvport,ro,async,noatime,nodiratime <mount-point> <local-path>
```

Table 2-5 Performance optimization options

Parameter	Description
rsize	Maximum number of bytes in each read request that the client can receive when reading data from a file on the server. The actual data size is less than or equal to this parameter setting. The value of **rsize** must be a positive integer that is a multiple of **1024**. If the entered value is smaller than **1024**, the value is automatically set to **4096**. If the entered value is greater than **1048576**, the value is automatically set to **1048576**. By default, this parameter is set through a negotiation between the server and the client. You are advised to set this parameter to the maximum value **1048576**.
wsize	Maximum number of bytes in each write request that the client can send when writing data to a file on the server. The actual data size is less than or equal to this parameter setting. The value of **wsize** must be a positive integer that is a multiple of **1024**. If the entered value is smaller than **1024**, the value is automatically set to **4096**. If the entered value is greater than **1048576**, the value is automatically set to **1048576**. By default, this parameter is set through a negotiation between the server and the client. You are advised to set this parameter to the maximum value **1048576**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
soft/hard	Value **soft** indicates soft mounts. With **soft** specified, if an NFS request times out, the client returns an error to the calling program. Value **hard** indicates hard mounts. With **hard** specified, if an NFS request times out, the client continues to request until the request is successful. **hard** is used by default.
retrans	The number of retransmission times before the client returns an error. The recommended value is **1**.
tcp/udp	If **mountproto** is not specified, the client will mount the file system using UDP first. If the UDP network cannot be connected, the client will mount the file system using TCP after freezing for several seconds. The security group of the file system does not allow inbound traffic over UDP ports, so you need to specify **mountproto=tcp**.
ro/rw	• **ro** indicates that the file system is mounted as read-only. • **rw** indicates that the file system is mounted as read/write. **rw** is used by default. If neither **ro** nor **rw** is specified, the file system will be mounted as read/write.
sync/async	**sync** indicates that data is written to the server immediately. **async** indicates that data is first written to the cache and then to the server. Value **async** is recommended. Synchronous writes require that an NFS server returns a success message after all data is written to the server, which brings long latency.
noatime	If you do not need to record the file access time, set this parameter. This prevents overheads caused by frequent access to modify the time.
nodiratime	If you do not need to record the directory access time, set this parameter. This prevents overheads caused by frequent access to modify the time.

NOTE

You are advised to use the default values for the parameters with no recommendations provided.

Step 6 View the mounted general-purpose file system.

```
mount -l
```

If the command output contains the following information, the file system has been mounted. You can access the file system from the ECS to read or write data.

```
<mount-point> on </local-path> type nfs (rw,vers=3,timeo=600,nolock,addr=)
```

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

> **NOTE**
>
> The maximum size of a file that can be written to a general-purpose file system is 240 TB.

Step 7 (Optional) Configure the **fstab** file to configure file system auto mount upon system startup.

After a client ECS is restarted, it loses the file system mount information. You can configure auto mount in the **fstab** file to ensure that the ECS automatically mounts the file system when it restarts.

1. Open the **/etc/fstab** file.

```
vi /etc/fstab
```

At the end of the file, add the file system information, for example:

```
<mount-point> </local-path> nfs vers=3,timeo=600,nolock,tcp 0 0
```

Replace *<mount-point>* and *</local-path>* with actual values. You can obtain the mount point from the file system list. Each record in the **/etc/fstab** file represents a mount. Each record has six fields, as described in **Table 2-6**.

Table 2-6 Mount fields

Field	Description
<mount-point>	The address or location of the general-purpose file system you are mounting. Set it to the mount point in the **mount** command in **Step 5**.
/local_path	A directory on the ECS used to mount the file system. Set it to the local path in the **mount** command in **Step 5**.
nfs	The file system or partition mount type. Set it to **nfs**.
vers=3,timeo=600,nolock,tcp	The mount options. Use commas (,) to separate multiple options. • **vers**: The file system version. Value **3** indicates NFSv3. • **timeo**: The waiting time before the NFS client retransmits a request. The unit is 0.1 second. The recommended value is **600**. • **nolock**: Whether to use the NLM protocol to lock files on the client. If **nolock** is specified, the lock is valid only for applications on the same client. It is invalid for applications on any other clients. **nolock** is recommended. If this parameter is not specified, **lock** is used. Then, other clients cannot write data to the file system. • **tcp**: The TCP transmission protocol.
0	Choose whether to use **dump** to back up the file system. • **0**: Dump backup is not used. • An integer greater than zero means that dump backup is used. A file system with a smaller integer is dumped earlier than one with a larger integer.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Field	Description
0	Choose whether to use fsck to check file systems when the ECS is starting and specify the sequence for checking file systems. • **0**: File systems are not checked. • By default, this field is set to **1** for the root directory. The values for other directories start from **2**, and one with a smaller integer is checked earlier than one with a larger integer.

NOTICE

For optimal system performance, configure file system information based on the mount example provided. If needed, you can customize certain mount options. Note that the customization may affect system performance.

2. Press **Esc**, enter **:wq**, and press **Enter** to save and exit.
3. (Optional) View the content of the **/etc/fstab** file after the update.

```
cat /etc/fstab
```

NOTE

If auto mount fails due to a network issue, add the **sleep** option and a time in front of the mount command in the **rc.local** file, and mount the file system after the NFS service is started.

```
sleep 10s && sudo mount -t nfs -o vers=3,timeo=600,noresvport,nolock,tcp <mount-point>|<local-path>
```

----End

FAQs

SFS Mount

2.4 Viewing a General-Purpose File System

You can search for a general-purpose file system by name, AZ, protocol type, mount point, or creation time, and then view the file system basic information.

NOTE

Viewing details of general-purpose file systems depends on the VPC service. Ensure that the required role or policy has been configured.

The permissions of the **SFS3 ReadOnlyAccess** policy already include the permissions of **VPC ReadOnlyAccess**, which are required for querying general-purpose file system details. An IAM user assigned the **SFS3 ReadOnlyAccess** policy does not need to have the **VPC ReadOnlyAccess** policy assigned explicitly.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Procedure

Step 1 Log in to the SFS console.

Step 2 In the general-purpose file system list, view the file systems you have created. Table 2-7 describes the file system parameters.

Table 2-7 General-purpose file system parameters

Parameter	Description
Name	Name of the general-purpose file system, for example, **sfs-name-001**
Availability Zone	Availability zone (AZ) where the file system resides
Protocol Type	File system protocol, which is **NFS**
Used Capacity	File system space already used for data storage **NOTE** This information is refreshed every hour.
Uploaded Files (Count)	Number of files that have been uploaded to the file system **NOTE** This information is refreshed every hour.
Standard Storage Used Capacity	Total standard storage used in the general-purpose file system
Uploaded Standard Files (Count)	Total number of files that use standard storage in the general-purpose file system
Warm Storage Used Capacity	Total infrequent access storage used in the general-purpose file system
Uploaded Warm Files (Count)	Total number of files that use infrequent access storage in the general-purpose file system
Mount Point	Address or location of the general-purpose file system. The format is *<file-system-domain-name>:/<file-system-name>*, for example, **example.com:/sfs-name-001**. **NOTE** • If the mount point is too long to display completely, you can adjust the column width. • Hover over the mount point to view the full mount command.
Tag	Tag information of the general-purpose file system
Created	Time when the file system was created
Operation	The **Configure Limits** and **Delete** buttons are available.

Step 3 Click the name of the general-purpose file system. On the **Basic Information** tab, view more information about the file system, as shown in Figure 2-6.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 2-6 Details of a general-purpose file system

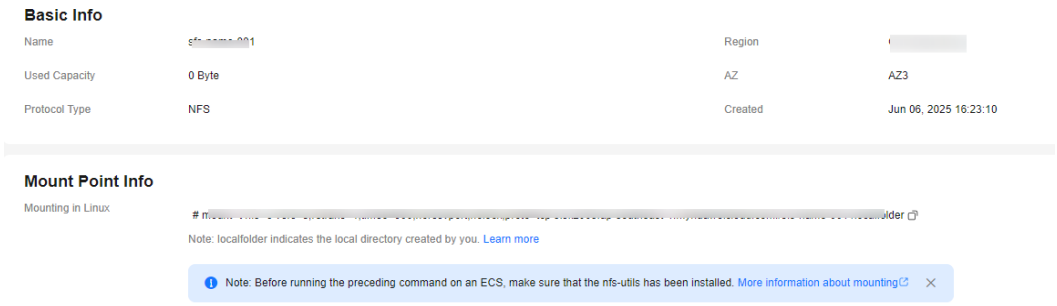

----End

2.5 Managing General-Purpose File System Limits

SFS does not limit the capacity of each general-purpose file system. To enable users to properly allocate and manage capacity and resources, SFS supports limits management for general-purpose file systems. You can configure and remove file system limits as required.

You can configure a capacity limit and the maximum number of files for a general-purpose file system.

Constraints

- SFS takes about 10 to 20 minutes to update the general-purpose file system's used capacity, so the used capacity displayed on the console may not be the latest. For this reason, the actual used capacity may surpass the limit you configured, or the displayed used capacity may not decrease right after some data is deleted from the file system.
- After limits are configured, when the general-purpose file system's used capacity reaches the configured limit, new files or directories cannot be created in the file system and append operations to the file system will fail.
- Configuring limits brings certain risks, so you are advised to evaluate and fully test and verify services before doing so.

Configuring Limits

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

Step 2 In the navigation pane on the left, choose **General Purpose File System** > **File Systems**.

Step 3 In the file system list, find the target general-purpose file system and click **Configure Limits**.

Step 4 On the **Limits Management** tab, click **Configure** to open the page shown in **Figure 2-7**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 2-7 Configure Limits

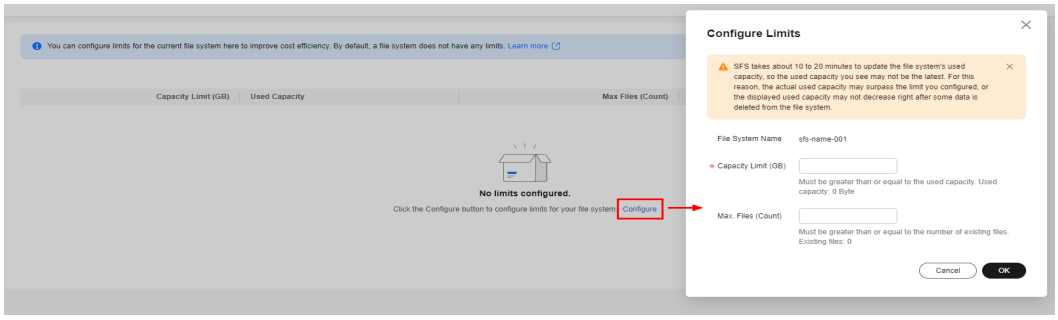

Step 5 Configure limits for the general-purpose file system.

Capacity Limit (GB): This is a required field. Enter a value greater than 0 and at least equal to the used capacity, in GB.

Max. Files (Count): This field is optional. Enter a value greater than 0 and at least equal to the number of existing files.

Step 6 Click **OK**. View the limits details in the list.

Figure 2-8 Limits details

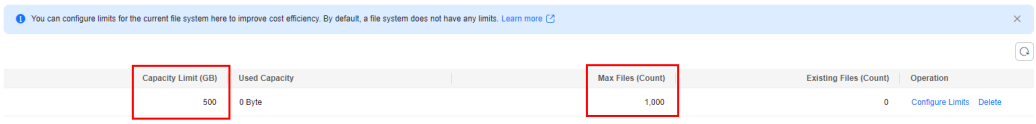

----End

Modifying Limits

Step 1 On the **Limits Management** tab, click **Configure Limits** to open the page shown in **Figure 2-9**.

Figure 2-9 Modify Limits

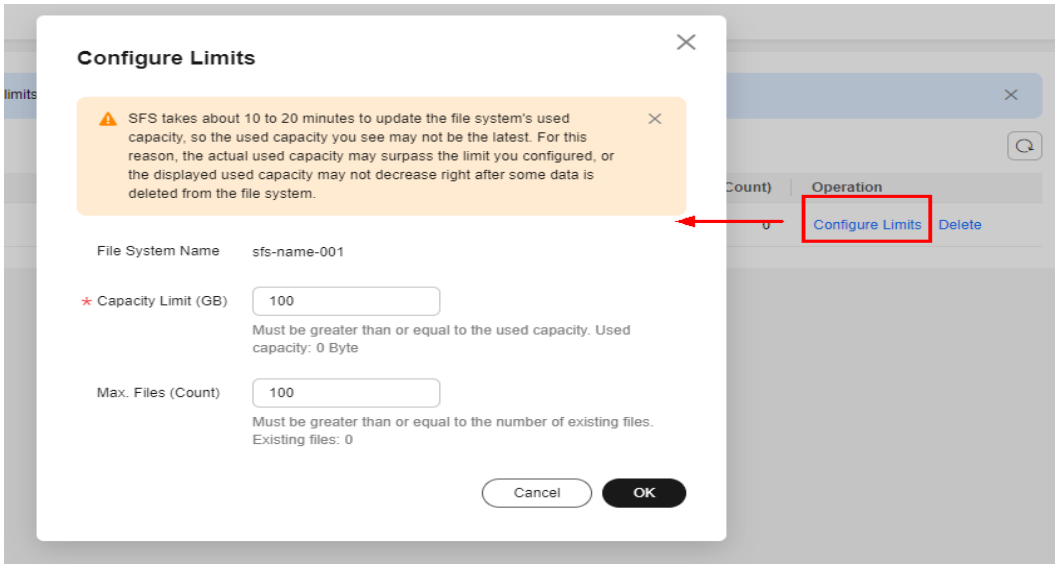

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Step 2 Modify the limits.

Capacity Limit (GB): This is a required field. Enter a value greater than 0 and at least equal to the used capacity, in GB.

Max. Files (Count): This field is optional. Enter a value greater than 0 and at least equal to the number of existing files.

Step 3 Click **OK**. View the limits details in the list.

----End

Removing Limits

Step 1 On the **Limits Management** tab, click **Delete** to open the page shown in **Figure 2-10**.

Figure 2-10 Removing limits

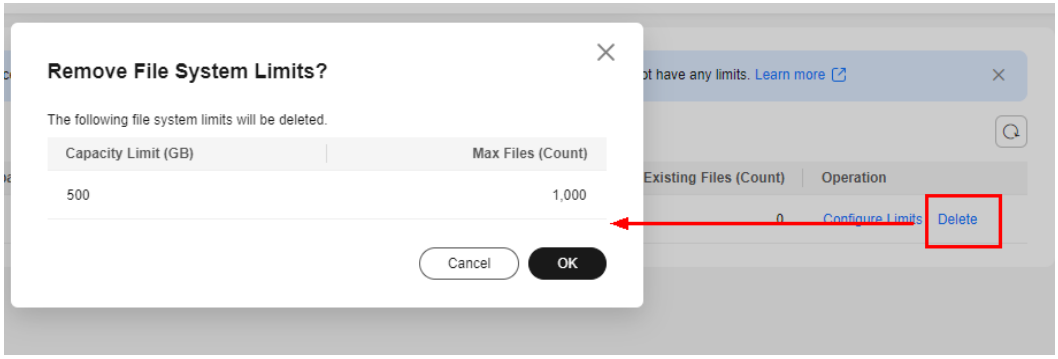

Step 2 Click **OK**.

----End

2.6 Unmounting a General-Purpose File System

If a general-purpose file system is no longer used and you want to delete it, you are advised to unmount the file system and then delete it.

Prerequisites

Stop the process and read/write operations before you unmount a general-purpose file system.

Linux OS

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

Step 2 Unmount the file system.

umount *<local-path>*

Variable *Local path* is an ECS local directory where the general-purpose file system is mounted, for example, **/local_path**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

NOTE

Before running the **umount** command, stop all read and write operations related to the general-purpose file system and exit from the local path. Otherwise, the unmounting will fail.

----End

2.7 Managing General-Purpose File System Tags

You can add tags to existing general-purpose file systems. You can also add tags when creating general-purpose file systems. For details, see **Creating a General-Purpose File System**.

Tags are used to identify and classify general-purpose file systems. For more information about tags, see **Tag Management Service Documentation**.

Constraints

- A tag is composed of a key-value pair.
  - A tag key can contain a maximum of 128 characters. It can contain letters, digits, and spaces representable in UTF-8 and special characters (_.:=+-@). It cannot start or end with a space and cannot be left empty. Tag keys starting with _sys_ are system tags, and you cannot start a tag key with _sys_.
  - A tag value can contain a maximum of 255 characters. It can contain letters, digits, and spaces representable in UTF-8 and special characters (_.:=+-@) and can be left empty. It cannot start or end with a space.
- You can add a maximum of 20 tags to a general-purpose file system.
- Tag keys of the same general-purpose file system must be unique.

Scenarios

Tags help you to identify your cloud resources. When you have many cloud resources of the same type, you can use tags to classify them by dimension (for example, by purpose, owner, or environment) for usage or cost analysis.

In SFS, you can use tags to identify and classify file systems.

If you add tags to a file system, the bills generated for this file system will contain these tags. You can activate the tags and classify bills by tag for cost analysis.

Procedure

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

Step 2 Choose **General-Purpose File System** > **File Systems**. In the file system list, find the general-purpose file system you want to add tags and click its name to go to its details page.

Step 3 In the navigation pane on the left, choose **Tag Management**, as shown in **Figure 2-11**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 2-11 Tag Management

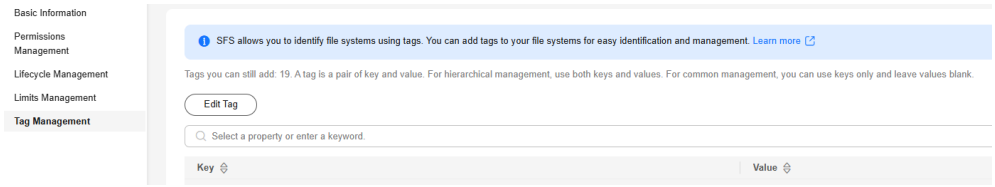

Step 4 Click **Edit Tag** to open the **Edit Tag** page.

Figure 2-12 Edit Tag

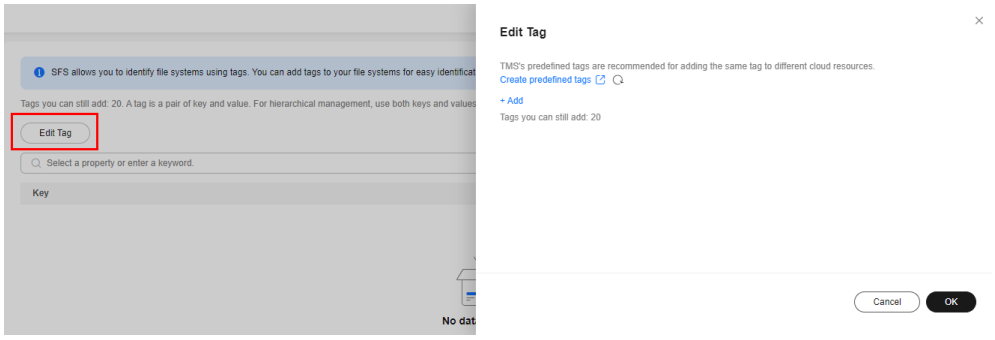

Step 5 Add tag keys and values and click **OK**.

- Tag key: This parameter is mandatory.
- Tag value: This parameter is optional.

Step 6 Return to the tag list. You can see the tags you have just added. Edit or delete the tags if needed.

----End

2.8 Deleting a General-Purpose File System

Data in a deleted general-purpose file system cannot be recovered. Ensure that files in a general-purpose file system have been properly stored or backed up before you delete the file system.

Prerequisites

You are advised to unmount the general-purpose file system before deleting it. For how to unmount general-purpose file systems, see **Unmounting a General-Purpose File System**.

Procedure

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

Step 2 In the general-purpose file system list, find the file system you want to delete and click **Delete** in the **Operation** column.

Step 3 In the displayed dialog box shown in **Figure 2-13**, confirm the information, enter **DELETE** in the text box, and click **OK**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

NOTE

A general-purpose file system can only be deleted when the file system's used capacity and the number of files in the file system are both zero.

Figure 2-13 Deleting a general-purpose file system

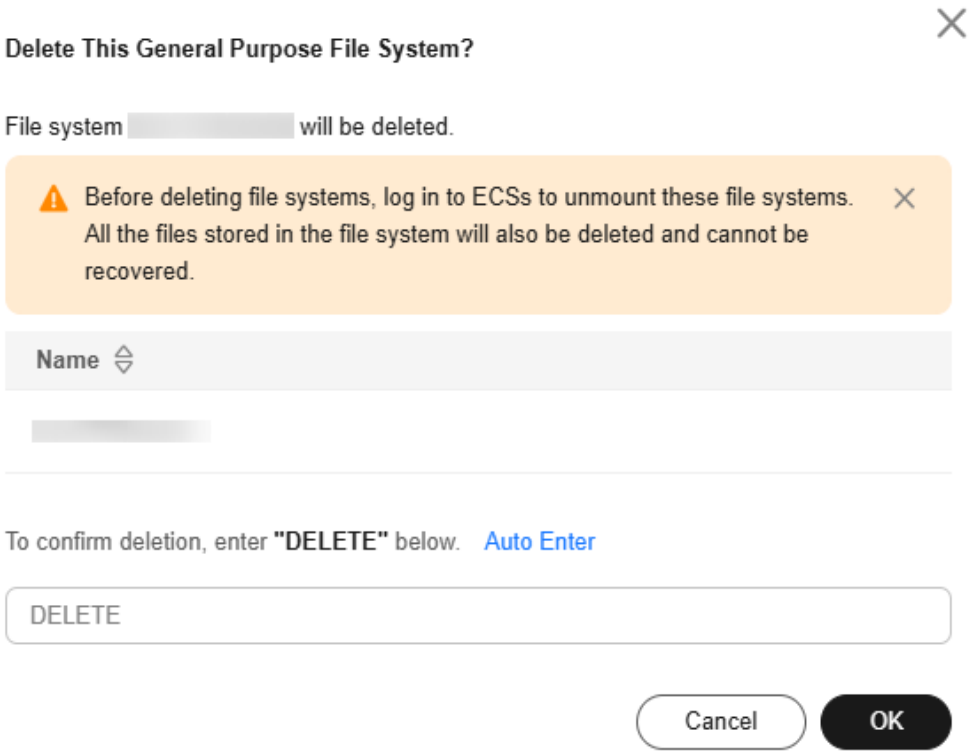

Step 4 Check whether the general-purpose file system is successfully deleted.

If the file system disappears from the file system list, the deletion is successful. If you can still see it in the list, the deletion fails. In this case, **submit a service ticket**.

----End

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

3 Network Management

3.1 Configuring Multi-VPC Access

A VPC enables you to provision logically isolated, configurable, and manageable virtual networks for ECSs, improving the security of cloud resources and simplifying network deployment. When using SFS to share files, a general-purpose file system and the cloud servers need to be in the same VPC.

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

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

Scenarios

You can configure multiple VPCs for a general-purpose file system so that cloud servers in different VPCs can share the same file system, as long as the VPCs are added as authorized VPCs of the file system or the server IP addresses are added as authorized IP addresses of the VPC.

This section describes how to access a general-purpose file system from different VPCs.

Constraints

- If a VPC added to a general-purpose file system has been deleted from the VPC console, the IP addresses or IP address ranges of this VPC can still be seen as activated in the file system's VPC list. But this VPC can no longer be used and you are advised to remove it from the list.
- Before adding an authorized VPC for a general-purpose file system, you need to create a VPC endpoint to establish communication between the compute resources and the file system.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

- You need to configure a VPC endpoint for each VPC you want to add as an authorized VPC of a general-purpose file system. Or, the file system will fail to be mounted.

Procedure

Step 1 Log in to the SFS console.

Step 2 In the general-purpose file system list, click the name of the desired file system to go to its details page.

Step 3 In the navigation pane on the left, choose **Permissions Management**.

Step 4 If no VPCs are available, create a VPC first. Click **Add Authorization Rule**. A dialog box is displayed, as shown in Figure 3-1.

Table 3-1 describes the parameters.

Figure 3-1 Add Authorization Rule

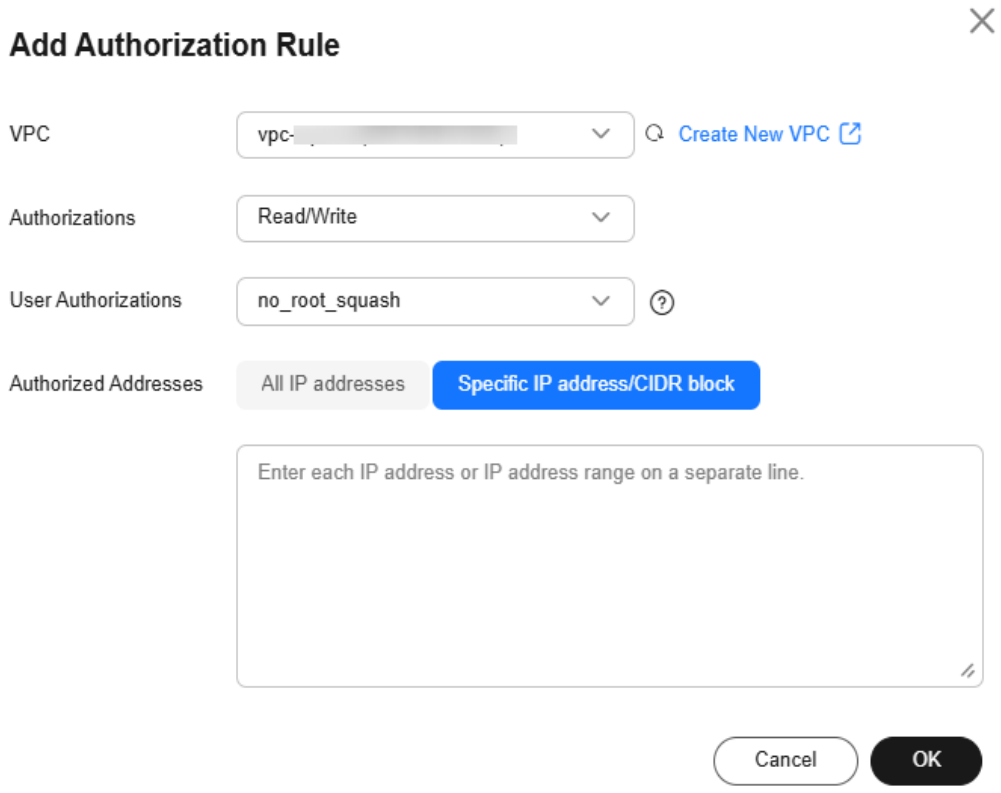

Table 3-1 Parameter description

Parameter	Description
VPC	Select the VPC you want to add, for example, **vpc-30e0**. If no VPC is available, create one.
Authorizations	You can select **Read/Write** or **Read-only**. **Read/Write** is preselected.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
User Authorizations	You can select **no_root_squash**, **root_squash**, or **all_squash**. • **no_root_squash** allows the **root** user on the client to access the general-purpose file system as **root**. • **root_squash** allows the **root** user on the client to access the general-purpose file system as the **nobody** user. • **all_squash** allows any user on the client to access the general-purpose file system as the **nobody** user, and the user can modify and delete the file system.
Authorized Addresses	You can select **All IP addresses** or **Specific IP address/CIDR block**. **All IP addresses** is preselected. **NOTE** If you select **Specific IP address/CIDR block**, you can add multiple IP addresses or CIDR blocks. Enter each one on a separate line. After the authorized addresses are added, you can click the number shown under **Authorized Addresses** in the permissions management list to check their information.

Step 5 Click **OK**. View the added VPC in the list.

Step 6 On the **VPC Endpoints** page, click **Buy VPC Endpoint**.

The **Buy VPC Endpoint** page is displayed.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 3-2 Buy VPC Endpoint

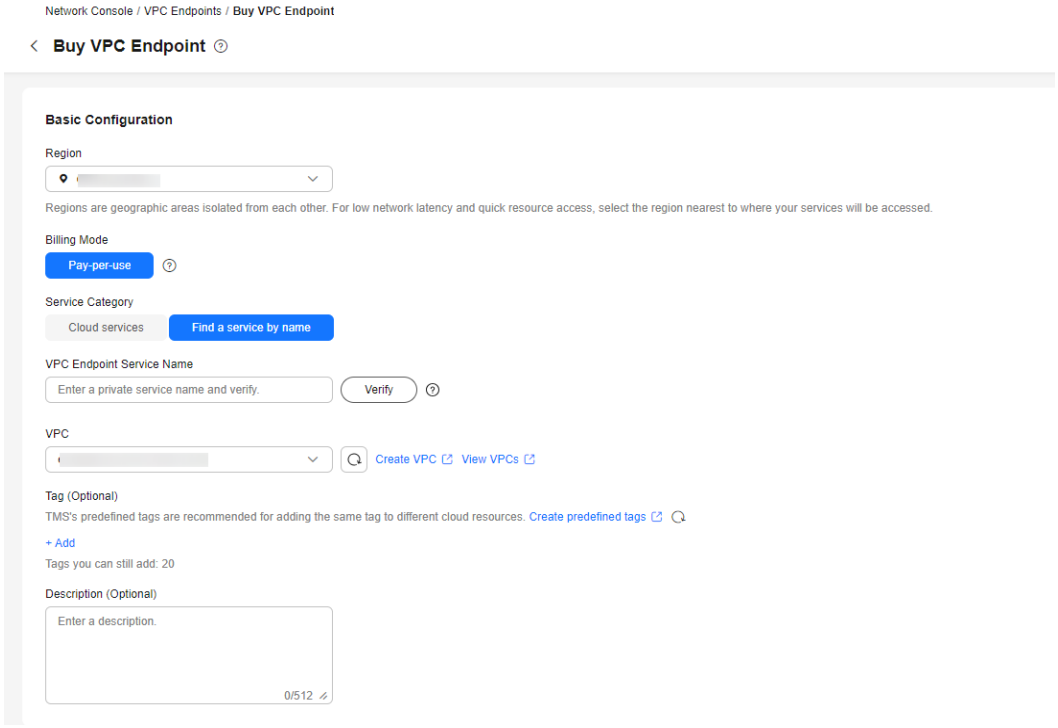

Step 7 Configure VPC endpoint parameters.

Table 3-2 VPC endpoint parameters

Parameter	Description
Region	Region where the VPC endpoint is located. Ensure that this region is the same as the one where the planned general-purpose file system resides.
Billing Mode	**Pay-per-use** is preselected by default, but you will not be billed for the endpoint purchased for general-purpose file systems.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
Service Category	Select **Find a service by name**. Enter a VPC endpoint service name based on the region selected. • If the CN North-Beijing4 region is selected, enter **cn-north-4.com.myhuaweicloud.v4.storage.lz13**. • If the CN South-Guangzhou region (AZ1) is selected, enter **cn-south-1.com.myhuaweicloud.v4.obsv2**. **NOTE** General-purpose file systems created in AZ1 of the CN South-Guangzhou region cannot be mounted to containers. • If the CN South-Guangzhou region (AZ6) is selected, enter **cn-south-1.com.myhuaweicloud.v4.obsv2.storage.lz06**. • If the CN East-Shanghai1 region is selected, enter **cn-east-3.com.myhuaweicloud.v4.storage.lz07**. • If the CN Southwest-Guiyang1 region is selected, enter **com.myhuaweicloud.cn-southwest-2.ipv4.sfs**. • If the CN-Hong Kong region is selected, enter **ap-southeast-1.com.myhuaweicloud.v4.obsv2.storage.lz005**. After entering the service name, click **Verify**. If **Service name found** is displayed, proceed with subsequent steps. If **Service name not found** is displayed, check whether the entered service name is correct. If the problem persists, **submit a service ticket**.
VPC	VPC where the planned general-purpose file system and ECSs reside.
Tag	Optional VPC endpoint tags. Each tag consists of a key and a value. You can add a maximum of 10 tags to a VPC endpoint. Tag keys and values must meet the requirements listed in **Table 3-3**. **NOTE** If you have created a predefined tag in TMS, you can select the corresponding tag key and value. For details about predefined tags, see **Predefined Tags**.

Table 3-3 describes the tag parameters.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Table 3-3 Tag parameters

Parameter	Description	Example Value
Tag key	Each tag has a unique key. You can customize a key or select the key of an existing tag created in TMS. A tag key: • Cannot be an empty string. • Must be unique for each resource. • Can contain a maximum of 128 characters. • Can contain letters, digits, spaces, and the following characters: _.:=+-@; cannot start or end with a space, or start with _sys_.	Key_0001
Tag value	A tag value can be repetitive or left blank. A tag value: • Can be an empty string. • Can contain a maximum of 255 characters. • Can contain letters, digits, spaces, and the following characters _.:/=+-@; cannot start or end with a space.	Value_0001

Step 8 Click **Next**.

- If you do not need to modify the configuration, click **Submit**.
- If you need to modify the configuration, click **Previous**, modify the configuration as needed, and then click **Submit**.

Step 9 Go back to the VPC endpoint list and check whether the status of the VPC endpoint changes to **Accepted**. If so, the VPC endpoint has been connected to the VPC endpoint service.

----End

Verification

After a VPC is added to the general-purpose file system, mount the file system to a server in the added VPC. If the mount is successful and the server can access the file system, the multi-VPC access configuration is successful.

3.2 Configuring a DNS Server for Domain Name Resolution

A DNS server is used to resolve domain names of general-purpose file systems. For details about DNS server IP addresses, see **What Are Huawei Cloud Private DNS Server Addresses?**

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Scenarios

By default, the IP address of the DNS server is automatically configured on ECSs when ECSs are created. No manual configuration is needed except when the resolution fails due to a change in the DNS server IP address.

Procedure (Linux)

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

Step 2 Run **vi /etc/resolv.conf** to edit the **/etc/resolv.conf** file. Add the DNS server IP address above the existing nameserver information.

Figure 3-3 Configuring DNS

The format is as follows:

```
nameserver 100.125.1.250
```

Step 3 Press **Esc**, enter **:wq**, and press **Enter** to save and exit.

Step 4 Check whether the IP address is successfully added.

cat /etc/resolv.conf

Step 5 Check whether the domain name of the general-purpose file system can be resolved.

nslookup *<domain-name-of-the-general-purpose-file-system>*

NOTE

Obtain the domain name of a general-purpose file system from its mount point. For example, if the mount point of a general-purpose file system is *xxx*:**/sfs-name-001**, *xxx* is the file system domain name, and **sfs-name-001** is the file system name.

Step 6 (Optional) If DHCP is configured for the ECS, 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. Lock the file.

   chattr +i /etc/resolv.conf

   NOTE

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

2. Check whether the file is locked.

   lsattr /etc/resolv.conf

   If the information shown in **Figure 3-4** is displayed, the file is locked.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 3-4 File locked

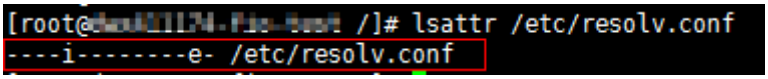

----End

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

4 Data Management

4.1 Configuring a Lifecycle Rule

Infrequent Access Storage

General-purpose file systems allow you to configure lifecycle rules to transition inactive files to infrequent access storage to reduce costs.

Infrequent access storage has the following advantages:

- Simple configuration (no need to compile scripts or migrate data)

  All you need to do is configure lifecycle rules, then general-purpose file systems will automatically transition files that meet the rules to infrequent access storage. No complex or high-risk operation is involved.

- Low costs

  Infrequent access storage saves more money than standard storage.

  NOTE

  For details about the billing of infrequent access storage, see **Billed Items**.

- Normal data access after transition

  After files are transitioned to infrequent access storage, the content and structure of general-purpose file systems remain unchanged and applications can access such files normally. You do not need to modify applications or suspend services.

Configuring a Lifecycle Rule

You can configure lifecycle rules for a general-purpose file system or a specific directory in a file system. Files that meet the rules will be transitioned from standard storage to infrequent access storage.

You can configure up to 20 lifecycle rules for a general-purpose file system.

You can copy, enable, disable, modify, and delete lifecycle rules. Perform the following steps to create a rule:

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

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

Step 2 In the general-purpose file system list, click the name of the desired file system to go to its details page.

Step 3 On the **Lifecycle Management** tab, click **Create Rule**, as shown in **Figure 4-2**.

Figure 4-1 Lifecycle Management

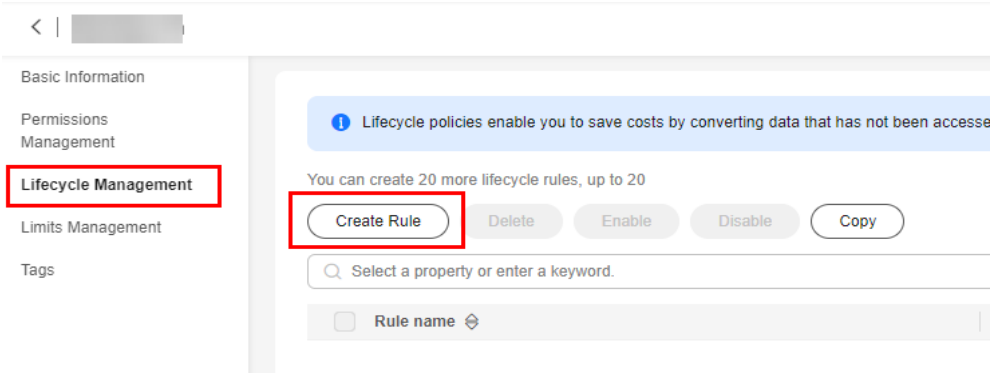

Figure 4-2 Create Lifecycle Rule

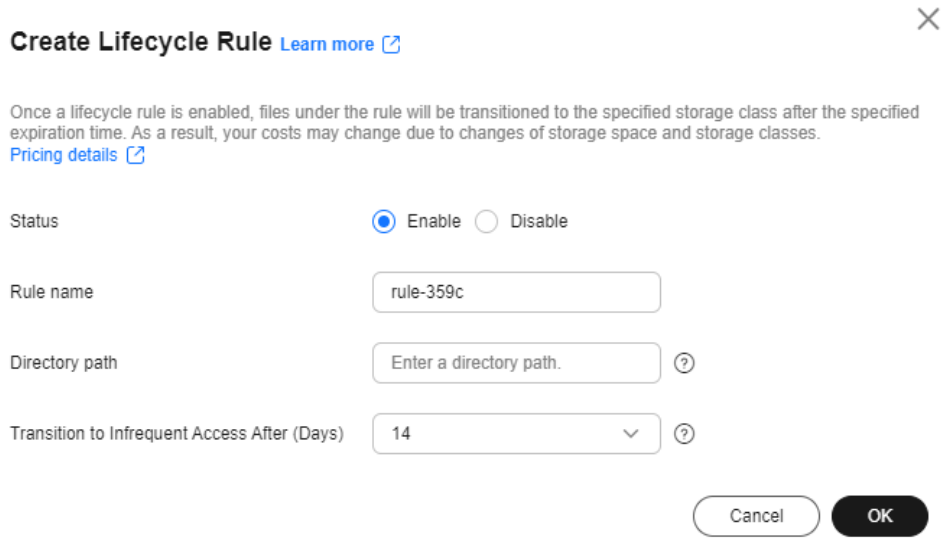

Step 4 Configure rule parameters.

- **Status**: Selecting **Enable** enables this rule after it is created.
- **Rule name**: Enter a rule name, which can contain only letters, digits, periods (.), underscores (_), or hyphens (-).
- **Directory path**: Enter the path of a directory on which the rule will be applied. If not specified, the rule will be applied to the entire file system. The path cannot start with a slash (/), contain two adjacent slashes (//), or contain the following special characters: \:*?"<>|
- **Transitioned to Infrequent Access After**: defines the number of days that must pass for files to transition to infrequent access storage after their last access. There are four options: 14 days, 30 days, 60 days, and 90 days. When the specified directory is not accessed for the specified number of days, files in this directory will be transitioned to infrequent access storage.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Step 5 Click **OK**.

----End

Replicating a Lifecycle Rule

In addition to creating lifecycle rules, you can replicate rules from other general-purpose file systems. Perform the following steps to replicate a rule:

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

Step 2 In the general-purpose file system list, click the name of the desired file system to go to its details page.

Step 3 On the **Lifecycle Management** tab, click **Copy** to open the page shown in **Figure 4-3**.

Figure 4-3 Replicate Lifecycle Rule

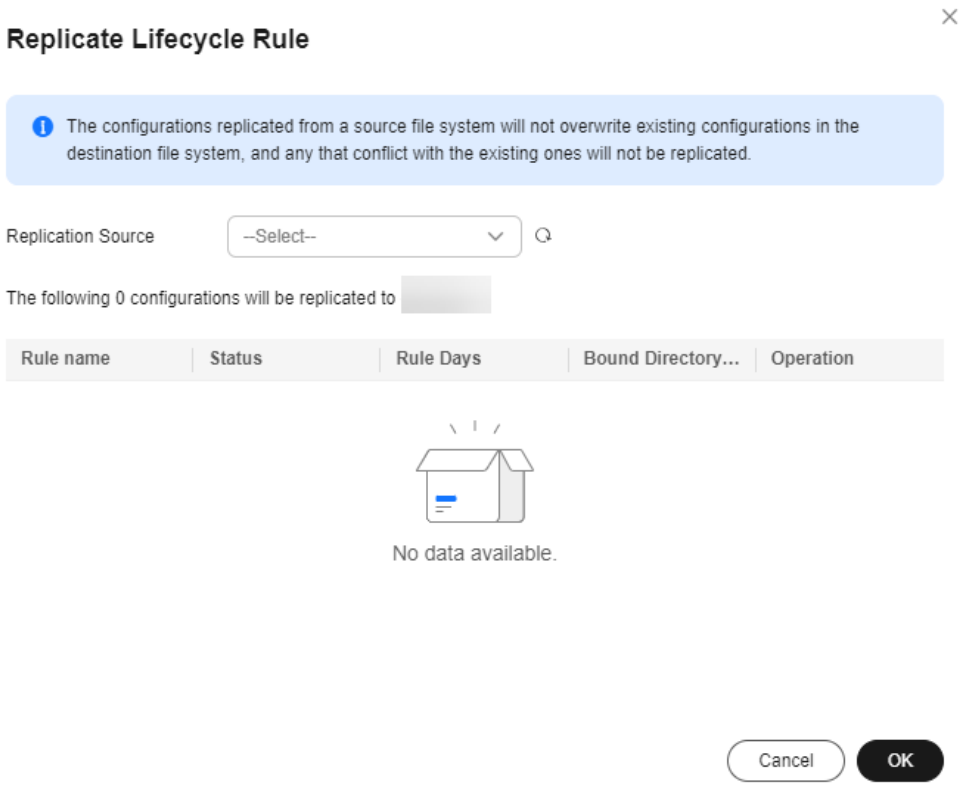

Step 4 Select a replication source, which is the general-purpose file system whose lifecycle rules you want to replicate.

NOTE

- Lifecycle rules replicated from a source file system will not overwrite existing rules in the destination file system, and any rules that conflict with the existing ones will not be replicated.
- You can remove rules that you do not want to replicate.

Step 5 Click **OK**.

----End

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Other Operations

- Modifying a lifecycle rule: Locate the rule you want to modify and click **Edit** in the **Operation** column. For details about the rule parameters, see **Step 4**.
- Enabling or disabling a lifecycle rule: Locate the target rule and click **Enable** or **Disable** in the **Operation** column.

  NOTE

  To batch enable lifecycle rules, ensure that all target rules are disabled. To batch disable lifecycle rules, ensure that all target rules are enabled.

  Figure 4-4 Disable This Lifecycle Rule

  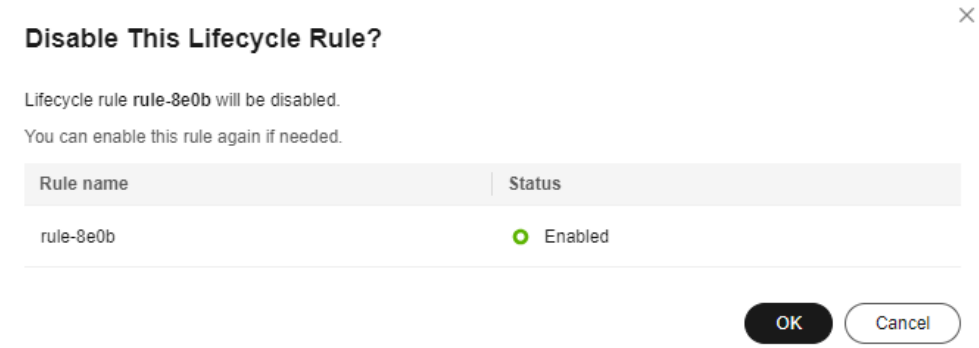

  Figure 4-5 Enable This Lifecycle Rule

  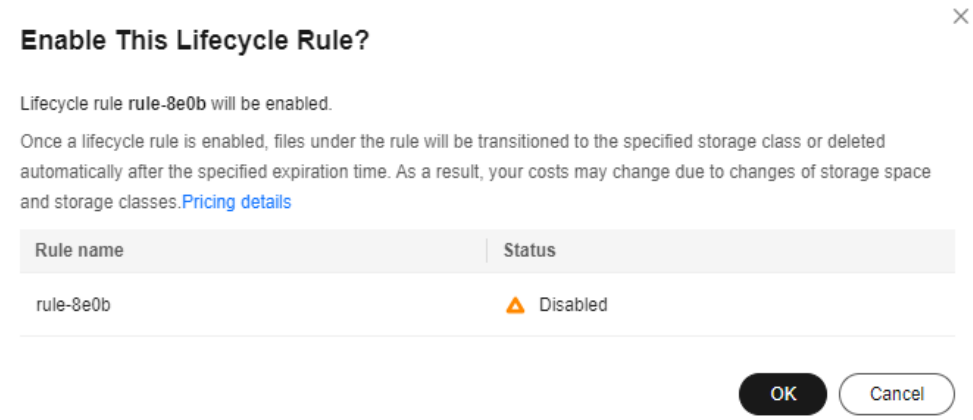

- Batch enabling or disabling lifecycle rules: Select the desired rules and click **Enable** or **Disable** above the rule list to perform the corresponding operation.
- Deleting a lifecycle rule: Locate the target rule and click **Delete** in the **Operation** column. Or, click the checkbox in front of the rule name and click **Delete** above the rule list. You can also delete rules in a batch.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 4-6 Delete This Lifecycle Rule

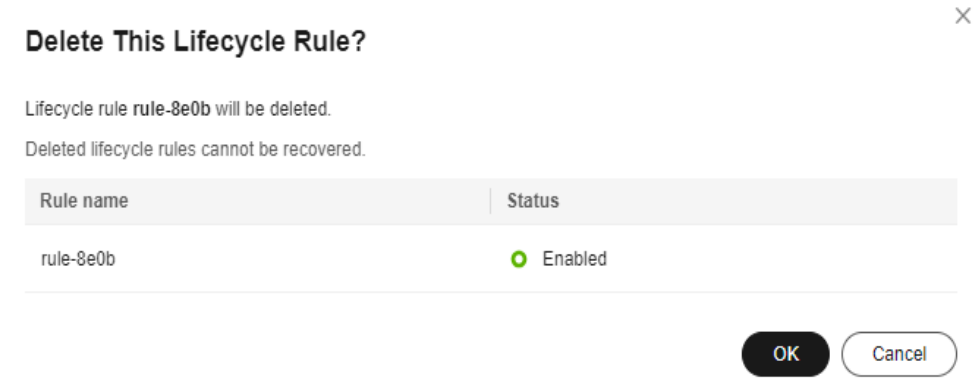

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

5 Monitoring and Auditing

5.1 Monitoring General-Purpose File Systems Using Cloud Eye

5.1.1 SFS Metrics

Function

This section describes the SFS metrics reported to Cloud Eye as well as their namespaces and dimensions. You can use the Cloud Eye console or **APIs** provided by Cloud Eye to query the SFS metrics and alarms.

NOTE

Cloud Eye can monitor dimensions nested to a maximum depth of four levels (levels 0 to 3). 3 is the deepest level.

Namespace

SYS.SFS

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Metrics

Table 5-1 General-Purpose File System metrics

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
capacity_standard	Capacity Class Storage Usage	Storage space used by standard storage	≥ 0	bytes	1024 (IEC)	bucket_name	30 minutes
capacity_infrequent access	Infrequent-Access Class Storage Usage	Storage space used by infrequent access storage	≥ 0	bytes	1024 (IEC)	bucket_name	30 minutes
read_bandwidth	File system read bandwidth	Read bandwidth of a file system within a monitoring period	≥ 0	bytes/s	1024 (IEC)	share_id	4 minutes
write_bandwidth	File system write bandwidth	Write bandwidth of a file system within a monitoring period	≥ 0	bytes/s	1024 (IEC)	share_id	4 minutes
read_tps	File system read TPS	Number of read operations of a file system within a monitoring period	≥ 0	count	N/A	bucket_name	4 minutes

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
write_tps	File system write TPS	Number of write operations of a file system within a monitoring period	≥ 0	count	N/A	bucket_name	4 minutes

NOTE

General-purpose file system capacities cannot be monitored using the **used_capacity** metric.

Monitoring data of general-purpose file systems are only displayed when there are service accesses.

Dimension

Key	Value
share_id	General-purpose file system
bucket_name	General-purpose file system

Viewing Monitoring Statistics

Step 1 **Log in to the Cloud Eye console**.

Step 2 Choose **Cloud Service Monitoring** > **Scalable File Service SFS**.

Step 3 Click the name in the **ID** column to go to the resource instance page.

Step 4 View the SFS file system monitoring data by metric or monitored duration.

Figure 5-1 shows the monitoring graphs. For more information about Cloud Eye, see the *Cloud Eye User Guide*.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 5-1 Monitoring graphs of a general-purpose file system

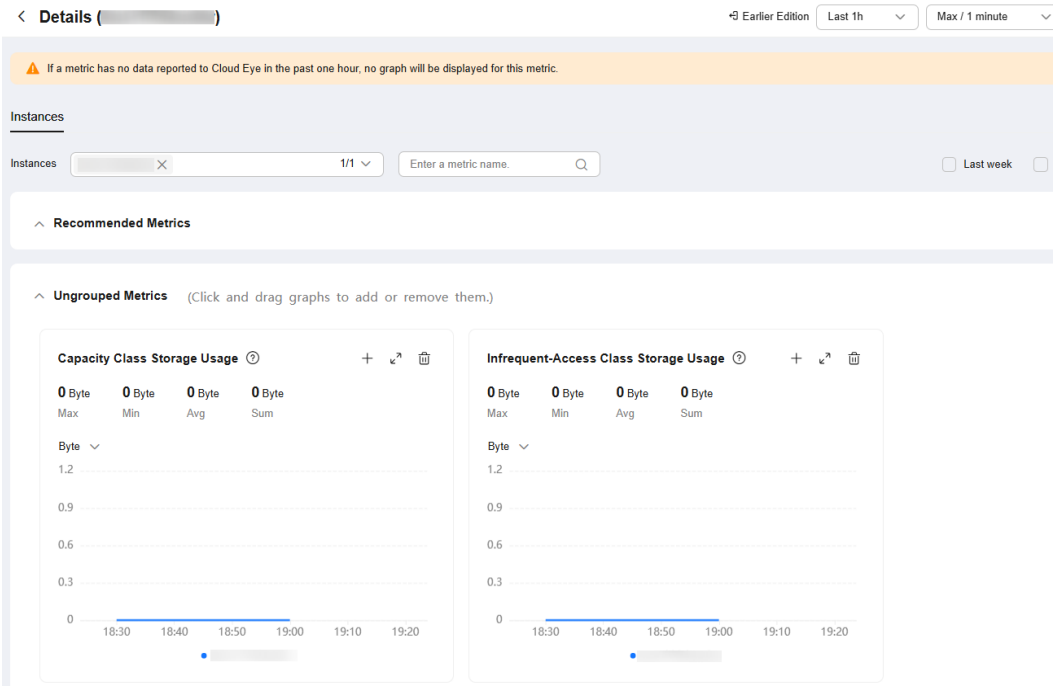

----End

5.1.2 Creating an Alarm Rule

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 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 need to enable **Alarm Notification** 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 Cloud Eye console**.
2. In the navigation pane on the left, choose **Alarm Management** > **Alarm Rules**.
3. Click **Create Alarm Rule** in the upper right corner.
4. On the **Create Alarm Rule** page, configure the parameters.
   a. Configure the basic information for the alarm rule.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Figure 5-2 Basic information

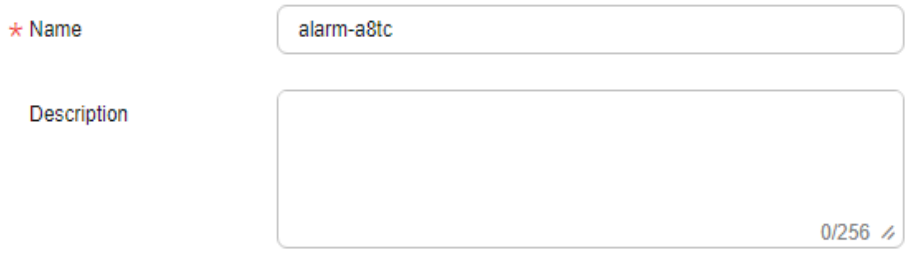

Table 5-2 Parameters for configuring basic information

Parameter	Description	Example Value
Name	Name of the alarm rule. Cloud Eye will generate a random name, but you can modify it.	alarm-b6al
Description	Description of the alarm rule. This parameter is optional.	-

b. Select monitored objects and configure alarm parameters.

Figure 5-3 Configuring alarm rule parameters

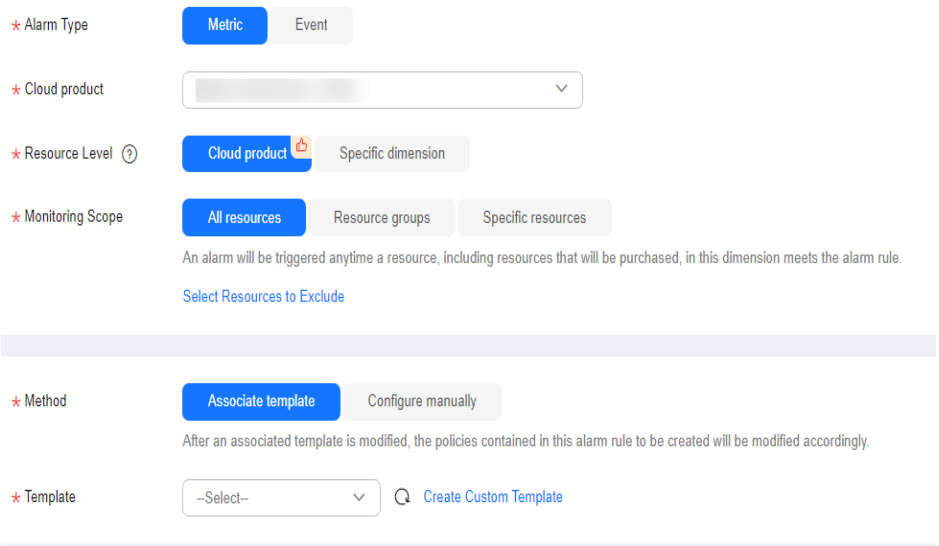

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Table 5-3 Configuring alarm rule parameters

Parameter	Description	Example Value
Alarm Type	Alarm type to which the alarm rule will apply. The type can be **Metric** or **Event**.	Metric
Cloud Product	Select the cloud product you want to monitor. This parameter is only available if you select **Metric** for **Alarm Type**.	Scalable File Service - General Purpose File System
Resource Level	Select the resource level of the alarm rule. This parameter is only available if you select **Metric** for **Alarm Type**. You can select **Cloud product** (recommended) or **Specific dimension**.	Cloud product
Monitoring Scope	Select the resource scope that the alarm rule will apply to. This parameter is only available if you select **Metric** for **Alarm Type**. You can select **All resources**, **Resource groups**, or **Specific resources**. **NOTE** • **All resources**: An alarm will be triggered if any resource of the current cloud product meets the alarm policy. To exclude resources that do not require monitoring, click **Select Resources to Exclude** to select resources. • **Resource groups**: An alarm will be triggered if any resource in the resource group meets the alarm policy. To exclude resources that do not require monitoring, click **Select Resources to Exclude** to select resources. • **Specific resources**: Click **Select Specific Resources** to select resources.	All resources
Group	This parameter is only available if you select **Metric** for **Alarm Type** and **Resource groups** for **Monitoring Scope**.	-
Instance	This parameter is only available if you select **Metric** for **Alarm Type** and **Specific resources** for **Monitoring Scope**.	-
Event Type	Select either **System event** or **Custom event**. This parameter is only available if you select **Event** for **Alarm Type**.	System event

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description	Example Value
Event Source	This parameter is only available if you select **Event** for **Alarm Type**. • If you select **System event** for **Event Type**, select a cloud service from which the event originates. • If you select **Custom event** for **Event Type**, specify the event source. Ensure that the event source is the same as that of the reported fields and is written in the service.item format.	-
Method	• **Configure manually**: If you select **Event** for **Alarm Type** and **Custom event** for **Event Type**, only **Configure manually** can be set for **Method**. • **Associate template**: If you select this option, any modification made to the template will also be synchronized to the policies of the alarm rule that the template is associated. **NOTE** • When **Resource Level** is set to **Cloud product**, only modifications made to the policies of the specified cloud product in the associated template will be automatically synchronized. • When **Resource Level** is set to **Specific dimension**, only modifications made to the policies of the specified dimension in the associated template will be automatically synchronized.	Configure manually
Template	If you select **Metric** for **Alarm Type** and **Associate template** for **Method**, or select **Event** for **Alarm Type**, **System event** for **Event Type**, and **Associate template** for **Method**, you need to select a template. You can select a default template or create a custom template.	-

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description	Example Value
Alarm Policy	If you select **Event** for **Alarm Type** and **Custom event** for **Event Type**, you need to set **Alarm Policy**. If you select **Custom event** for **Event Type**, as long as an event occurs, an alarm will be triggered. For example, if the running status is abnormal, an alarm will be triggered. **NOTE** You can add up to 50 alarm policies to an alarm rule. If any of these alarm policies are met, an alarm will be triggered.	-
Alarm Severity	Severity of an alarm. Valid values are **Critical**, **Major**, **Minor**, and **Warning**.	Major

c. Configure alarm notifications.

Figure 5-4 Configuring alarm notifications

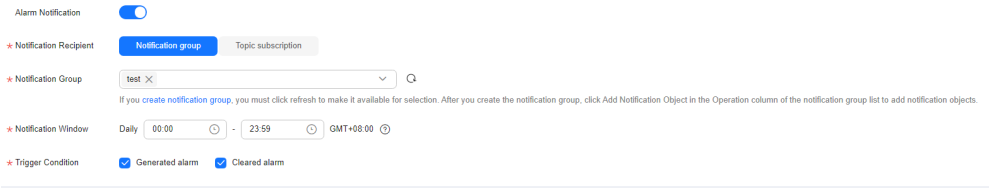

Table 5-4 Parameters for configuring alarm notifications

Parameter	Description
Alarm Notifications	Whether to send notifications to users via SMS, email, voice notification, HTTP, HTTPS, FunctionGraph (function), FunctionGraph (workflow), WeCom chatbot, DingTalk chatbot, Lark chatbot, or WeLink chatbot.
Notified By	You can select **Notification policies**, **Notification groups** or **Topic subscriptions**. • **Notification policies**: Flexible alarm notifications are sent by severity. There are many notification channels. • **Notification groups**: Notification templates are configured on the Cloud Eye console. • **Topic subscriptions**: Notification templates are configured on the SMN console.
Notification Policies	This parameter is only available if you select **Notification policies** for **Notified By**. Select one or more notification policies. You can specify the notification group, window, template, and other parameters in a notification policy.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Parameter	Description
Notification Group	This parameter is only available if you select **Notification groups** for **Notified By**. Select the notification groups to which alarm notifications will be sent.
Recipient	Recipient of alarm notifications. You can select the account contact or a topic name. This parameter is only available if alarm notification is set to **Topic subscriptions**. • **Account contact** is the mobile number and email address of the registered account. • **Topic** is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one and add subscriptions to it. For details, see **Creating a Topic** and **Adding Subscriptions**.
Notification Template	This parameter is only available if you select **Notification groups** or **Topic subscriptions** for **Notified By**. You can select an existing template or create a new one.
Notification Window	This parameter is only available if you select **Notification groups** or **Topic subscriptions** for alarm notifications. Cloud Eye sends notifications only within the notification window you specified. If **Notification Window** is set to **08:00-20:00**, Cloud Eye sends notifications only from 08:00 to 20:00.
Trigger Condition	This parameter is only available if you select **Notification groups** or **Topic subscriptions** for alarm notifications. You can select either **Generated alarm** or **Cleared alarm**, or both. **NOTE** When the alarm type is **Event**, you can only select **Generated alarm** for **Trigger Condition**.

d. Configure **Enterprise Project** and **Tags**.

Figure 5-5 Advanced Settings

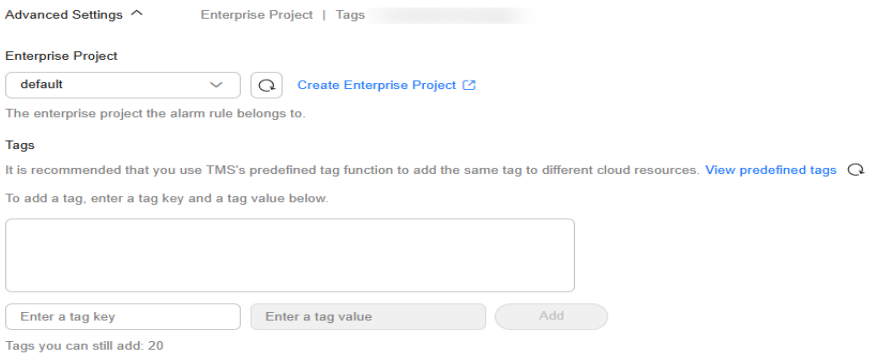

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Table 5-5 Parameters for configuring advanced settings

Parameter	Description
Enterprise Project	Enterprise project that the alarm rule belongs to. Only users who have the permissions for the enterprise project can manage this alarm rule. For how to create an enterprise project, see **Creating an Enterprise Project**.
Tags	Tags are key-value pairs. You can tag cloud resources to easily categorize and search for them. You are advised to create predefined tags in TMS. For how to create predefined tags, see **Creating Predefined Tags**. If your organization has enabled tag policies and has a Cloud Eye-related tag policy attached, you must comply with the tag policy rules when creating alarm rules, otherwise alarm rules may fail to be created. Contact the organization administrator to learn more about tag policies. • A key can contain up to 128 characters, and a value can contain up to 225 characters. • You can add up to 20 tags.

e. Click **Create**.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

6 Typical Applications

6.1 High-performance Computing

Context

A high-performance computing (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 general-purpose file systems:

- Industrial design: In automobile manufacturing, CAE and CAD simulation software is widely used. When the software is operating, compute nodes need to communicate with each other closely, which requires general-purpose file systems that can provide high bandwidth and low latency.
- Bioscience: General-purpose file systems 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 general-purpose file systems to provide large memory and high bandwidth.
- Image rendering: Image processing, 3D rendering, and frequent processing of small files require high read/write performance, large capacity, and high bandwidth of general-purpose file systems.
- Heterogeneous computing: Compute elements may have different instruction set architectures, requiring general-purpose file systems to provide high bandwidth and low latency.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

SFS is a shared storage service based on general-purpose 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. 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 cost and O&M cost are greatly reduced, the service rollout time is shortened, and efficiency is boosted.

Configuration Process

1. Prepare the files of DNA sequencing to be uploaded.
2. **Log in to the SFS console** and create a general-purpose file system to store the prepared files.
3. Log in to the cloud servers that function as the head node and compute node, and mount the general-purpose file system on them.
4. On the head node, upload the files to the general-purpose file system.
5. On the compute node, edit the files.

Prerequisites

- A VPC has been created.
- Cloud servers that function as the head node and compute node have been created and are in the created VPC.
- SFS has been enabled.

Example Configuration

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

Step 2 In the upper right corner of the page, click **Create File System**.

Step 3 On the page for creating a general-purpose file system, configure parameters as instructed.

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

For how to mount a file system to Linux ECSs, see **Mounting a General-Purpose File System to Linux ECSs**.

Step 5 Log in to the head node and upload the files to the general-purpose file system.

Step 6 Start gene sequencing. The compute node obtains the gene sequencing file from the mounted general-purpose file system for calculation.

----End

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

6.2 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 general-purpose file systems:

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

SFS is a shared storage service based on general-purpose 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.

For example, 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. Raw materials are then uploaded, through the upload workstations, to the general-purpose file system. In this way, the editing workstations concurrently edit the materials in the file system.

Configuration Process

1. Prepare the material files to be uploaded.
2. **Log in to the SFS console** and create a general-purpose file system to store the prepared files.
3. Log in to the ECSs that function as upload and editing workstations, and mount the general-purpose file system.
4. On the upload workstation, upload the material files to the general-purpose file system.
5. On the editing workstation, edit the material files.

Prerequisites

- A VPC has been created.
- ECSs that function as upload and editing workstations have been created and are in the created VPC.
- SFS has been enabled.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Example Configuration

Step 1 Log in to the SFS console.

Step 2 In the upper right corner of the page, click **Create File System**.

Step 3 On the page for creating a general-purpose file system, configure parameters as instructed.

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

For how to mount a file system to Linux ECSs, see **Mounting a General-Purpose File System to Linux ECSs**.

Step 5 Log in to the upload workstation and upload the material files to the general-purpose file system.

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

----End

6.3 Log Printing

Context

General-Purpose File System 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:

- Sharing: A general-purpose file system is mounted to multiple service hosts and logs are printed concurrently.
- Large file size and small I/Os: The size of a single log file is large, but the I/O of each log write is small.
- Intensive write I/Os: Most service I/Os are write I/Os of small blocks.

Configuration Process

1. Log in to the SFS console and create a general-purpose file system to store the prepared files.
2. Log in to the cloud server that functions as the compute node and mount the file system.
3. Configure the file system path as the log directory. It is recommended that each host use different log files.
4. Start applications.

Prerequisites

- A VPC has been created.
- Cloud servers that function as the head node and compute node have been created and are in the created VPC.
- SFS has been enabled.

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.

Example Configuration

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

Step 2 In the upper right corner of the page, click **Create File System**.

Step 3 On the page for creating a file system, configure parameters as instructed.

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

For how to mount a file system to Linux ECSs, see **Mounting a General-Purpose File System to Linux ECSs**.

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

Step 6 Start applications.

----End

 Copyright © Huawei Cloud Computing Technologies Co., Ltd.