Alibaba Cloud
Network Attached Storage

User Guide

Issue: 20200709

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
6. Please contact Alibaba Cloud directly if you discover any errors in this document.

Document conventions

Style	Description	Example
	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 **Danger:** Resetting will result in the loss of user configuration data.
	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 **Warning:** Restarting will cause business interruption. About 10 minutes are required to restart an instance.
	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	 **Notice:** If the weight is set to 0, the server no longer receives new requests.
	A note indicates supplemental instructions, best practices, tips, and other content.	 **Note:** You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings** > **Network** > **Set network type**.
Bold	Bold formatting is used for buttons, menus, page names, and other UI elements.	Click **OK**.
Courier font	Courier font is used for commands.	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`

Style	Description	Example
{} or {a|b}	This format is used for a required value, where only one item can be selected.	switch {active|stand}

Contents

1 Manage permissions

1.1 Use RAM to manage users' access to resources

You can create RAM user accounts to manage users and their access to Aspara File Storage NAS resources.

Context

You can create and manage multiple RAM user accounts with a single Alibaba Cloud account. You can grant different permissions for each RAM user account. This allows each RAM user account to have different access permissions on Alibaba Cloud resources. With RAM, you do not need to share an AccessKey with another account. You can assign minimal permissions to each user to reduce your data security risks.

Create a RAM user

1. Log on to the RAM console by using an Alibaba Cloud account.
2. In the left-side navigation pane, choose **Identities** > **Users**, and click **Create User**.
3. Configure the user account information.
4. Select **Console Password Logon** and **Programmatic Access** under Access Mode.
5. Select **Custom Logon Password** under Console Password, enter a password, and select **Required at Next Logon** under Password Reset.
6. Optional. Select Required to Enable MFA under Multi-factor Authentication and click **OK**.
7. Save the new account, logon password, AccessKey ID, and AccessKey secret.

> **Note:**
> We recommend that you save the AccessKey information in a timely manner and keep all details strictly confidential.

Create a user group

If you attempt to create multiple RAM user accounts, you can group RAM user accounts with identical responsibilities into the same group and authorize the group. This makes it easier to manage users and their permissions.

1. Log on to the RAM console by using an Alibaba Cloud account.
2. In the left-side navigation pane, choose **Identities** > **Groups**, and click **Create Group**.

3. Enter a group name and display name, and click **OK**.

Grant permissions to a RAM user or group

By default, a new RAM user or group does not have any permissions. You need to grant permissions to the RAM user or group to ensure that the user or group can access resources by using the console or API operations. The following steps take a RAM user account as an example to grant permissions.

Alibaba Cloud provides two system polices for you to manage access to Aspara File Storage NAS resources. You can grant one of the following policies to a RAM user account as required.

- AliyunNASFullAccess: This policy grants a RAM user account full access to Aspara File Storage NAS resources.
- AliyunNASReadOnlyAccess: This policy grants a RAM user account read-only access to Aspara File Storage NAS resources.

Note:
If these two system policies cannot meet your business requirements, you can create custom policies. For more information, see Create a custom policy.

1. On the **Users** page, select a RAM user account to be authorized, and click **Add Permissions**.
2. In the **Add Permissions** dialog box, select the required NAS permission and grant the permission to the RAM user account.

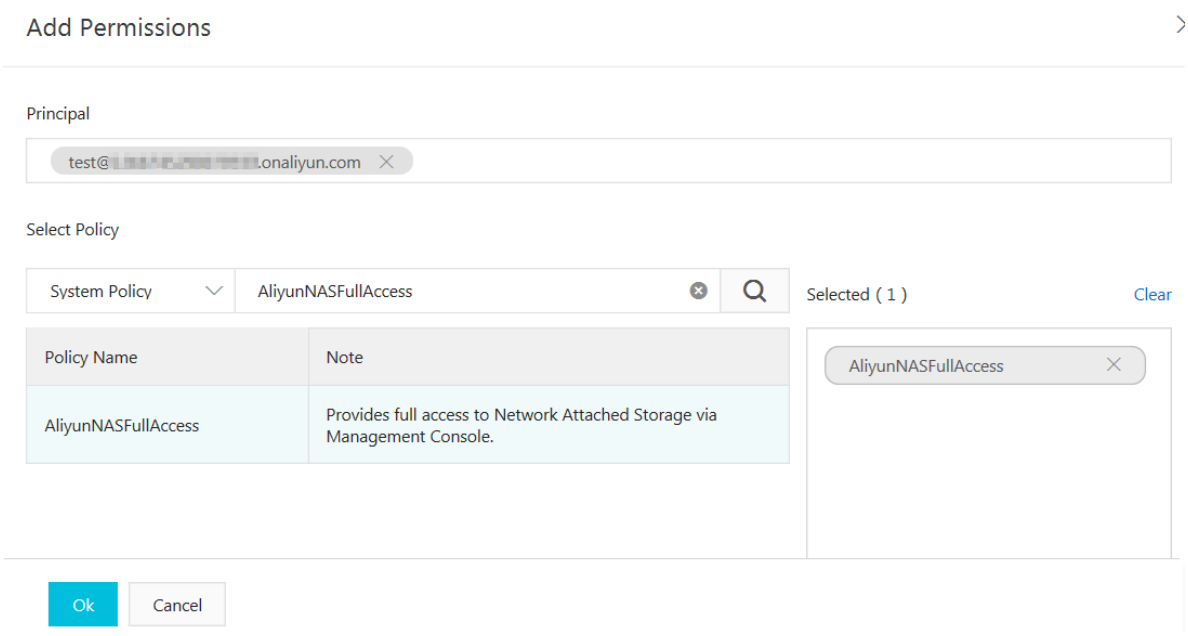

1.2 Create a custom policy

This topic describes how to create a custom policy and grant the policy to a RAM user account. Custom policies can better satisfy your specific requirements and help better manage access to your Apsara File Storage NAS resources.

Procedure

1. Log on to the RAM console by using an Alibaba Cloud account.
2. In the left-side navigation pane, select **Policies**, click **Create Policy**, and follow the instructions to create a policy. The following takes the NASReadOnlyAccess policy as an example. This policy allows read-only access to all Aspara File Storage NAS resources. For more information about the script syntax, see #unique_6.

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "nas:Describe*",
      "Resource": "*"
    }
  ],
  "Version": "1"
}
```

The following table lists the API operations that you can call to manage Apsara File Storage NAS file systems.

Operation	Description
DescribeFileSystems	Lists all file systems.
DescribeMountTargets	Lists all mount targets of a file system.
DescribeAccessGroup	Lists all permission groups.
DescribeAccessRule	Lists all rules added to a permission group .
CreateMountTarget	Adds a mount target for a file system.
CreateAccessGroup	Creates a permission group.
CreateAccessRule	Adds a rule to a permission group.
DeleteFileSystem	Deletes a file system.
DeleteMountTarget	Deletes a mount target.
DeleteAccessGroup	Deletes a permission group.

Operation	Description
DeleteAccessRule	Deletes a rule that is added to a permission group.
ModifyMountTargetStatus	Enables or disables a mount target.
ModifyMountTargetAccessGroup	Modifies the permission group of a mount target.
ModifyAccessGroup	Modifies a permission group.
ModifyAccessRule	Modifies a rule added to a permission group.

The following table shows the accessible Apsara File Storage NAS resources.

Resource	Description
*	All Apsara File Storage NAS resources

3. After the policy is created, go to the **Users** page.
4. Select a RAM user account to be authorized, click **Add Permissions**, select the required NAS permission, and grant the permission to the RAM user account.

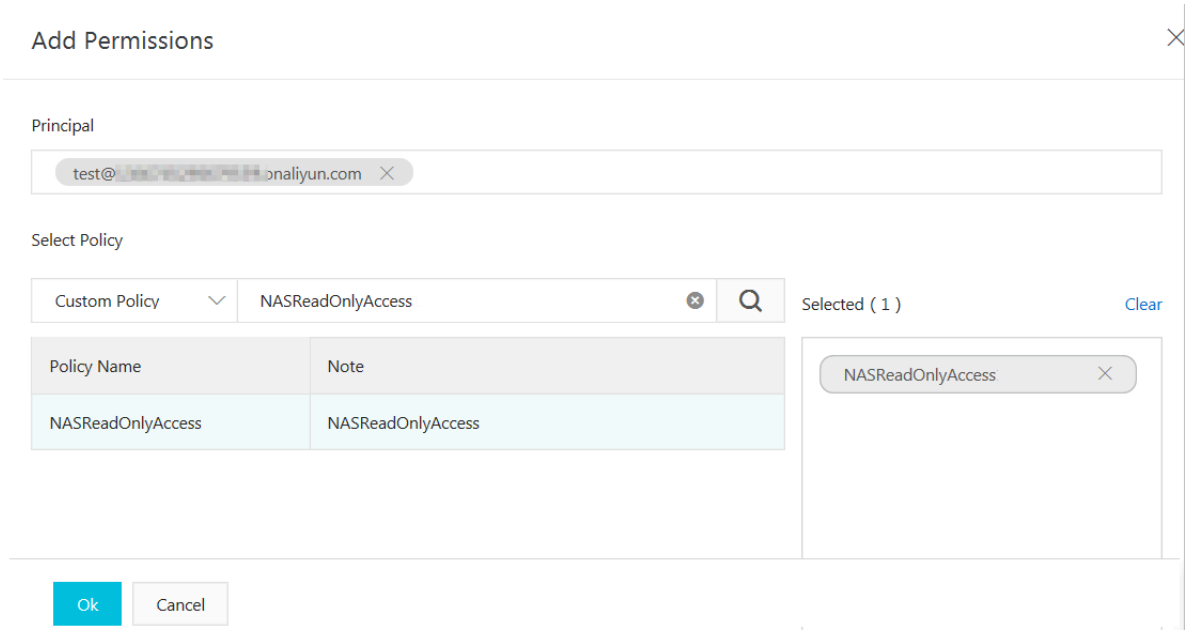

1.3 Manage permission groups

This topic describes how to manage permission groups in the Apsara File Storage NAS console. You can create and delete permission groups and rules. You can also view a list of permission groups and a list of rules.

Context

In NAS, each permission group represents a whitelist. You can add rules to a permission group to allow access to a file system from specific IP addresses or CIDR blocks. You can also grant different access permissions to different IP addresses or CIDR blocks.

After you activate NAS, a permission group named VPC default permission group (all allowed) is created. The default permission group allows read/write access to a file system from all IP addresses in a VPC. No limits are specified for root users.

Note:

- We recommend that you add rules only for required IP addresses and CIDR blocks to ensure data security.
- You cannot delete or modify the default permission group and its rules.
- You can create up to 10 permission groups for an Alibaba Cloud account.

Create a permission group and add rules

1. Log on to the NAS console.

2. Create a permission group.

   a) In the left-side navigation pane, choose **File System** > **Access Group** > **General Purpose NAS**. On the page that appears, click **Create Permission Group**.

   b) In the **Create a New Permission Group** dialog box, set the required parameters.

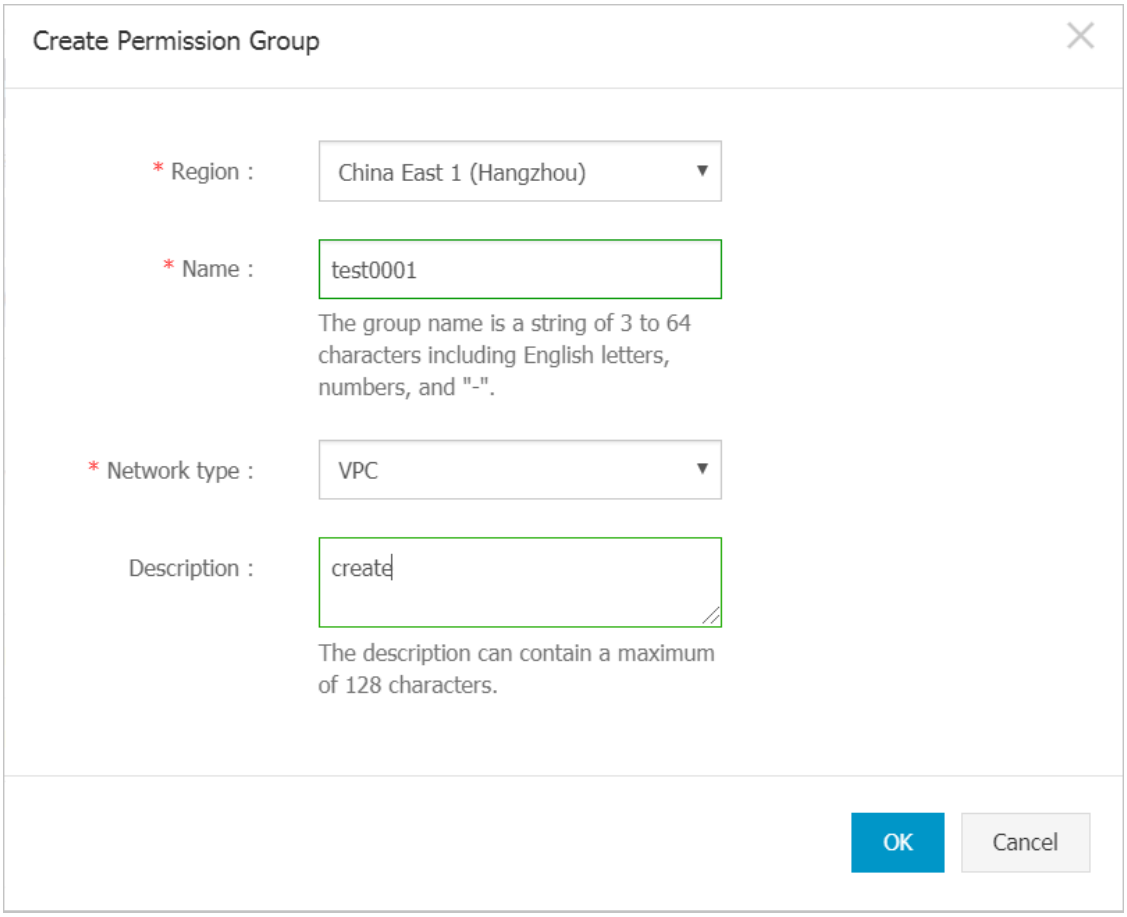

The following table describes the required parameters.

Parameter	Description
Name	The name of the permission group.
Network type	The network type. Valid values: VPC and Classic Network.

3. Add rules to the permission group.

   a) Find the permission group and click **Management Rules** in the Operations column.

   b) On the **List of Rules** page, click **Add Rules**.

   c) Set the required parameters.

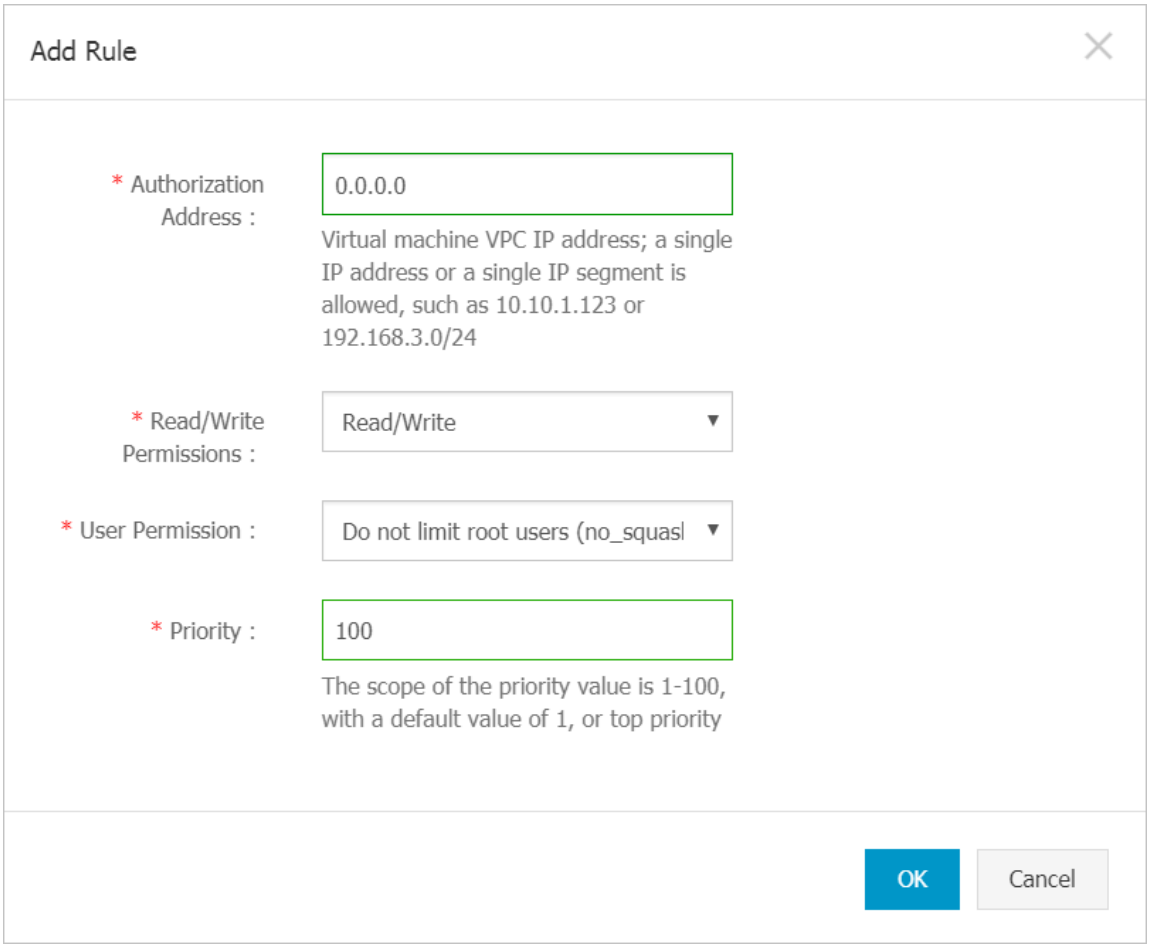

Parameter	Description
Authorized Address	Specifies the authorized object to which the rule is applied. **Note:** If you add a rule to a permission group of the classic network type, you can specify an IP address rather than a CIDR block for the parameter.
Read/Write Permission	Specifies whether to allow read-only or read/write access to the file system from the authorized object. Valid values: **Read-only** and **Read and write**.

Parameter	Description
User Permission	Specifies whether to limit the access to the file system from a Linux server. • **All Users Are Not Anonymous (No_Squash)**: allows the root user to access the file system. • **Root User anonymity (root_squash)**: denies access to the file system from the root user. The root user is treated as a nobody user. • **All Users Anonymous (All_Squash)**: denies access from all users. All users are treated as nobody users. A nobody user has the least permissions. To ensure high security, the nobody user can access only the open content of the server.
Priority	Specifies the priority of the rule. When multiple rules are applied to an authorized object, the rule with the highest priority takes effect. Valid values: 1 to 100 (1 indicates the highest priority).

d) Click **OK**.

What to do next?

On the **Access Group** page, you can perform the following operations.

Operation	Description
View a list of permission groups and the details of a permission group.	View the list of permission groups in a region and the details of a permission group. The details include the network type, number of rules, and number of associated file systems.
Modify a permission group.	Find the permission group and click **Edit** in the Operations column to edit the description of the permission group.
Delete a permission group.	Find the permission group and click **Delete** in the Operations column to delete the permission group.
View a list of rules.	Find the permission group and click **Management Rules** in the Operations column to view the list of rules in the permission group.
Modify a rule.	Click **Management Rules**. On the page that appears, find the rule, and click **Edit** in the Operations column to edit the Authorized Address, Read And Write Permissions, User Permissions, and Priority fields.

Operation	Description
Delete a rule.	Click **Management Rules**. On the page that appears, find the rule, and click **Delete** in the Operations column to delete the rule.

1.4 Apsara File Storage NAS NFS ACLs

1.4.1 Overview

Apsara File Storage NAS supports NFSv4 access control lists (ACLs) and Portable Operating System Interface (POSIX) ACLs. This topic describes POSIX ACLs and NFSv4 ACLs. It also lists precautions for using these ACLs.

Access control and user management are important for enterprise-level users who want to share files between different users and groups by using a shared file system. You can grant users and groups different types of access to specified files and directories. NAS provides Network File System (NFS) ACLs to allow you to meet specific requirements. An ACL consists of one or more access control entries (ACEs) that each grant a user or group one or more permissions to access a file or directory.

The NFSv3 protocol includes extended support for POSIX ACLs. POSIX ACLs extend the support for access control over file mode creation masks. You can grant permissions for specific users and groups besides users of the owner, group, and other classes. Permissions can also be inherited from parent objects. For more information, see acl - Linux man page.

The NFSv4 protocol includes extended support for NFSv4 ACLs that provide more fine-grained access control than POSIX ACLs do. For more information, see nfs4_acl - Linux man page.

You can use the NFSv3 protocol to mount a file system that has NFSv4 ACLs applied. These NFSv4 ACLs will then be converted into POSIX ACLs. You can also use the NFSv4 protocol to mount a file system that has POSIX ACLs applied. These POSIX ACLs will then be converted into NFSv4 ACLs. If you use NFS ACLs, we recommend that you mount NFSv4 file systems and control access by using NFSv4 ACLs rather than file mode creation masks and POSIX ACLs. The recommendation is based on the following aspects: NFSv4 ACLs and POSIX ACLs are not fully compatible. The interoperability between ACLs and file mode creation masks is not in an ideal state. The file systems that are mounted by using the NFSv3 do not support locks. For more information about NFS ACL features, see Features.

Note:

The NFS ACL feature is available only for NFS file systems in the following regions: China (Zhangjiakou-Beijing Winter Olympics), China (Beijing), China (Hohhot), China (Hangzhou), China (Shanghai), China (Chengdu), China (Hong Kong), Australia (Sydney), Indonesia (Jakarta), US (Silicon Valley), US (Virginia), Germany (Frankfurt), UK (London), and India (Mumbai). If the region where your file system resides does not support the NFS ACL feature, submit a ticket.

Precautions for using POSIX ACLs

- Configure ACLs
  - We recommend that you use the default inheritance method that allows a subdirectory or file to inherit the same ACL from the parent directory. This allows you to avoid configuring another ACL when you create a new file or subdirectory in the parent directory.
  - Use caution when you configure ACLs by using the recursive method (`setfacl -R`). Large amounts of metadata are produced when you perform a recursive operation on a directory that contains a large number of files and subdirectories. This may affect your businesses.
  - Before you configure ACLs, we recommend that you manage groups and related permissions. For example, you can add a user to one or more groups. If you want to add, remove, or modify permissions for a user, move the user to a group that has the required permissions. You do not need to modify the ACL of a group as long as the structure of groups remains unchanged. We recommend that you configure ACLs for groups rather than single users. This provides a simple and effective time-saving method to control access and ensure the better organization of permissions.
  - You can apply a POSIX ACL to multiple objects that resides on different clients. In such cases, you must ensure that the ACL you apply to each object is the same. Apsara File Storage NAS stores user IDs (UIDs) and group IDs (GIDs) at the backend. You must ensure that the mappings between a username or group name and a UID or GID are the same.
- Use ACLs
  - We recommend that you retain a minimum number of ACEs because a file system needs to scan all ACEs each time it performs permission verification. Abuse of ACLs may diminish the performance of file systems.

- Grant permissions to the other class
  - We recommend that you grant the least permissions to the other class because all users have the permissions that are granted to the other class. A potential security vulnerability may be exposed if the other class has more permissions than any ACE.
  - We recommended that you grant the least permissions to the other class. Before you create files or directories, you can use the `umask 777` command to configure the file mode creation mask. This command sets the file mode creation mask to 000 when the mask is used as a parameter to create a new file or directory. This ensures that the new file or directory has the least permissions. For more information, see umask and the default file mode creation mask.
  - We recommended that you grant the least permissions to the other class. Before creating files or directories, you can use the `umask 777` command to configure the file mode creation mask. This command sets the file mode creation mask to 000 when the mask is used as a parameter to create a new file or directory. This ensures that the new file or directory has the least permissions. For more information, see umask and the default file mode creation mask.
  - After you enable POSIX ACLs, the semantics of the other class for the POSIX ACL are equal to the semantics of the EVERYONE@ principal. The semantics of the other class for the file mode creation mask are also equal to the semantics of the EVERYONE@ principal. When a system performs permission verification, the system treats the other class the same as the EVERYONE@ principal.

Precautions for using NFSv4 ACLs

- Configure ACLs
  - Use UIDs or GIDs such as UID 1001 to configure ACLs.
  - We recommend that you use the default inheritance method that allows a subdirectory or file to inherit the same ACL from the parent directory. This allows you to avoid configuring another ACL when you create a new file or subdirectory in the parent directory.
  - Use caution when you configure ACLs by using the recursive method (`nfs4_setfacl -R`). Large amounts of metadata are generated when you perform a recursive operation on a directory that contains a large number of files and subdirectories. This may affect your businesses.

- Use ACLs
  - We recommend that you retain a minimum number of ACEs because a file system needs to scan all ACEs each time it performs permission verification. Abuse of ACLs may diminish the performance of file systems.
- Add ACEs to ACLs
  - We recommend that you do not configure the file mode creation mask after you configure an NFSv4 ACL.
  - The `nfs4_setfacl` command provides -a, -x, -m, and other options. You can use these options to add, remove, or modify ACEs. However, we recommend that you use `nfs4_setfacl -e <file>` the command to edit an ACL in an interactive mode.
  - NFSv4 ACLs have fine-grained permissions. In most cases, it is unnecessary to subdivide permissions at such a fine-grained level. For example, if you have the write (w) access to a file but do not have the append-only (a) access, an error may occur when you write data to the file. The same issue occurs on a directory. To avoid unexpected permission errors, we recommend that you specify a capital w (W) as a parameter when you use the `nfs4_setfacl` command to configure an ACL. The `nfs4_setfacl` command converts W to a full write access permission. For a file, W is expanded to wadT. For a directory, W is expanded to wadTD.
  - Before you configure ACLs, we recommend that you manage groups and related permissions. For example, you can add a user to one or more groups. If you want to add, remove, or modify permissions for a user, move the user to a group that has the required permissions. You do not need to modify the ACL of a group as long as the structure of groups remains unchanged. We recommend that you configure ACLs for groups rather than single users. This provides a simple and effective time-saving method to control access and ensure the better organization of permissions.
  - We recommend that you configure the least permissions for the EVERYONE@ principal because NFSv4 ACLs only support allow rather than deny ACEs. A potential security vulnerability may be exposed if the EVERYONE@ principal has more permissions than other ACEs.

1.4.2 Features

This topic describes the features of NFSv4 access control lists (ACLs) and POSIX ACLs.

Note:

The NFS ACL feature is available only for NFS file systems in the following regions: China (Zhangjiakou-Beijing Winter Olympics), China (Beijing), China (Hohhot), China (Hangzhou), China (Shanghai), China (Chengdu), China (Hong Kong), Australia (Sydney), Indonesia (Jakarta), US (Silicon Valley), US (Virginia), Germany (Frankfurt), UK (London), and India (Mumbai). If the region where your file system resides does not support the NFS ACL feature, submit a ticket.

Features of Apsara File Storage NAS NFSv4 ACLs

- Only access control entries (ACEs) of the allow type are supported. The following types of ACEs are not supported: deny, audit, and alarm.

  Deny ACEs increase the complexity of configuring permissions. In most cases, complexity leads to confusion and increases potential security risks. As agreed by the industry, we recommend that you avoid using deny ACEs. For more information about why deny ACEs are not recommended, see FAQ.

  Audit and alarm ACEs are not applicable to Apsara File Storage NAS NFS file systems. Instead, you can audit file systems and configure alerts based on auditing results in the Apsara File Storage NAS console.
- If no ACL is specified for a file or a directory, the default ACL that corresponds to the predefined file mode creation mask is applied.

```
touch file
```

```
[root@vbox test]# ls -l file
-rw-r--r--. 1 root root 0 May  6 14:27 file
```

```
[root@vbox test]# nfs4_getfacl file
# file: file
A::OWNER@:rwatTnNcCy
A::GROUP@:rtncy
```

```
A::EVERYONE@:rtncy
```

- An ACL is an ordered list that contains and deduplicates ACEs. This scheme ensures that permissions defined in an ACL are clear and informative.

  If you apply a new ACE and an existing ACE to the same object and the existing ACE is inherited from the parent object, the permissions of the new ACE override the permissions of the existing ACE. For example:

  - In an ACL that contains an ordered list of ACEs, ACEs that include the following principals are queued in sequence and take precedence over other ACEs: OWNER@, GROUP@, and EVERYONE@.

    ```
    [root@vbox test]# nfs4_getfacl file
    # file: file
    A::OWNER@:rwaxtTnNcCy
    A::GROUP@:rtncy
    A::EVERYONE@:rtncy
    A::1001:rwaxtTNcCy
    ```

  - Add an ACE of the read and write permissions to the following ACL for a user principal named 1009. The ACE is placed after the ACE that is defined for a user principal named 1001 based on the predefined order.

    ```
    [root@vbox test]# nfs4_setfacl -a A::1009:X file
    [root@vbox test]# nfs4_getfacl file
    # file: file
    A::OWNER@:rwaxtTcCy
    A::GROUP@:rwatcy
    A::EVERYONE@:tcy
    A::1001:rwaxtTNcCy
    A::1009:xtcy
    ```

  - Add a new ACE that includes the execute permission to the ACL. The system automatically merges the execute permission into the existing ACE for the 1009 user principal.

    ```
    [root@vbox test]# nfs4_setfacl -a A::1009:W file
    [root@vbox test]# nfs4_getfacl file
    # file: file
    A::OWNER@:rwaxtTcCy
    A::GROUP@:rwatcy
    A::EVERYONE@:tcy
    A::1001:rwaxtTNcCy
    A::1009:waxtTncCy
    ```

  - When you add the f and d inheritance flags to an ACE that includes a user principal named 1009, the system splits the ACE into two ACEs. One ACE has an extra inheritance flag named i specified, which indicates an inherit-only ACE. The other ACE only applies to the file object without any inheritance flag. If the inheritance type of an existing ACE matches the type for one of the two ACEs, the system combines the

existing ACE with the ACE from the two ACEs. The two matching ACEs are converted into one ACE.

```
[root@vbox test]# nfs4_setfacl -a A:fd:1009:R file
[root@vbox test]# nfs4_getfacl file
# file: file
A::OWNER@:rwaxtTcCy
A::GROUP@:rwatcy
A::EVERYONE@:tcy
A::1001:rwaxtTNcCy
A::1009:rwaxtTNcCy
A:fdi:1009:r
```

- All ACEs can be inherited.

  1. For example, the OWNER@ principal has the write access, the GROUP@ principal has the read access, and the EVERYONE@ has no access to the dir directory.

     ```
     [root@vbox nfs]# nfs4_getfacl dir
     # file: dir
     A::OWNER@:rwaDxtTnNcCy
     A::GROUP@:rxtcy
     A::EVERYONE@:tncy
     ```

  2. Add an ACE that grants a user principal named 1000 the read, write, and execute access to the dir directory. The f and d inheritance flags are also specified for the ACE.

     ```
     [root@vbox nfs]# nfs4_setfacl -a A:fd:1000:rwx dir
     [root@vbox nfs]# nfs4_getfacl dir
     # file: dir
     A::OWNER@:rwaDxtTcCy
     A::GROUP@:rxtcy
     A::EVERYONE@:tcy
     A::1000:rwx
     A:fdi:1000:rwx
     ```

  3. When you create a file or subdirectory in the dir directory, the file or the subdirectory automatically inherits all ACEs from the dir directory.

     ```
     [root@vbox nfs]# touch dir/file
     [root@vbox nfs]# nfs4_getfacl dir/file
     # file: dir/file
     A::OWNER@:rwatTcCy
     A::GROUP@:rwatcy
     A::EVERYONE@:rwatcy
     A::1000:rwx
     ```

     ```
     [root@vbox nfs]# mkdir dir/subdir
     [root@vbox nfs]# nfs4_getfacl dir/subdir
     # file: dir/subdir
     A::OWNER@:rwaDxtTcCy
     A::GROUP@:rwaDxtcy
     A::EVERYONE@:rwaDxtcy
     ```

```
A:fdi:1000:rwx
```

Note:

- We recommend that you grant the least privileges to the EVERYONE@ principal. Before you perform the following steps, we recommend that you run the `umask 777` command. This command ensures that no access to a file or directory is granted when the file or directory is created. For more information, see Why doesn't umask change execute permissions on files? [duplicate].
- When Linux calls functions to create files or directory, the predefined file mode creation mask is used as a request parameter. You can obtain the final ACL for a child object from the overlap of the inherited ACL (parent to child) and the file mode creation mask, as specified in the RFC7530 standard. When you modify the group bits of a file mode creation mask based on the standard, permissions included in an ACL for each group must be less than or equal to permissions defined in group bits. However, this scheme results in an invalid inheritance for groups. For example, you create a file and the file attempts to inherit A:RWX from a parent object. However, the predefined file mode creation mask sets the group bits to R. The final permission for the file becomes A:R. In actual practice, we recommend that you only modify file mode creation masks for ACLs that include the following principals: OWNER@, GROUP@, and EVERYONE@. This prevents against potential issues and ensures that semantics are clear. To remove permissions for a group, you only need to remove the ACE that relates to the group.

- You need to maintain mappings between usernames or group names and user IDs (UIDs) or group IDs (GIDs) across multiple independent instances.

  Apsara File Storage NAS NFS adopts IP security groups rather than usernames to authenticate users. When you configure NFSv4 ACLs, UIDs or GIDs that are included in ACEs are stored in Linux. When you print an ACL for an object in a shell, Linux automatically loads the /etc/passwd file and converts UIDs or GIDs into usernames or group names. You need to maintain mappings between usernames or group names and UIDs or GIDs across multiple instances. You must ensure a username or group name is mapped to its UID or GID.
- NFSv4 ACLs can be printed by using extended attributes.

```
[root@vbox nfs]# getfattr -n system.nfs4_acl file
# file: file
system.nfs4_acl=0sAAAABgAAAAAAAAAABYBhwAAAAZPV05FUkAAAAAAAAAAAAAA
ABIAhwAAAAZHUk9VUEAAAAAAAAAAAAAABIAhwAAAAlFVkVSWU9ORUAAAAAA
```

```
AAAAAAAAAAAAAAAEAAAAEMTAwMAAAAAAAAALAAAAAwAAAAQxMDAwAAAAAAAA
AEAAFgGQAAAABTEwMDAxAAAA
```

- Tools such as cp are supported for migrating NFSv4 ACLs.

  Apsara File Storage NAS supports migrating NFSv4 ACLs by using the cp, tar, and rsync tools. For more information, see How to preserve NFS v4 ACLs via extended attributes when copying file.

  The following `cp --preserve=xattr file1 file2` command makes a copy of the file1 file as the file2 file while making a copy of the ACL of the file1 file for the file2 file. The `cp -ar dir1 dir2` command makes a copy of the dir1 directory as the dir2 directory while making a copy of the ACL of the dir1 directory for the dir2 directory.

  Note:

  You may fail to migrate NFSv4 ACLs if the version of the rsync tool is earlier than 3.1.2.

  ```
  [root@vbox nfs]# nfs4_getfacl file1
  # file: file1
  A::OWNER@:rwatTcCy
  A::GROUP@:rwatcy
  A::EVERYONE@:rwatcy
  A::1000:rtncy
  [root@vbox nfs]# cp --preserve=xattr file1 file2
  ```

  ```
  [root@vbox nfs]# nfs4_getfacl file2
  # file: file2
  A::OWNER@:rwatTcCy
  A::GROUP@:rwatcy
  A::EVERYONE@:rwatcy
  A::1000:rtncy
  [root@vbox nfs]# cp -ar dir1 dir2
  ```

- Interoperation between NFSv4 ACLs and file mode creation masks is supported. The modification for the ACL of an object may change the file mode creation mask of the object. The modification for the file mode creation mask of an object may change the ACL of the object.

  For example, the file mode creation mask of the file object is 0666.

  ```
  [root@vbox nfs]# ls -l file
  -rw-rw-rw-. 1 root root 0 May  3  2019 file
  [root@vbox nfs]# nfs4_getfacl file
  # file: file
  A::OWNER@:rwatTcCy
  A::GROUP@:rwatcy
  ```

```
A::EVERYONE@:rwatcy
```

- If you add the execute permission to the file mode creation mask by modifying the owner bits, the execute permission is also added to the ACE that includes the OWNER@ principal.

```
[root@vbox nfs]# chmod u+x file
[root@vbox nfs]# ls -l file
-rwxrw-rw-. 1 root root 0 May  3  2019 file
[root@vbox nfs]# nfs4_getfacl file
# file: file
A::OWNER@:rwaxtTcCy
A::GROUP@:rwatcy
A::EVERYONE@:rwatcy
```

- If you add the execute permission to an ACE that includes the GROUP@ principal, the execute permission is also added to the related file mode creation mask.

```
[root@vbox nfs]# nfs4_setfacl -a A::GROUP@:x file
[root@vbox nfs]# ls -l file
-rwxrwxrw-. 1 root root 0 May  3  2019 file
```

Note:

- In the interaction between ACLs and file mode creation masks, the EVERYONE@ principal is equal to the others class. When you modify the others class, the change also applies to the EVERYONE@ principal. This operation causes a slight impact on the semantics of permissions. For example, the current file mode creation mask is 177. After you run the chmod o+r command, all users that include the file owner and group members have the read permission. This occurs because the read permission is added to the related ACE that includes the EVERYONE@ principal. If no change is applied to the default file mode creation mask, the owner and group classes still have no read permission after you run the chmod o+r command.
- If no change is applied to NFSv4 ACLs, the others class of the file mode creation mask keeps the same semantics. If an NFSv4 ACL is changed, the semantics of the others class changes to the semantics of the EVERYONE@ principal and keeps the semantics . We recommend that you do not use file mode creation masks after using NFSv4 ACLs.

- Interactions between NFSv4 ACLs and POSIX ACLs are supported.

  You can mount file systems that have NFSv4 ACLs applied by using the NFSv3 protocol . These NFSv4 ACLs will then be converted into POSIX ACLs. You can also mount file

systems that have POSIX ACLs applied by using the NFSv4 protocol. These POSIX ACLs will then be converted into NFSv4 ACLs.

Note:

The semantics of POSIX ACLs are different from the semantics of NFSv4 ACLs. For example, the inheritance rules that apply to POSIX ACLs do not differentiate files and directories. NFSv4 ACLs have more diverse permissions than POSIX ACLs, which have only read, write, and execute permissions. We recommend that you use either NFSv4 ACLs or POSIX ACLs to prevent against potential issues.

For example, you configure an NFSv4 ACL for the dir0 directory. The permissions are listed as follows.

```
[root@vbox test] sudo nfs4_getfacl dir0
A::OWNER@:tTnNcCy
A::GROUP@:tncy
A::EVERYONE@:tncy
A:fdi:EVERYONE@:tncy
A:fdi:OWNER@:tTnNcCy
A:fdi:GROUP@:tncy
A:g:19064:rxtncy
A:g:19065:rwaDxtTnNcCy
A:fdig:19064:rxtncy
A:fdig:19065:rwaDxtTnNcCy
```

You configure a POSIX ACL for the dir0 directory. The permissions are listed as follows.

```
[root@vbox test] sudo getfacl dir0
user::---
group::---
group:players:r-x
group:adminis:rwx
mask::rwx
other::---
default:user::---
default:group::---
default:group:players:r-x
default:group:adminis:rwx
default:mask::rwx
default:other::---
```

For example, you configure an NFSv4 ACL for the dir0/file file. The permissions are listed as follows.

```
[root@vbox test] sudo nfs4_getfacl dir0/file
A::OWNER@:tTnNcCy
A::GROUP@:tncy
A::EVERYONE@:tncy
A:g:19064:rxtncy
```

```
A:g:19065:rwaxtTnNcCy
```

For example, you configure a POSIX ACL for the dir0/file file. The permissions are listed as follows.

```
[root@vbox test] sudo getfacl dir0/file
user::---
group::---
group:players:r-x
group:adminis:rwx
mask::rwx
other::---
```

- The number of NFSv4 ACLs is limited.

  Apsara File Storage NAS supports a maximum of 100,000 ACLs that are different from one another in each file system by default. Each ACL contains a maximum of 500 ACEs.

  Note:
  We recommend that you do not abuse ACLs and ACEs. This reduces the time and resources consumed for verifying permissions.

Features of Apsara File Storage NAS POSIX ACLs

- Permissions that are specified for the other class apply to all.

  Everyone includes the owner, group, and users that are related to each ACE. The other class is equal to the EVERYONE@ principal of an NFSv4 ACL.

  Note:
  We recommend that you grant the least permissions to the other class for all cases.

  For example, the following ACL is configured for the myfile file. Although the ACE contains a user named alice who does not have the write permission, the write permission propagates to the ACE because the permission is specified for the other class.

```
[root@vbox 3]# getfacl myfile
# file: myfile
# owner: root
# group: root
user::rw-
user:alice:r--
group::r--
mask::r--
```

```
other::rw-
```

- Permissions that are configured by ACLs will not be changed after you run the chmod command.

  Note:
  We recommend that you avoid modifying the file mode creation mask of a file that has a POSIX ACL applied. You can configure permissions for the file by modifying the POSIX ACL.

  1. For example, an ACE that grants the players group the read and write access to the myfile file.

     ```
     [root@vbox 3]# getfacl myfile
     # file: myfile
     # owner: root
     # group: root
     user::rw-
     user:player:rw-
     group::rw-
     group:players:rw-
     mask::rw-
     other::---
     ```

  2. The chmod g-w myfile or chmod u-w myfile command does not change the permissions that are granted to the player user and the players group, which is different from the POSIX ACL standard. However, this ensures that permissions that are granted by POSIX ACLs to non-reserved users are the same after you modify permissions by using file mode creation masks. The non-reserved users include all users except for the users of the owner, group, and other classes.

     ```
     [root@vbox 3]# getfacl myfile
     # file: myfile
     # owner: root
     # group: root
     user::r--
     user:player:rw-
     group::r--
     group:players:rw-
     mask::rw-
     ```

```
other::---
```

- If the execute permission is not granted to the group and other classes of an ACL, the ACL has no execute permission.

  The rule is predefined in Linux. The execute action is allowed by the backend of Apsara File Storage NAS. However, to make the execute permission in the ACL effective, you must grant the execute permission to the group or other class.

  For example, if the group and other classes do not have the execute access to the myfile file, the player user cannot execute the file.

  ```
  [root@vbox 3]# getfacl myfile
  # file: myfile
  # owner: root
  # group: root
  user::rw-
  user:player:r-x
  group::r--
  mask::r-x
  other::r--
  ```

  If you grant the execute permission to the group class, the execute permission also propagates to the player user.

  ```
  [root@vbox 3]# getfacl myfile
  # file: myfile
  # owner: root
  # group: root
  user::rw-
  user:player:r-x
  group::r-x
  mask::r-x
  other::r--
  ```

- If you configure inheritable NFSv4 ACLs for directories, these settings may not conform to the POSIX ACL standard when these directories reside in NFSv3 file systems.

  Inheritance rules that apply to files are different from those that apply to directories in NFSv4 ACLs. The same inheritance rules apply to both files and directories in POSIX ACLs.

  Note:
  We recommend that you apply either NFS4 ACLs or POSIX ACLs to an NFS file system to prevent against potential issues.

- File mode creation masks cannot be modified.

  The file mode creation mask of a POSIX ACL is yielded by the combination and interaction of permissions from all users and groups. The mask has no practical meaning and cannot be changed.

- You need to maintain mappings between usernames or group names and user IDs (UIDs) or group IDs (GIDs) across multiple instances.

  Apsara File Storage NAS NFS adopts IP security groups rather than usernames to authenticate users. When you configure POSIX ACLs, UIDs or GIDs that are included in ACEs are stored in Linux. When you print an ACL for an object in a shell, Linux automatically loads the /etc/passwd file and converts UIDs or GIDs into actual usernames or group names. You need to maintain mappings between usernames or group names and UIDs or GIDs across multiple instances. You must ensure a username or group name is mapped to its related UID or GID.
- POSIX ACLs can be printed by using extended attributes.

```
[root@vbox nfs]# getfattr -n system.posix_acl_access file
# file: file
system.posix_acl_access=0sAgAAAAEAAAD/////AgAFACAEAAAEAAAA/////xAABQD/////
IAABAP////8=
```

- POSIX ACLs can be migrated by using tools such as cp.

  Apsara File Storage NAS supports migrating POSIX ACLs by using the cp, tar, and rsync tools. For more information, see How to preserve NFS v4 ACLs via extended attributes when copying file.

  The following `cp --preserve=xattr file1 file2` command makes a copy of the file1 file as the file2 file while making a copy of the ACL of the file1 file for the file2 file. The `cp -ar dir1 dir2` command makes a copy of the dir1 directory as the dir2 directory while making a copy of the ACL of the dir1 directory for the dir2 directory.

  Note:
  You may fail to migrate POSIX ACLs if the version of the rsync tool is earlier than 3.1.2.

```
[root@vbox nfs]# getfacl file1
user::---
user:player:r-x
group::---
mask::r-x
other::--x
[root@vbox nfs]# cp --preserve=xattr file1 file2
```

```
[root@vbox nfs]# getfacl file2
# file: file2
user::---
user:player:r-x
group::---
mask::r-x
other::--x
```

```
[root@vbox nfs]# cp -ar dir1 dir2
```

- The number of POSIX ACLs is limited.

  Apsara File Storage NAS supports a maximum of 100,000 ACLs that are different from one another in each file system by default. Each ACL contains a maximum of 500 ACEs.

  Note:
  We recommend that you do not abuse ACLs and ACEs. This reduces the time and resources consumed for verifying permissions.

FAQ

Why are deny ACEs not supported?

- The position of an ACE that resides in an ACL is important.

  The sequence for ACEs that reside in an NFSv4 ACL is random. A deny ACE may be placed in any position of an NFSv4 ACL. For example, an ACL contains two ACEs: A::Alice:r and D::Alice:r. The position of the ACEs determines whether the Alice user has the write permission.

  Note:
  When you configure an ACL, you must consider the position of each ACE.

- The number of ACEs in an ACL experience a sharp increase.

  You may have difficulties to combine and deduplicate ACEs in an ACL because the sequencing for ACEs is not mandatory. The number of ACEs may increase up to tens or hundreds over a long period of time. To manage the final permissions that are produced by these ACEs, you need to check each ACE. The process to check is strenuous and time-consuming.

- The interactions between file mode creation masks and ACLs become more complex after deny ACEs are applied because deny features do not exist in file mode creation masks.

  - If deny ACEs are available, you may need to add several ACEs to an ACL when the file mode creation mask is changed. For example, if you change the file mode creation mask to -rw-rw-rw, you need to add the following ACEs to an ACL. You must add the ACEs in sequence at the beginning of the ACL.

    ```
    A::OWNER@:rw
    D::OWNER@:x
    A::GROUP@:rw
    ```

```
D::GROUP@:x
A::EVERYONE@:rw
D::EVERYONE@:x
```

- If deny ACEs are unavailable, you can sequence and deduplicate ACEs. You do not need to differentiate the EVERYONE@ principal and the other class. You can modify an ACL with ease when the file mode creation mask is changed. In such cases, you only need to find ACEs that contain the OWNER@, GROUP@, and EVERYONE@ principals and modify these ACEs as follows.

```
A::OWNER@:rw
A::GROUP@:rw
A::EVERYONE@:rw
```

- Conversions between NFSv4 ACLs and POSIX ACLs are not supported in some cases.

  POSIX ACLs do not support deny ACEs. If one or more deny ACEs are included in an NFSv4 ACL, you cannot convert the ACL into a POSIX ACL.

1.4.3 Use POSIX ACLs to control access

This topic describes how to configure Portable Operating System Interface (POSIX) access control lists (ACLs). You can use POSIX ACLs to control access to files and directories that reside in an NFSv3 file system.

Prerequisites

An NFSv3 file system is mounted. For more information, see Mount an NFS file system.

Note:

The NFS ACL feature is available only for NFS file systems in the following regions: China (Zhangjiakou-Beijing Winter Olympics), China (Beijing), China (Hohhot), China (Hangzhou), China (Shanghai), China (Chengdu), China (Hong Kong), Australia (Sydney), Indonesia (Jakarta), US (Silicon Valley), US (Virginia), Germany (Frankfurt), UK (London), and India (Mumbai). If the region where your file system resides does not support the NFS ACL feature, submit a ticket.

Commands

Before you configure POSIX ACLs, we recommend that you familiarize yourself with the related commands.

Command	Description
getfacl <filename>	Shows the ACL that applies to the specified file.

Command	Description
setfacl -m g::w <filename>	Grants the owing group the write access.
setfacl -m u:player:w <filename>	Grants the player user the write access.
setfacl -m g:players:rwx <filename>	Grants the players group the read, write, and execute access.
setfacl -x g:players <filename>	Removes permissions from the players group
getfacl file1 | setfacl --set-file=- file2	Copies the ACL for the file1 file to the file2 file.
setfacl -b file1	Removes all extended ACEs from the file1 file. The base ACEs of the owner, group, and others are retained.
setfacl -k file1	Removes all default ACEs from the file1 file.
setfacl -R -m g:players:rw dir	Grants the players group the read and write access to files and subdirectories in the dir directory.
setfacl -d -m g:players:rw dir1	Grants the players group the read and write access to the new files and subdirectories in the dir1 directory.

Procedure

To control access to files and directories by configuring NFS ACLs, follow these steps.

1. Create users and groups.

   In this example, the following users are created: player, admini, and anonym. The following groups are created: players and adminis. The player user is added to the players group and the admini user is added to the adminis group.

   ```
   sudo useradd player
   sudo groupadd players
   sudo usermod -g players player
   sudo useradd admini
   sudo groupadd adminis
   sudo usermod -g adminis admini
   sudo useradd anonym
   ```

2. Configure POSIX ACLs to control access to files and directories.

   Use the following commands to complete the operations: create a directory named dir0 and grant the players group the read-only access, the adminis group the read,

write, and execute permissions, and the others class no access to all the files in the dir0 directory.

```
sudo umask 777
sudo mkdir dir0
sudo setfacl -m g:players:r-x dir0
sudo setfacl -m g:adminis:rwx dir0
sudo setfacl -m u::--- dir0
sudo setfacl -m g::--x dir0
sudo setfacl -m o::--- dir0
sudo setfacl -d -m g:players:r-x dir0
sudo setfacl -d -m g:adminis:rwx dir0
sudo setfacl -d -m u::--- dir0
sudo setfacl -d -m g::--x dir0
sudo setfacl -d -m o::--- dir0
```

Use the `sudo getfacl dir0` command to verify the result after the configuration is complete.

```
# file: dir0
# owner: root
# group: root
user::---
group::--x
group:players:r-x
group:adminis:rwx
mask::rwx
other::---
default:user::---
default:group::--x
default:group:players:r-x
default:group:adminis:rwx
default:mask::rwx
default:other::---
```

3. Verify the ACL configuration.

a) Use the following command to verify that the admini user has read and write access to the dir0/file file.

```
[root@vbox test] sudo su admini -c 'touch dir0/file'
[root@vbox test] sudo su admini -c 'echo 123 > dir0/file'
```

b) Use the following command to verify the read-only access of the player user.

```
[root@vbox test] sudo su player -c 'touch dir0/file'
touch: cannot touch ‘dir0/file’: Permission denied
[root@vbox test] sudo su player -c 'cat dir0/file'
123
[root@vbox test] sudo su player -c 'echo 456 >> dir0/file'
bash: dir0/file: Permission denied
[root@vbox test] sudo su player -c 'getfacl dir0/file'
# file: dir0/file
# owner: admini
# group: adminis
user::---
group::---
group:players:r-x
```

```
group:adminis:rwx
mask::rwx
other::---
```

c) Use the following command to verify that the anonym user does not have access to the dir0/file file.

```
[root@vbox test] sudo su anonym -c 'ls dir0'
ls: cannot open directory dir0: Permission denied
[root@vbox test] sudo su anonym -c 'cat dir0/file'
cat: dir0/file: Permission denied
[root@vbox test] sudo su anonym -c 'getfacl dir0/file'
getfacl: dir0/file: Permission denied
```

Related operations

If you want to remove user permissions, use the following method.

When you use NFSv4 ACLs, we recommend that you sort each user into different groups. This allows you to configure permissions for a group rather than a separate user. To disable access to an object from a user, you can remove the user from a group that has access to the object. For example, the following commands remove the admini user from the adminis group and add the user to the adminis2 group.

```
[root@vbox test] sudo groupadd adminis2
[root@vbox test] sudo usermod -g adminis2 admini
[root@vbox test] id admini
uid=1057(admini) gid=1057(admini) groups=1061(adminis2)
[root@vbox test] sudo su admini -c 'ls dir0'
ls: cannot open directory dir0: Permission denied
[root@vbox test] sudo su admini -c 'cat dir0/file'
cat: dir0/file: Permission denied
[root@vbox test] sudo su admini -c 'getfacl dir0/file'
getfacl: dir0/file: Permission denied
```

1.4.4 Use NFSv4 ACLs to control access

This topic describes how to configure NFSv4 access control lists (ACLs) and apply these ACLs to NFSv4 file systems to control access to files and directories.

Prerequisites

An NFSv4 file system is mounted. For more information, see Mount an NFS file system.

Note:

The NFS ACL feature is available only for NFS file systems in the following regions: China (Zhangjiakou-Beijing Winter Olympics), China (Beijing), China (Hohhot), China (Hangzhou), China (Shanghai), China (Chengdu), China (Hong Kong), Australia (Sydney), Indonesia (Jakarta), US (Silicon Valley), US (Virginia), Germany (Frankfurt), UK (London), and India

(Mumbai). If the region where your file system resides does not support the NFS ACL feature, submit a ticket.

Context

You can mount an NFSv4 file system on an Elastic Compute Service (ECS) instance that runs Linux and install the Linux-specific nfs4-acl-tools tool on the instance. You can use the standard nfs4_getfacl and nfs4_setfacl tools to configure NFSv4 ACLs after the installation is complete.

Description

Before you configure NFSv4 ACLs, we recommend that you familiarize yourself with the related commands.

Command	Description
nfs4_getfacl <filename>	Views the access permissions for the specified file.
nfs4_setfacl -a A::GROUP@:W <filename>	Adds an access control entry (ACE) that grants the GROUP@ principal the write access to the specified file.
nfs4_setfacl -a A::1000:W <filename>	Adds an ACE that grants a user principal named 1000 the write access to the specified file.
nfs4_setfacl -a A:g:10001:W <filename>	Adds an ACE that grants a group principal named 10001 the write access to the specified file.
nfs4_setfacl -e <filename>	Configures an ACL in an interactive mode.
nfs4_getfacl <filename> > saved_acl.txt	Saves a list of permissions for the specified file as a TXT file.
nfs4_setfacl -S saved_acl.txt <filename>	Configures permissions for the specified file by using a TXT file that includes a list of ready-made permissions.
nfs4_setfacl -m A::1001:rwaxTNcCy A::1001:rxtcy file1	Modifies the permission of an ACE that applies to the file1 file.
nfs4_getfacl file1 | nfs4_setfacl -S - file2	Copies the permissions for the file1 file to the file2 file.

Command	Description
`nfs4_getfacl file1 | grep @ | nfs4_setfacl -S - file1`	Deletes all ACEs that apply to the file1 file except for ACEs that include the following principals: OWNER@, GROUP@, and EVERYONE@.
`nfs4_setfacl -R -a A:g:10001:rW dir`	Adds an ACE that grants a group principal named 10001 the read and write access to files and subdirectories in the dir directory.
`find dir -type f -exec sh -c 'for ace in $( nfs4_getfacl \{} | grep "^A.*\:1005\:"); do nfs4_setfacl -x $ace \{}; done' \;`	Deletes ACEs that grant a user principal named 1005 any access to files in the dir directory.
`nfs4_setfacl -a A:fdg:10001:rW dir1`	Adds an ACE that grants a group principal named 10001 the read and write access to all newly created files and subdirectories in the dir1 directory.
`nfs4_setfacl -a A:fg:10001:rx dir1`	Adds an ACE that grants a group principal named 10001 the read and write access to all newly created files in the dir1 directory.

Procedure

You can configure NFSv4 ACLs to control access to files and directories by performing the following steps.

1. Create users and groups.

   In this example, the following users are created: player, admini, and anonym. The following groups are created: players and adminis. The player user is added to the players group and the admini user is added to the adminis group.

   ```
   sudo useradd player
   sudo groupadd players
   sudo usermod -g players player
   sudo useradd admini
   sudo groupadd adminis
   sudo usermod -g adminis admini
   ```

```
sudo useradd anonym
```

2. Install the related tools that are used to configure NFSv4 ACLs.

   If you have installed these tools, skip this step.

   ```
   sudo yum -y install nfs4-acl-tools
   ```

3. Obtain the group IDs of the players and adminis groups.

   Open the /etc/group file. The group IDs of the players and adminis groups are displayed as follows:

   ```
   players:x:19064:player
   adminis:x:19065:admini
   ```

4. Configure NFSv4 ACLs for files and directories.

   Use the following commands to complete the operations: create a directory named dir0 and add ACEs that grant the players group the read-only access, the adminis group the read, write, and execute access, and other users no access to all the files in the dir0 directory.

   ```
   sudo umask 777
   sudo mkdir dir0
   sudo nfs4_setfacl -a A:fdg:19064:RX dir0
   sudo nfs4_setfacl -a A:fdg:19065:RWX dir0
   sudo nfs4_setfacl -a A:fdg:OWNER@: dir0
   sudo nfs4_setfacl -a A:fdg:GROUP@: dir0
   sudo nfs4_setfacl -a A:fdg:EVERYONE@: dir0
   ```

   Use the `sudo nfs4_getfacl dir0` command to verify the configuration.

   ```
   A::OWNER@:tTnNcCy
   A::GROUP@:tncy
   A::EVERYONE@:tncy
   A:fdi:EVERYONE@:tncy
   A:fdi:OWNER@:tTnNcCy
   A:fdi:GROUP@:tncy
   A:g:19064:rxtncy
   A:g:19065:rwaDxtTnNcCy
   A:fdig:19064:rxtncy
   A:fdig:19065:rwaDxtTnNcCy
   ```

5. Verify the configuration of the ACL.

   a) Use the following commands to verify the read and write access of the admini user.

   ```
   [root@vbox test] sudo su admini -c 'touch dir0/file'
   [root@vbox test] sudo su admini -c 'echo 123 > dir0/file'
   ```

   b) Use the following command to verify the read-only access of the player user.

   ```
   [root@vbox test] sudo su player -c 'touch dir0/file'
   touch: cannot touch ‘dir0/file’: Permission denied
   [root@vbox test] sudo su player -c 'echo 456 >> dir0/file'
   ```

```
bash: dir0/file: Permission denied
[root@vbox test] sudo su player -c 'cat dir0/file'
123
[root@vbox test] sudo su player -c 'nfs4_getfacl dir0/file'
A::OWNER@:tTnNcCy
A::GROUP@:tncy
A::EVERYONE@:tncy
A:g:19064:rxtncy
A:g:19065:rwaxtTnNcCy
```

c) Use the following command to verify that the anonym user has no access to the /dir0/file file.

```
[root@vbox test] sudo su anonym -c 'ls dir0'
ls: cannot open directory dir0: Permission denied
[root@vbox test] sudo su anonym -c 'cat dir0/file'
cat: dir0/file: Permission denied
[root@vbox test] sudo su anonym -c 'nfs4_getfacl dir0/file'
Invalid filename: di
```

Related operations

If you want to remove user permissions, use the following method.

We recommend that you sort each user into different groups when you use NFSv4 ACLs. Then, when you configure NFSv4 ACLs, you only need to configure permissions for a group rather than a separate user. You can disable access to an object from a user by removing the user from a group that has access to the object. For example, use the following commands to remove the admini user from the adminis group and add the user to the adminis2 group:

```
[root@vbox test] sudo groupadd adminis2
[root@vbox test] sudo usermod -g adminis2 admini
[root@vbox test] id admini
uid=1057(admini) gid=1057(admini) groups=1054(adminis2)
[root@vbox test] sudo su admini -c 'ls dir0'
ls: cannot open directory dir0: Permission denied
[root@vbox test] sudo su admini -c 'cat dir0/file'
cat: dir0/file: Permission denied
[root@vbox test] sudo su admini -c 'nfs4_getfacl dir0/file'
Invalid filename: dir0/file
```

2 Manage file systems

This topic describes how to manage file systems in the Apsara File Storage NAS console. In the console, you can create file systems, delete file systems, view a list of the file systems and their details.

Create a file system

1. Log on to the NAS console.
2. Log on to the Apsara File Storage NAS console. In the left-side navigation pane, choose **NAS** > **File System List**. On the page that appears, click **Create File System**.

   If you want to create an Extreme NAS file system, click **Package Year** or **Pay-as-you-go** in the **Extreme NAS** section.
3. In the section, click **Pay-as-you-go**. On the page that appears, set the required parameters.

Parameter	Description
Region	The region where the file system resides. **Note:** If an ECS instance and a file system reside in different regions, the ECS instance cannot access the file system. You can create a maximum of 20 file systems in a region by using an Alibaba Cloud account. Storage types and protocol types that are supported by NAS change based on the region. For more information, see #unique_15.
Storage Type	The storage type. Valid values: **Performance** and **Capacity**. The maximum size of a Performance NAS file system is 1 PB. The maximum size of a Capacity NAS file system is 10 PB. The pay-as-you-go billing method is adopted.

Parameter	Description
Protocol Type	The protocol type. Valid values: **NFS (including NFSv3 and NFSv4)** and **SMB (2.1 and later)**. The Network File System (NFS) protocol is suitable for file sharing among Linux ECS instances. The Server Message Block (SMB) protocol is suitable for file sharing among Windows ECS instances.
Available Zone	The zone where the file system resides. Each region consists of multiple zones. Each zone has an independent power supply and network. If a file system and an ECS instance reside in different zones of the same region, they can still communicate with each other. Select a zone. We recommend that you select the zone where the ECS instance resides. Otherwise, extra latency may occur if the file system and the ECS instance communicate across zones.
Storage plan	Attach a storage plan if one is available.
Encryption	You can use keys that are hosted by Key Management Service (KMS) to encrypt static data on a file system. For more information, see Encrypt data. **Note:** • This parameter is available only for NFS file systems. • You do not need to decrypt data when you read and write it. • Valid values: None and Encryption.

4. Click **OK**.

View the file system list

On the **File System List** page, you can view all the file systems in a region. In the **File System List**, you can modify the name of the file system that you created.

View the details of a file system

Find the file system, and click the file system ID or **Management** to go to the **File System Details** page. You can view the basic information, mounting use, and performance monitoring of the file system.

Delete a file system

Before you delete a file system, you must remove all the mount targets of the file system.

Find the file system, move the pointer over More, and click**Delete** to delete the file system.

Warning:

Use caution when you delete a file system. After a file system is deleted, the data on the file system cannot be restored. Ensure that all data is backed up.

3 Manage mount targets

This topic describes how to manage mount points in the Apsara File Storage NAS console. The management includes creating, deleting, enabling, and disabling mount points. It also includes viewing a list of mount points and modifying the permission group of a mount point.

Create a mount target

To mount a file system on an Elastic Compute Service (ECS) instance, you must add a mount target for the file system. To create a mount target in the NAS console, you can follow these steps.

Note:

NAS General-purpose file systems support mount points of the VPC and classic network types. You can create up to two mount points for each file system.

1. Log on to the NAS console.
2. Choose **File System** > **File System List**.
3. Find the target file system and choose **More** > **Add Mount Target**.

4. In the **Add Mount Point** dialog box, configure the required parameters.

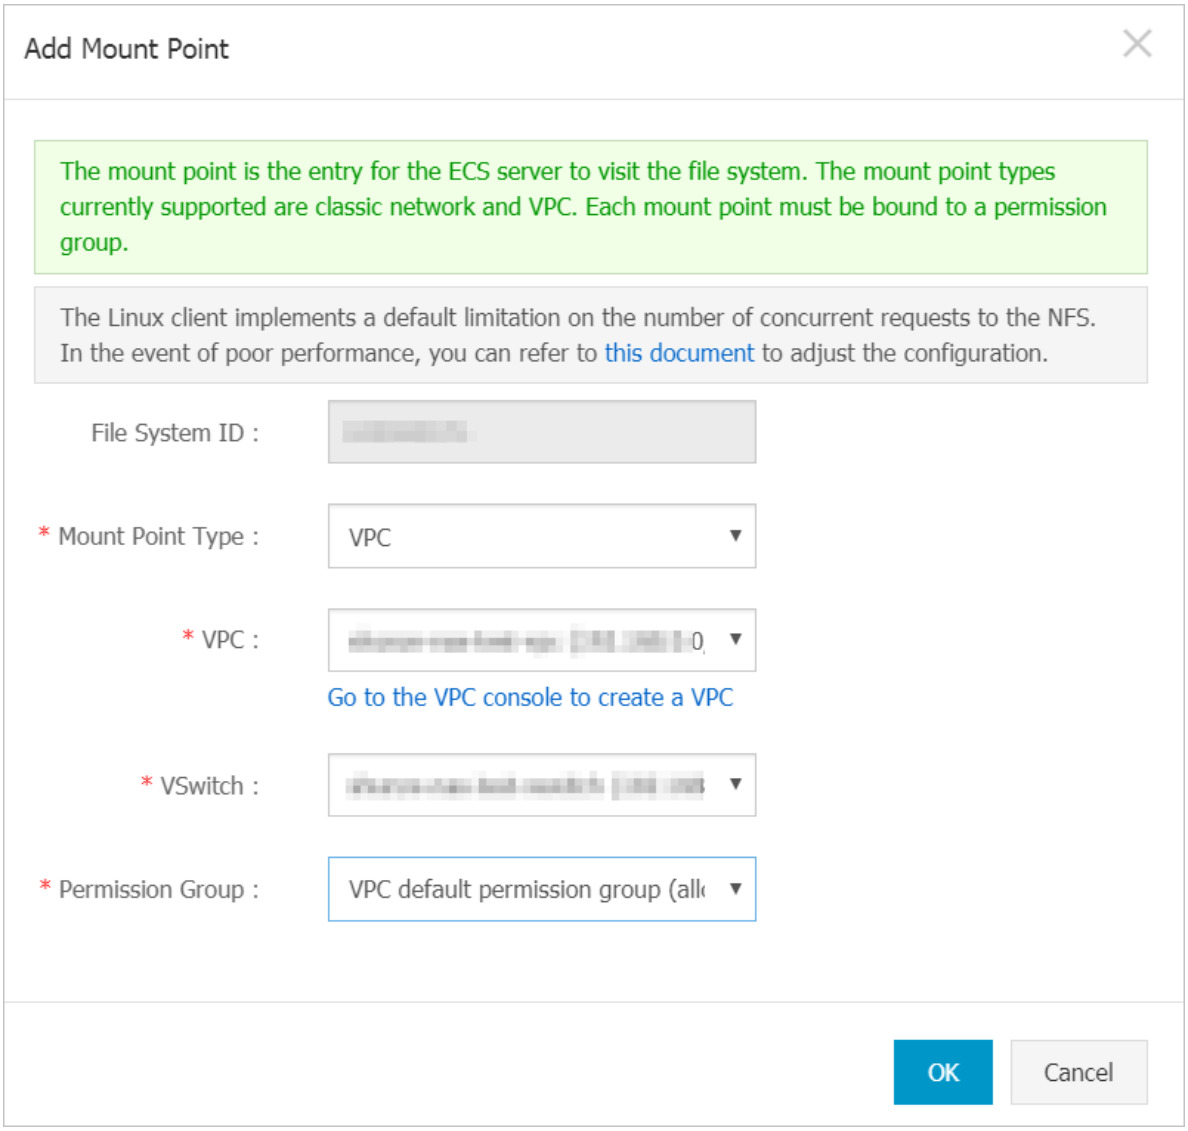

Mount Point Type: specifies the network type of mount point. Valid values: VPC and classic network.

- If you want to create a mount point of the Virtual Private Cloud (VPC) type, configure the parameters that are described in the following table.

Parameter	Description
VPC Network	The VPC where the mount point resides. If no VPC exists, create a VPC in the VPC console. **Note:** The VPC and VSwitch that you select must be the same as those of the ECS instance on which you mount a file system. An Elastic Compute Service (ECS) instance and a file system may reside in different VPCs. To enable communication between the ECS instance and the file system, you can use Cloud Enterprise Network (CEN) to connect VPCs. For more information, see #unique_18.
VSwitch	The VSwitch that resides in the VPC.
Permission Group	The permission group. The VPC default permission group is generated for each Alibaba Cloud account, which allows access to the file system through the mount point from all IP addresses of the VPC. For more information about how to create a permission group, see Create a permission group and add rules.

- If you want to create a mount point in a classic network, configure the parameter that is described in the following table.

Parameter	Description
Permission group	The permission group. **Note:** - You can add a mount point of the classic network type only in a region that resides in China. - You can attach a mount point of the classic network type only to an ECS instance.

5. After the configuration is complete, click **OK**.

View a list of mount targets

On the **File System List** page, find the target file system, and click **Manage** to open the **File System Details** page. In the **Mount Target** section, view a list of mount targets.

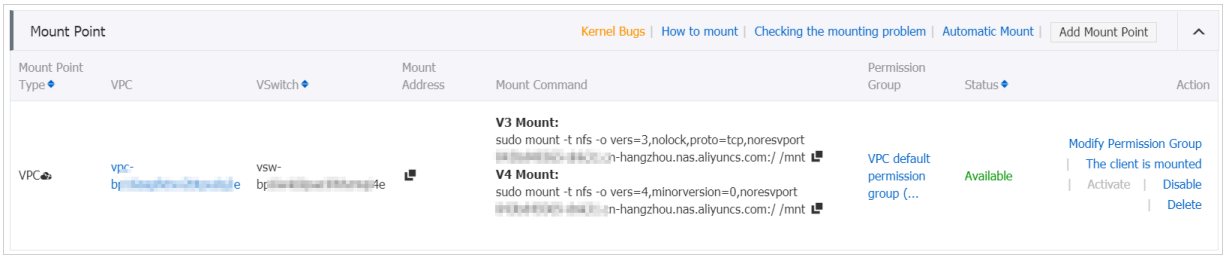

View a list of clients on which a file system is mounted

To view a list of clients on which a file system is mounted, click **Clients**. The IP address of each client is also displayed.

Note:

The list shows clients that have used the file system within the last minute. Some clients that have the file system attached but never used it may be excluded from the list.

Enable or disable a mount target

To control access to the mount target from clients, perform the following operations.

- To deny access to the mount target from clients, click **Disable**.
- To allow access to the mount target from clients, click **Enable**.

Delete a mount target

To delete a mount target, click **Delete**.

Warning:

Use caution when you delete a mount target. After you delete a mount target, the mount target cannot be restored.

Modify the permission group of a mount target

To modify the permission group of a mount target, click **Modify Permission Group**. For more information about permission groups, see Manage permission groups.

4 Manage file system quotas

This topic describes how to use the Alibaba Cloud quota_tool tool to manage Apsara File Storage NAS (NAS) quotas on ECS instances that have NAS file systems mounted. You can configure, view, and cancel quotas on these ECS instances.

Prerequisites

An NFS file system of the NAS Capacity or NAS Performance type is mounted on an ECS instance. For more information, see #unique_12.

Context

NAS allows you to view and manage directory-level quotas with ease. Directory-level quotas specify the maximum number of files in each folder and the maximum storage space that is allowed for these files.

From the perspective of the application scope, quotas are sorted into quotas for all users and quotas for a single user or group. Quotas for all users specify the maximum storage space allowed for files that all users can create in a directory. Quotas for a single user or group specify the maximum storage space allowed for files that a user or group can create in a directory.

From the perspective of limits, quotas are sorted into quotas for statistics and quotas for restriction. Quotas for statistics only collect the usage of the storage space. It provides an easy method to retrieve and view statistical data. Quotas for restriction specify the maximum capacity of storage space for files that you can create in a directory. If the limit is exceeded, you may fail to create a file or subdirectory, append data to a file, or perform other operations.

Notice:

- You can only configure quotas for statistics.
- You can only configure quotas for file systems that reside in China (Zhangjiakou-Beijing Winter Olympics), China (Hohhot), Australia (Sydney), Malaysia (Kuala Lumpur), and US (Silicon Valley).
- NAS performs asynchronous calculation for quotas at the backend. When you use the quota_tool tool to retrieve statistical data about quotas, the process requires a period of time to complete. In most cases, the time period is about 5 to 15 minutes.

Configure quotas

The following uses the /mnt directory as an example.

1. Log on to the Elastic Compute Service (ECS) instance with the root account.

   You can use the quota_tool tool on an ECS instance that has a NAS file system mounted. You must run the tool with the root permissions. The following describes how to use the quota_tool tool on an ECS instance.

2. Use the following command to download the quota_tool tool.

   ```
   wget https://nasimport.oss-cn-shanghai.aliyuncs.com/quota_tool_v1.0 -O quota_tool
   ```

3. Use the following command to grant the execute permission to the quota_tool tool.

   ```
   sudo chmod a+x quota_tool
   ```

4. Configure quotas.

   Note:
   For each file system, you can only configure quotas for a maximum of 10 directories.

   The syntax of the command that you use to configure quotas is `sudo ./quota_tool set --dir [DIR] [OPTION]`.

Parameter	Description
--dir [DIR]	Specifies the directory for which you want to configure quotas. For example, --dir /mnt/data/.

Parameter	Description
OPTION	Specifies the required options for the OPTION parameter. **Note:** When you specify options for the parameter, you must follow these rules. the --accounting option is required, one of the following options is required: --alluser, --uid, and --gid, . • --accounting: specifies a quota for statistics. • --alluser: specifies a directory-level quota for all users. • --uid: introduces the UID of a user. For example, --uid 505 only collects a quota for a user whose UID is 505. • --gid: introduces the GID of a group. For example, --gid 1000 only collects a quota for a group whose GID is 1000.

The following examples describe how to configure quotas.

- Use the following command to configure a quota for statistics for the /mnt/data/ directory to retrieve the total number of files that reside in a directory.

```
sudo ./quota_tool set --dir /mnt/data/ --accounting --alluser
```

- Use the following command to configure a quota for statistics for the /mnt/data/ directory to retrieve the total number of files that are created by a user whose UID is 505.

```
sudo ./quota_tool set --dir /mnt/data/ --accounting --uid 505
```

Retrieve quotas

After you configure a quota for a NAS directory, you can retrieve statistical data about the quota for the directory.

1. Log on to the ECS instance with the root account.
2. Use the following command to retrieve quotas.

```
sudo ./quota_tool get --dir /mnt/data/
```

Use the `sudo ./quota_tool get --dir /mnt/data/ --all` command to retrieve all quotas you specify for the /mnt/data/ directory.

Note:

- During the process of retrieving a quota for the first time, a state called **Initializing** appears. After the initialization process is complete, you can receive the quota and a result showing **success** appears. The duration of the initialization process is based on the number of files and subdirectories that are included in a directory.
- Before the expected FileCountReal and SizeReal are displayed, a delay of 5 to 10 minutes may occur when you perform daily retrieval of quotas after the initialization process is complete. This occurs due to the asynchronous calculation for quotas at the backend.

The following table lists parameters that are included in a response in the JSON format.

Parameter	Description
Path	Indicates a directory for which you retrieve a quota.
Report	Includes all information about a quota that is specified for a directory, for example, UID and GID.
ReportStatus	The state for the retrieval of a quota.
FileCountLimit	Indicates the limit for the number of files. A value of Empty indicates no limit.
FileCountReal	Indicates the total number of files including subdirectories, files, and special files that reside in a directory.

Parameter	Description
QuotaType	Accounting indicates a quota for statistics .
Uid	Indicates the UID of a user. A value of All indicates all users.
Gid	Indicates the GID of a group. The value of All indicates all groups.
SizeLimit	Indicates the maximum capacity of files that reside in a directory . The value of Empty indicates no limit.
SizeReal	Indicates the total capacity of files that reside in a directory.

Cancel quotas

After you configure a quota, you can also cancel the quota.

1. Log on to the Elastic Compute Service (ECS) instance with the root account.
2. Cancel quotas.

   The syntax of the command that you can use to cancel quotas is `sudo ./quota_tool cancel --dir [DIR] [OPTION]`.

Parameter	Description
--dir [DIR]	Specifies the directory for which you want to cancel quotas, for example, -dir /mnt/data/.

Parameter	Description
OPTION	Configure the required option for the OPTION parameter.  **Note:** When you configure the OPTION parameter, one of the following options is required: --alluser, --uid, and --gid. • --alluser: cancels a quota for all users. • --uid: introduces the UID of a user. For example, --uid 505 cancels a quota for a user whose UID is 505. • --gid: introduces the GID of a group. For example, --gid 505 cancels a quota for a group whose GID is 505.

The following examples describes how to cancel quotas.

- You configure a quota for the /mnt/data/ directory. Use the following command to cancel the quota for a user whose UID is 100.

  ```
  sudo ./quota_tool cancel --dir /mnt/data/ --uid 100
  ```

- You configure a quota for the /mnt/data/ directory. Use the following command to cancel the quota for all users.

  ```
  sudo ./quota_tool cancel --dir /mnt/data/ --alluser
  ```

5 Manage snapshots

You can create snapshots for Extreme NAS file systems. This topic describes how to manage snapshots in the Apsara File Storage NAS console. You can create or delete snapshots. You can also create, apply, or delete snapshot policies.

Prerequisites

An Extreme NAS file system is created. For more information, see Mount a NAS Extreme file system.

Context

Snapshots are commonly used to ensure data backup and recovery. To eliminate the risk of data loss, you can create a snapshot for the file system before you perform an operation. You can create manual snapshots for file systems. You can also use automatic snapshot policies to create auto snapshots.

> **Note:**
> Snapshots are only available for Extreme NAS file systems.

Create a snapshot

1. Log on to the NAS console.
2. Log on to the NAS console. In the left-side navigation pane, choose **Data Service** > **Snapshot**. On the **Snapshot** tab, click **Manually Create a Snapshot**.

   > **Note:**
   > - You can create a maximum of 128 snapshots for a file system.
   > - You can create snapshots only for a file system that is in the Running state.
   > - You can create only one snapshot at a time.
   > - If the file system expires during the creation of a snapshot, the file system is released and the created snapshot is deleted.
   > - When you create a snapshot for the file system, the file system may exhibit a short-term decrease in I/O performance. We recommend that you create snapshots during off-peak hours.

- A snapshot is a backup of a file system at a specific point in time. During the process of creating a snapshot, incremental data that is generated by the operating system is not synchronized to the snapshot.
- Manual snapshots remain in a file system for persistent storage. We recommend that you delete snapshots that you do not need on a regular basis to reduce the extra costs that they incur.
- If your account is overdue for more than 15 days, manual snapshots are deleted.

3. In the **Manually Create a Snapshot** dialog box, set the parameters.

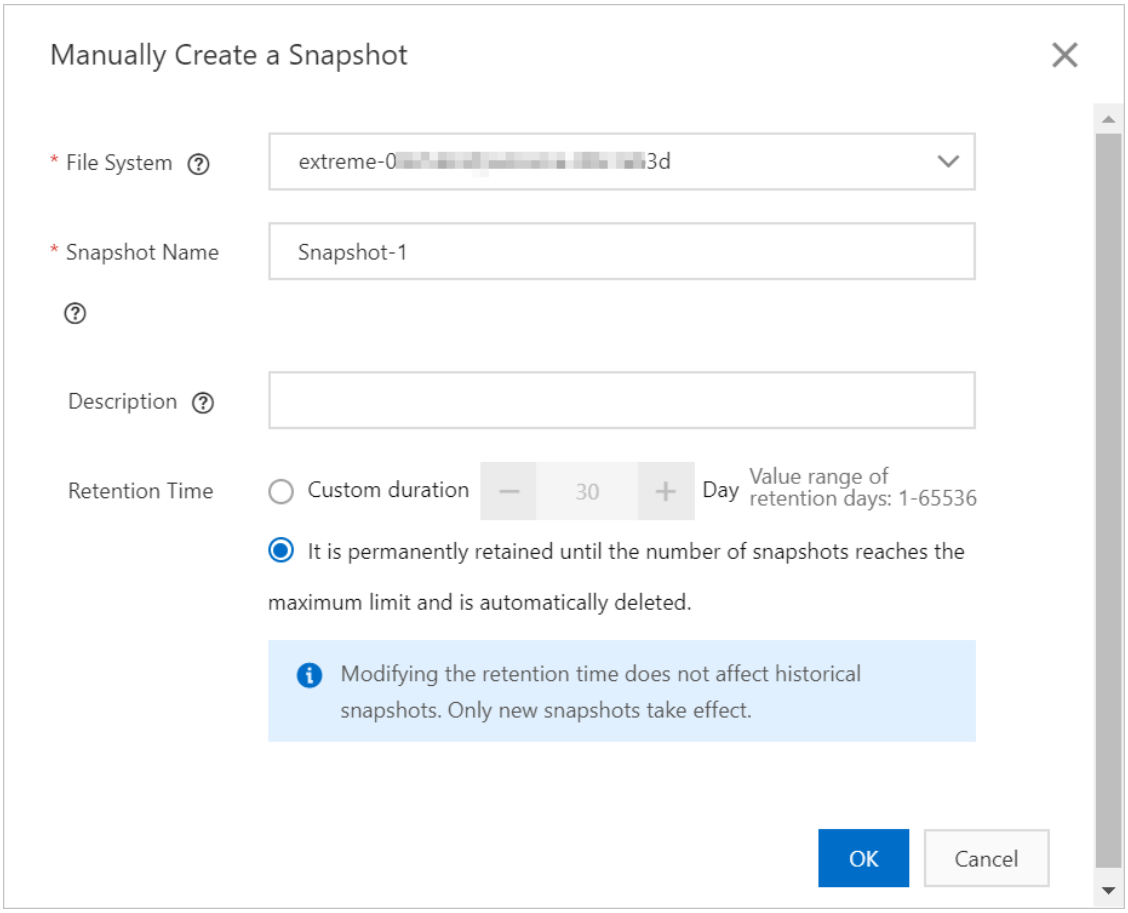

The following table describes the required parameters.

Parameter	Description
File System	Specify the ID of an Extreme NAS file system.
Retention Time	Select a retention period based on your needs: • **Custom duration** specifies a duration that ranges from 1 to 65536. Unit: day. • Select **It is permanently retained until the number of snapshots reaches the maximum limit and is automatically deleted.**.

4. Click **OK** to create the snapshot.

Create an auto snapshot

You can use an automatic snapshot policy to create auto snapshots.

1. Log on to the NAS console.

2. Create an automatic snapshot policy.
   a) Log on to the NAS console. In the left-side navigation pane, choose **Data Service** > **Snapshot**. On the **Automatic Snapshot Policy** tab, click **Create an Automatic Snapshot Policy**.
   b) In the **Create an Automatic Snapshot Policy** dialog box, set the required parameters.

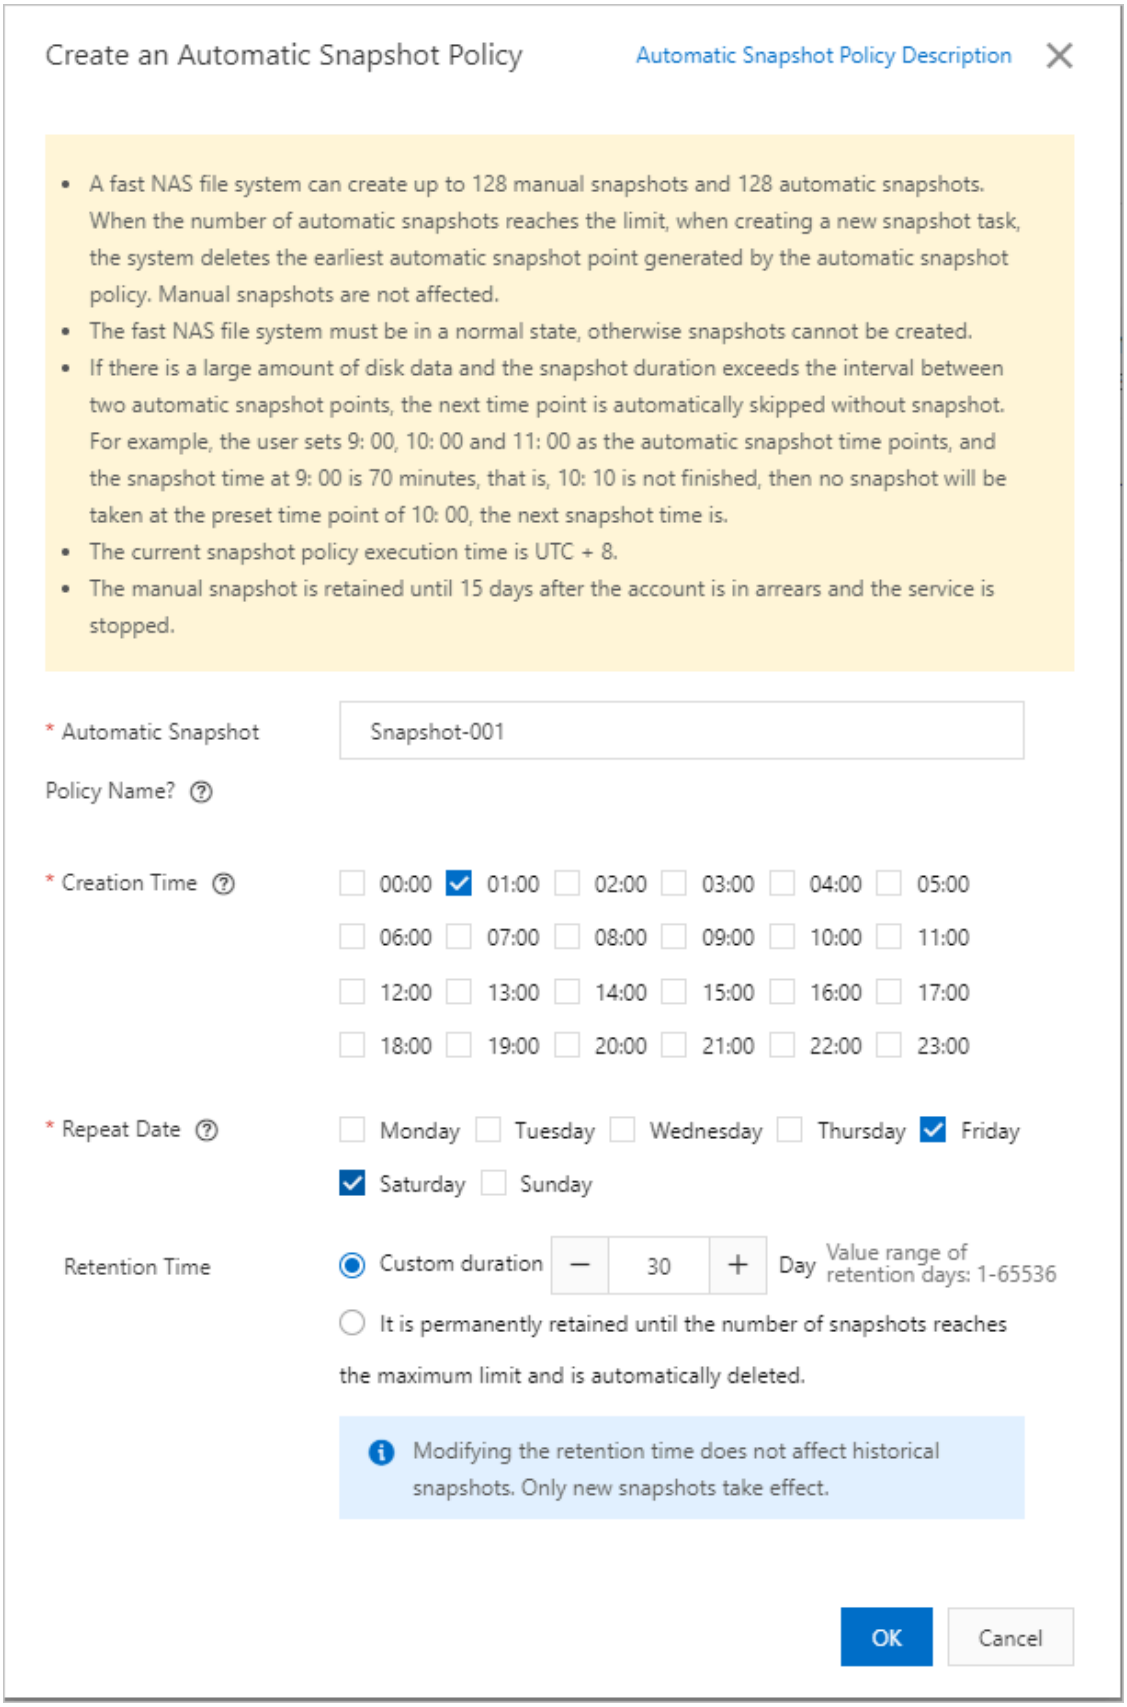

The following table describes the required parameters.

Parameter	Description
Creation Time	Specify the time when an auto snapshot is created. You can select one or more time points within the range of 00:00 to 23:00.
Repeat Date	Specify the frequency at which auto snapshots are created. You can select multiple days from Monday to Sunday.
Retention Time	• Select **Custom duration** and specify a durationwithin the range of 1 to 65536. Unit: days. • Select **It is permanently retained until the number of snapshots reaches the maximum limit and is automatically deleted.**.

Note:

- You can create a maximum of 100 automatic snapshot policies in each region for a single Alibaba Cloud account.
- You can apply an automatic snapshot policy to multiple file systems.
- If you modify the retention time of an automatic snapshot policy, the modificati on applies only to the snapshots that are created after the modification. The retention time of previous snapshots is not affected.

c) Click **OK**.

3. Apply the automatic snapshot policy.

a) Find the automatic snapshot policy, and click **Apply to The File System**.

b) In the **File System ID** section of the **Apply to The File System** dialog box, select the file systems and add the automatic snapshot policy to the **Apply to The File System** section.

Note:

- You can create a maximum of 100 automatic snapshot policies in each region for a single Alibaba Cloud account.
- If an auto snapshot is being created when the scheduled time for a new automatic snapshot arrives, the new snapshot creation is canceled. This may occur when the file system stores a large volume of data.

  For example, you have scheduled auto snapshots to be created at 09:00, 10:00, 11:00, and 12:00. The system starts to create a snapshot at 09:00 and completes

the process at 10:20. The process takes 80 minutes because the file system has a large volume of data. The system does not create a snapshot at 10:00, but creates a snapshot at 11:00.

- A file system has a maximum of 64 auto snapshots. If this limit is reached, the earliest auto snapshots are deleted. This rule does not apply to manual snapshots .
- If your account is overdue for more than 15 days, manual snapshots are deleted.
- If you modify the retention time of an automatic snapshot policy, the modificati on applies only to the snapshots that are created after the modification. The retention time of previous snapshots is not affected.
- If an automatic snapshot is being created for a file system, you cannot create a manual snapshot for the file system. You must wait after the automatic snapshot is created.
- You can only apply automatic snapshot policies to a file system that is in the Running state.
- All automatic snapshots are named by using the auto_yyyyMMdd_X format.

  For example, auto_20140418_1 indicates the name of the first automatic snapshot that was created on April 18, 2014. In the naming format, auto indicates that the snapshot is an auto snapshot which is different from a manual snapshot. yyyyMMdd indicates the date when a snapshot is created. yyyy stands for the year , MM the month, and dd the day of the month. X indicates the ordinal number of the snapshot.

c) Click **OK**.

After you apply an automatic snapshot policy to a file system, NAS creates automatic snapshots for the file system based on the policy.

Use a snapshot to create a file system

To create a file system, you can specify a snapshot when you call API operations.

1. Install Python and the SDK for Python.

```
pip install aliyun-python-sdk-corepip
pip install aliyun-python-sdk-bssopenapipip
```

```
pip install aliyun-python-sdk-nas
```

2. Customize the code that is provided in the SDK for Python and run the code to create a file system.

   The sample code creates a pay-as-you-go file system. If you want to create a subscripti on file system, you must customize the sample code.

   The sample code includes the following required parameters.

   - accessKeyId and accessSecret: Specify the AccessKey ID and AccessKey secret of your Alibaba Cloud account. For information about an AccessKey pair, see #unique_22.
   - set_parameters: Specify the required parameters of the file system.

```
#! /usr/bin/env python
# coding=utf-8

from aliyunsdkcore.client import AcsClient
from aliyunsdkcore.acs_exception.exceptions import ClientException
from aliyunsdkcore.acs_exception.exceptions import ServerException
from aliyunsdkbssopenapi.request.v20171214.GetPayAsYouGoPriceRequest import
GetPayAsYouGoPriceRequest
from aliyunsdkbssopenapi.request.v20171214.CreateInstanceRequest import
CreateInstanceRequest
from aliyunsdknas.request.v20170626.DescribeFileSystemsRequest import DescribeFi
leSystemsRequest

client = AcsClient('<accessKeyId>', '<accessSecret>', 'cn-hangzhou')

def Create():
    request = CreateInstanceRequest()
    request.set_accept_format('json')
    request.set_ProductCode("nas")
    ## Pay-as-you-go
    request.set_SubscriptionType("PayAsYouGo")
    request.set_ProductType("nas_extreme_post")
    ## Subscription
    # request.set_SubscriptionType("Subscription")
    # request.set_ProductType("nas_extreme")
    # request.set_Period(1) #The subscription duration. Unit: months.
    request.set_Parameters([
        {
            "Code": "Region",
            "Value": "cn-shanghai"
        },
        {
            "Code": "Zone",
            "Value": "cn-shanghai-g"
        },
        {
            "Code": "ProtocolType",
            "Value": "NFS"
        },
        {
            "Code": "StorageType",
            "Value": "standard"
        },
        {
            "Code": "Size",
```

```
            "Value": "100"
        },
        {
            "Code": "Throughput",
            "Value": "150"
        },
        {
            "Code": "SnapshotId",
            "Value": "s-extreme-xxxxxxxxxx"
        }
    ])
    response = client.do_action_with_exception(request)
    print response
if __name__ == '__main__':
    Create()
```

Related operations

Operation	Description
Cancel an automatic snapshot policy	Perform the following steps: 1. On the **File System List** page, find the file system, and choose **More** > **Snapshot** > **Set Snapshot Policy**. 2. In the **Set Snapshot Policy** dialog box, click **Cancel** to cancel the policy that is applied to the file system.
View a snapshot	On the **Snapshot** tab, view the list of snapshots and their details.
Roll back a file system	On the **Snapshot** tab, find the snapshot and click **Rollback**. You can roll back a file system to a previous snapshot.
Delete a snapshot	On the **Snapshot** tab, find the snapshot and click **Delete**.
View automatic snapshot policies	On the **Automatic Snapshot Policy** tab, view the list of automatic snapshot policies and their details.
View the file systems to which an automatic snapshot policy is applied	On the **Automatic Snapshot Policy** tab, find the automatic snapshot policy and click **Apply to The File System** to view the systems to which the policy is applied.
Modify an automatic snapshot policy	On the **Automatic Snapshot Policy** tab, find the automatic snapshot policy and click **Modify Policy**.
Delete an automatic snapshot policy	On the **Automatic Snapshot Policy** tab, find the automatic snapshot policy and click **Delete**.

6 Data backup

Data stored on Apsara File Storage NAS file systems cannot be directly backed up. You need to use Hybrid Backup Recovery (HBR) to back up Apsara File Storage NAS data.

You can back up Apsara File Storage NAS data in the Apsara File Storage NAS console and restore backups as needed.

> **Note:**
> Apsara File Storage NAS backups do not take up space within the current file system.

Prerequisites

A NAS file system that complies with the NFS or SMB protocol is created for backup. For more information, see Mount a file system on a Linux ECS instance and Mount a file system on a Windows ECS instance.

Preparations

Before backing up Apsara File Storage NAS data, you must activate the HBR service.

1. Log on to the Apsara File Storage NAS console.
2. Choose **Data Service** > **File Backup** and click **Apply HBR Service**.
3. Activate the HBR service as prompted on the page.

Step 1: Create a backup plan

> **Note:**
> We recommend that each created NAS backup job contain no more than 50 million files, and the total number of files and subdirectories in each directory be no more than 8 million.

1. Log on to the Apsara File Storage NAS console.
2. In the left-side navigation pane, choose**Data Service** > **File Backup**.
3. For more information about creating a back up plan, see Step 3 to Step 5 in Create a backup plan.

Step 2: Create a restoration job

For more information, see Create a restoration job.

Optional operations

For more information, see #unique_26/unique_26_Connect_42_section_lak_3jr_12s.

7 Migrate data

You can use the Data Transport service to migrate data between Network Attached Storage (NAS) file systems or between a NAS file system and an Object Storage Service (OSS) bucket.

Log on to the Data Transport console to migrate data.

- For more information about migrating data from a NAS file system to an OSS bucket, see Migrate data from NAS to OSS.
- For more information about migrating data from one NAS file system to another, see Migrate data between NAS file systems.
- For more information about migrating data from an OSS bucket to a NAS file system, see Migrate data from OSS to NAS.

8 Encrypt data

Apsara File Storage NAS allows you to enable data encryption at rest when you create a file system. This topic describes how to encrypt a file system and how data encryption at rest works.

Encrypt a file system

If you need to encrypt stored data and metadata of a file system, you can create a file system by performing the following steps.

1. Log on to the NAS console.
2. Log on to the Apsara File Storage NAS console. In the left-side navigation pane, choose **NAS**. On the page that appears, click .

   If you want to create an Extreme NAS file system, click **Package Year** or **Pay-as-you-go** in the **Extreme NAS** section.
3. On the buy page, select the **encryption type**. For more information about other parameters, see Create a file system.
4. Click **Buy Now** and follow the instructions on the page to complete the purchase.

Implementation of data encryption at rest

NAS uses the 256-bit advanced encryption standard (AES-256) to encrypt data that is stored in file systems and Key Management Service (KMS) to manage keys.

NAS uses customer master keys (CMKs) and envelope encryption to encrypt file systems. Each file system has a CMK and a data key. You can use only service keys that are provided by NAS as CMKs.

NAS encrypts data when the data is written to a file system for which data encryption at rest is enabled. When applications attempt to read data from the file system, NAS decrypts the data before sending the data to the applications. You do not need to modify your application code. The preceding operations that NAS performs are obscured.

9 Configure monitoring and alarm rules

This topic describes how to monitor Apsara File Storage NAS file systems and how to configure alarm rules in the CloudMonitor console.

Prerequisites

A file system is created. For more information, see Create a file system.

Context

You can use CloudMonitor to view performance metrics of Apsara File Storage NAS file systems and configure alarm rules for each metric. Supported performance metrics include read/write throughput, input/output operations per second (IOPS), latency, and metadata queries per second (QPS). Supported notification methods include voice calls, short message service (SMS) messages, and emails.

View performance metrics

1. Log on to the CloudMonitor console.
2. In the left-side navigation pane, choose **Dashboard** > **Dashboards**.
3. On the **Dashboards** page, select **NAS** from the Product drop-down list, and select a region and file system in the FileSystem field to view monitoring charts.

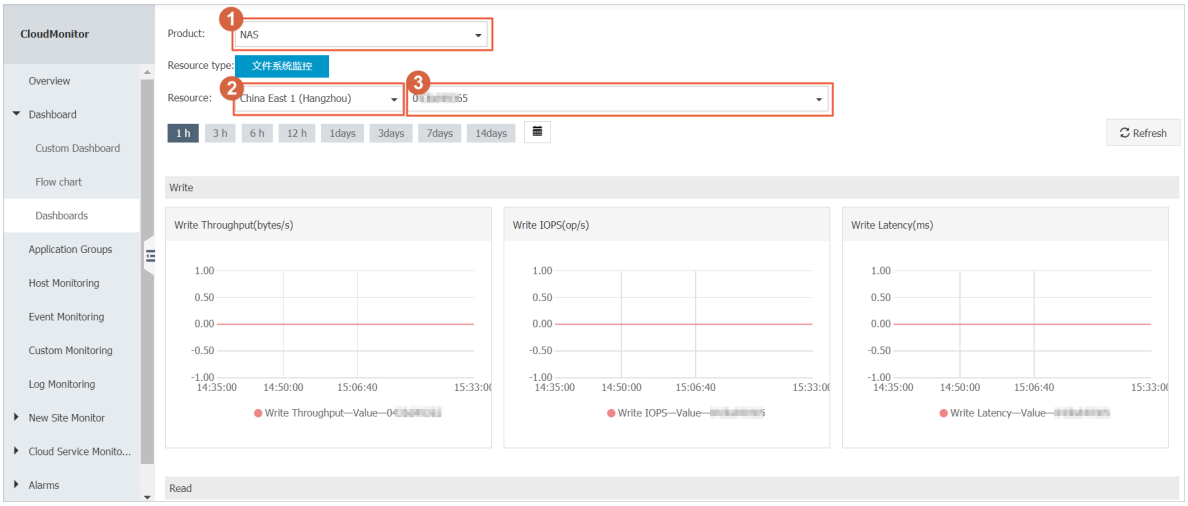

Note:

If the **No Data** message is displayed in a chart, it indicates that no request is sent from the specified file system to the associated backend server for a long period of time. If you want to simulate data to measure the write throughput of a file system, run the fio command on an Elastic Compute Service (ECS) instance on which the file

system is mounted. The following command is an example where /mnt is used as the mount directory: fio -numjobs=1 -iodepth=128 -direct=1 -ioengine=libaio -sync=1 -rw=randwrite -bs=1M -size=1G -time_based -runtime=600 -name=Fio -directory=/mnt.

Configure alarm rules

1. Log on to the CloudMonitor console.
2. Choose **Alarms** > **Alarm Rules**. On the Alarm Rules page, click **Create Alarm Rule**.
3. On the **Create Alarm Rule** page, specify the required parameters.

Parameter	Description
Related Resources	In the **Related Resource** section, complete the following settings: • From the **Product** drop-down list, select **NAS**. • From the **Resource Range** drop-down list, select **FileSystem**. • From the **Region** drop-down list, select the region where the target file system resides. • From the **FileSystem** drop-down list, select the target file system.
Set Alarm Rules	You can set multiple alarm rules based on your business requirements. For more information, see #unique_32.
Notification Methods	You can configure notification contacts and notification methods. Click **Quickly create a contact group** to create a contact group. For more information, see Create an alert contact and an alert contact group.

4. Click **Confirm** to enable the alarm rule.

   CloudMonitor sends alarms to the contacts when the value of a metric exceeds the specified threshold. This allows you to monitor the status of the file system in real time and resolve issues in a timely manner.

Monitor multiple file systems

If you want to monitor multiple file systems, you can create a group and add these file systems to the group. On the **Application Groups** page, you can group multiple file systems.

1. Log on to the CloudMonitor console.

2. Create a group.

   For more information, see #unique_34.

   a) In the left-side navigation pane, click **Application Groups**. On the page that appears, click **Create Group**.

   b) In the **Create Group** dialog box, set the required parameters.

Parameter	Description
Creation method	The method that is used to create the group. For example, select **Standard Group creation**.
Product Group Name	The name of the group.
Contact Group	The contact group that receives notifications. Click **Quickly create a contact group** to create a contact group. For more information, see Create an alert contact and an alert contact group.
MonitorAlarm	Select an alarm template from the **Select Template** drop-down list. Select an interval from the **Muted** drop-down list. Alarms will be sent at the selected interval. If you turn on the Initialize Agent Installation switch, CloudMonitor installs a CloudMonitor agent on all the instances in the application group. Then, the CloudMonitor agent can be used to collect monitoring data.
Subscribe Event Notification	After you select **Subscription Event Notification**, alarm notifications are sent if critical and warning events occur in the file systems within the group.

   c) Click **Create Group** to create a group.

3. Add a product.

   a) Click the name of the group to go to the Group Details page.

   b) In the left-side navigation pane, click **Group Resource**. On the page that appears, click **Add Product**.

   c) In the **AddResource** dialog box, select the required product and instance that you want to monitor.

   d) Click **Confirm** to add a resource.

4. View monitoring charts.

   a) Click the name of the group to go to the Group Details page.

   b) In the left-side navigation pane, click **Dashboards**, select the required file system, and view the related monitoring charts.

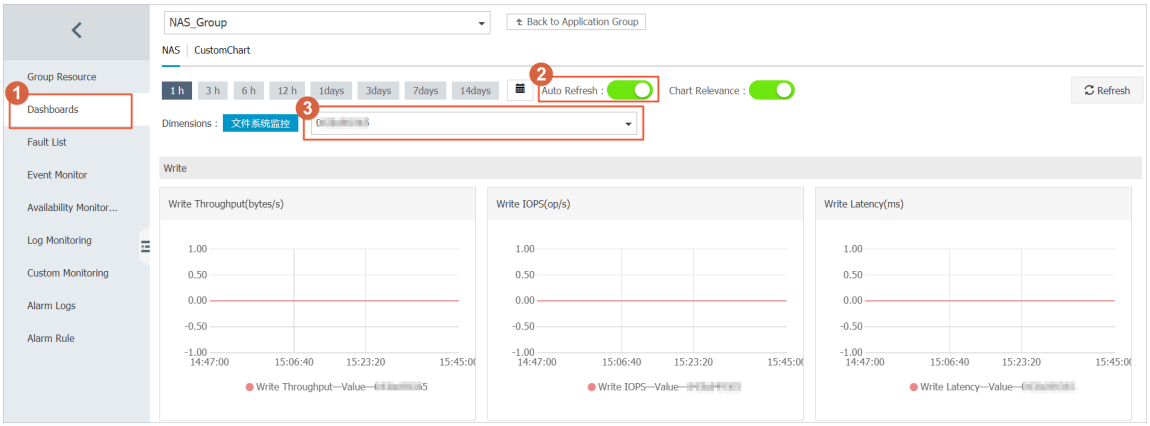

Note:

If the **No Data** message is displayed in a chart, it indicates that no request is sent from the specified file system to the associated backend server for a long period of time. If you want to simulate data to measure the write throughput of a file system, run the fio command on an Elastic Compute Service (ECS) instance on which the file system is mounted. The following command is an example where /mnt is used as the mount directory: `fio -numjobs=1 -iodepth=128 -direct=1 -ioengine=libaio -sync=1 -rw=randwrite -bs=1M -size=1G -time_based -runtime=600 -name=Fio -directory=/mnt`.

5. Configure an alarm rule.

   a) Click the name of the group to go to the Group Details page.

   b) In the left-side navigation pane, click **Alarm Rule**. On the **Threshold Value Alarm** tab, click **Create Alarm Rule**.

   c) In the dialog box that appears, click **Add Rules**, and configure the required settings. After the configuration is complete, click **OK**.

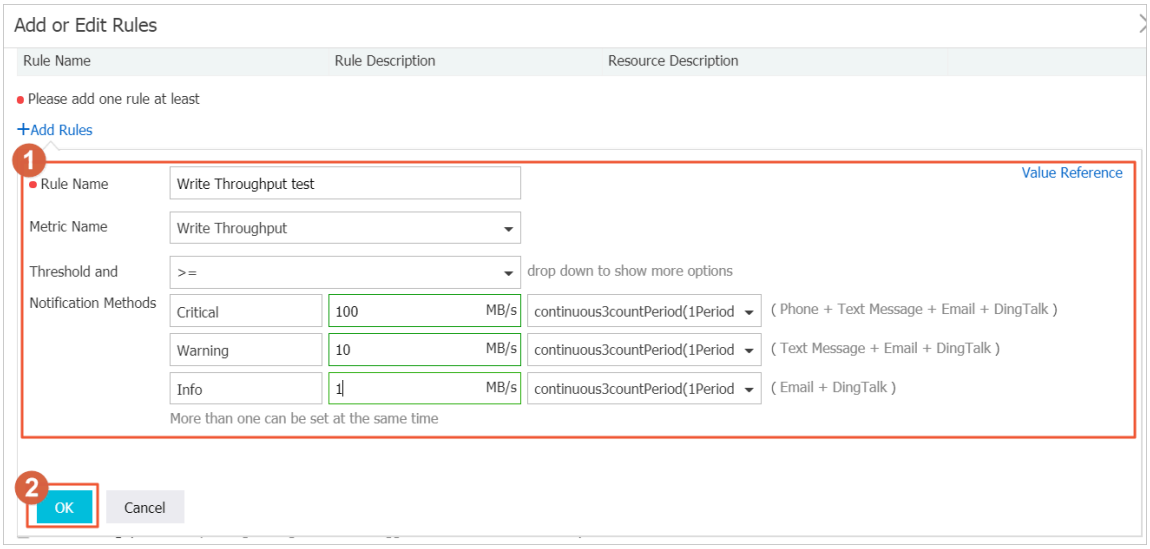

   d) Specify the **Muted** and **Contact Group** fields and click **Add**.

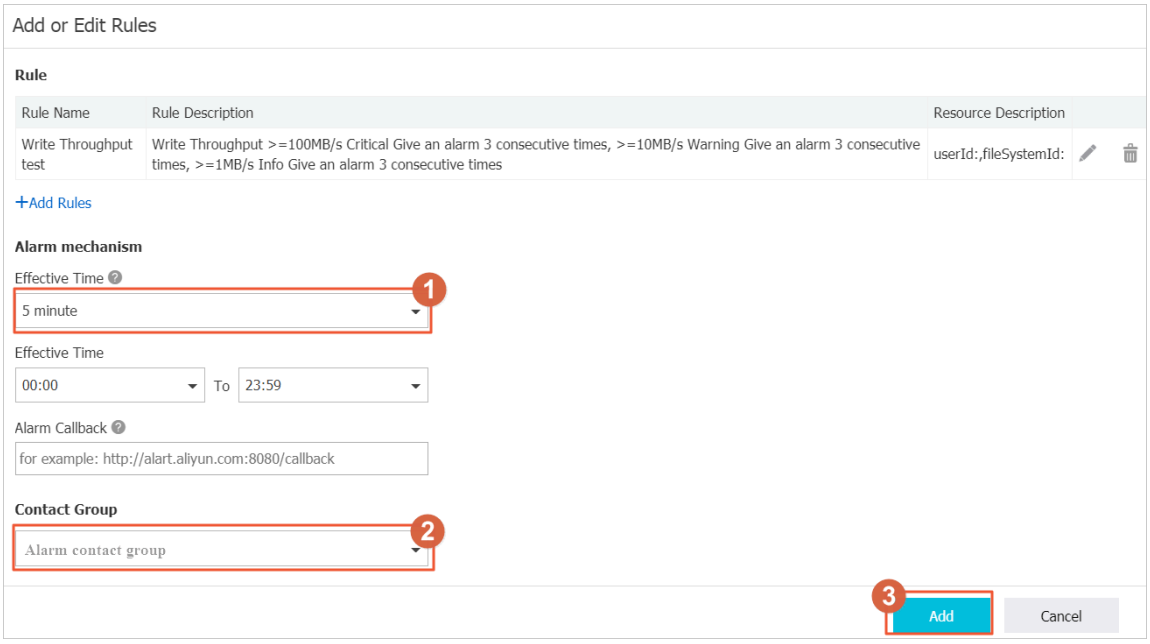

Use API operations to query metrics

You can use the following CloudMonitor API operations to query metrics of NAS file systems.

- #unique_35: queries the description of each time series metric that is available in CloudMonitor.

- #unique_36: queries the time series details of cloud services in a specified period of time.
- #unique_37: queries the latest monitoring data of an object.

The following table describes the required parameters.

Parameter	Value
Namespace	acs_nas
MetricName	IopsRead, IopsWrite, LatencyRead, LatencyWrite, QpsMeta, ThruputRead, and ThruputWriteIopsRead
Dimensions	{"userId":"xxxxxx","fileSystemId":"xxxxx"}

10 Use CPFS file systems

Cloud Paralleled File System (CPFS) is a parallel file system that offers high-performance data storage and data access.

In the CPFS console, you can perform the following operations:

- #unique_39
- #unique_40
- #unique_41
- #unique_42
- #unique_43