# Alibaba Cloud

# Apsara File Storage NAS File system mounting

Document Version: 20211230

C-J Alibaba Cloud

### 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.

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized 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 because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should 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 this document based on the "status quo", "being defective", and "existing functions" of its products and services. 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 take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, 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 and/or its affiliates 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 directly contact Alibaba Cloud for any errors of this document.

# **Document conventions**

Style	Description	Example
A Danger	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.
O Warning	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.
C) Notice	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.
⑦ Note	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 listinstanceid Instance_ID
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

# Table of Contents

1.Usage notes	06
2.Mount a file system on an ECS instance	08
2.1. Mount NAS file systems when you purchase an ECS instan	08
2.2. Mount an NFS file system on a Linux ECS instance	10
2.3. Mount an SMB file system on Windows	16
2.4. Mount an NFS file system on a Windows ECS instance	20
2.5. Mount an SMB file system on a Linux ECS instance	27
2.6. Troubleshoot and fix mount issues	32
3.Mount a file system on a pod in Kubernetes	34
3.1. Recommended mount methods	34
3.2. Use the CSI storage plug-in to mount a NAS instance	35
3.2.1. Overview	35
3.2.2. Install and upgrade the CSI plug-in	36
3.2.3. Mount a statically provisioned NAS volume	37
3.2.4. Mount a dynamically provisioned NAS volume	46
3.2.5. Set quotas on the subdirectories of NAS volumes	59
3.2.6. FAQ about NAS volumes	62
3.3. Use the Flexvolume storage plug-in to mount NAS	63
3.3.1. Use NAS volumes	64
3.3.2. Install and upgrade FlexVolume	65
3.3.3. Mount a statically provisioned NAS volume	68
3.3.4. Mount a dynamically provisioned NAS volume	73
3.3.5. Use NAS volumes for shared persistent storage	75
3.3.6. FAQ about NAS volumes	81
3.4. Mount SMB file systems to Windows containers	83
4.Mount a file system across VPCs or Alibaba Cloud accounts	86

4.1. Mount a file system across VPCs or regions	5
4.2. Enable a cross-account mount for a file system 8	7
5.Access file systems in on-premises data centers9	C
5.1. Access a NAS file system from a data center by using NAT 9	C
5.2. Access an Apsara File Storage NAS file system from a loca 9.	3
5.3. Access an SMB file system from a macOS client by using 9	5
6.Unmount a file system 10	1
6.1. Unmount a file system from a Linux ECS instance	1
6.2. Unmount a file system from an ECS instance running Win 10.	2
7.FAQ about mounting 10	3

# 1.Usage notes

This topic describes the usage notes and mounting methods of Apsara File Storage NAS file systems. Read the following sections before you mount a file system.

#### Usage notes

- If the mount target of a file system resides in a virtual private cloud (VPC), you can mount the file system on an Elastic Compute Service (ECS) instance only in the VPC. The private IP address of the ECS instance must be authorized in a rule of the permission group that is attached to the mount target.
- If the mount target of a file system resides in the classic network, you can mount the file system only on an ECS instance of the same Alibaba Cloud account. The private IP address of the ECS instance must be authorized in a rule of the permission group that is attached to the mount target. For more information, see Manage a permission group.

Onte If a mount target resides in the classic network, only the ECS instances in the classic network can access the mount target.

- You can mount a General-purpose NAS file system on an ECS instance that resides in a different zone of the same region. However, we recommend that you mount an Extreme NAS file system on an ECS instance that resides in the same zone as the file system. Otherwise, the performance of the Extreme NAS file system is compromised.
- You can create a mount target only in a VPC for an Extreme NAS file system. Extreme NAS file systems support only the Network File System (NFS) protocol.

#### Mounting methods

- The following topics describe how to mount a file system on an ECS instance.
  - Mount a file system by using the NAS console Mount NAS file systems when you purchase an ECS instance
  - Run the mount command to mount a file system
    - Mount an NFS file system on a Linux ECS instance
    - Mount an SMB file system on Windows
    - Mount an SMB file system on a Linux ECS instance
    - Mount an NFS file system on a Windows ECS instance

If you are unable to mount the file system on the ECS instance, you can use the scripts that are provided by NAS to troubleshoot issues. For more information, see Troubleshoot and fix mount issues.

- If you want to mount a file system by using Container Service for Kubernetes (ACK), we recommend that you use the methods described in Recommended mount methods. For more information, see the following topics:
  - Use the CSI driver to mount a statically provisioned NAS volume
  - Use the CSI driver to mount a dynamically provisioned NAS volume
  - Use the FlexVolume driver to mount a statically provisioned NAS volume
  - Use the FlexVolume driver to mount a dynamically provisioned NAS volume
  - Mount SMB file systems to Windows containers
- The following topics describe how to mount a file system across multiple regions or accounts by

using Cloud Enterprise Network (CEN).

- Mount a file system across VPCs or regions
- Enable a cross-account mount for a file system
- The following topics describe how to mount a file system on an on-premises machine.
  - Access an Apsara File Storage NAS file system from a local data center by using VPN Gateway
  - Access a NAS file system from a data center by using NAT Gateway
- The following topics describe how to upload data to or download data from a file system by using the SFTP client or rsync tool.
  - Upload data to and download data from an NFS file system
  - Upload data to and download data from an SMB file system
- The following topics describe how to unmount a file system.
  - Unmount a file system from a Linux ECS instance
  - Unmount a file system from an ECS instance running Windows

# 2.Mount a file system on an ECS instance2.1. Mount NAS file systems when you

### purchase an ECS instance

This topic describes how to mount one or more NAS file systems on an Elastic Compute Service (ECS) instance when you purchase the ECS instance.

#### Prerequisites

One or more NAS file systems are created. For more information, see Create a General-purpose NAS file system in the NAS console.

#### Context

After you create NAS file systems, you can purchase an ECS instance and mount the NAS file systems on the ECS instance on the buy page. This is the easiest way to mount NAS file systems. The following steps show you how to mount NAS file systems when you purchase an ECS instance. You can mount one or more NAS file systems on an ECS instance when you create the ECS instance. You can mount up to five NAS file systems on an ECS instance.

#### Mount a NAS file system

- 1. Log on to the ECS console.
- 2. Create an ECS instance. For more information, see Create an instance by using the wizard.

On the **buy page**, set the following parameters:

- **Region**: Select a region where the NAS file system resides. To achieve optimal performance, select a zone where the NAS file system resides.
- **Instance Type:** If you can specify a network type for the ECS instance in some regions, specify **VPC**. Specify other parameters based on your business requirements.
- **Image**: Select an image based on your business requirements. We recommend that you specify a Linux image of CentOS 7.6 or a Windows image of Windows Server 2019 Datacenter.
- **Storage**: Click **NAS File System**, click **Add NAS File System**, and then set the parameters. The following table describes the parameters.



Legend	Paramet er	Description
1	File system ID	<ul><li>Linux images support only NFS file systems.</li><li>Windows images support only SMB file systems.</li></ul>
2	Mount target	<ul> <li>You can use the mount target to establish a connection between the NAS file system and the ECS instance. The mount target and the ECS instance must reside in the same VPC.</li> <li>If no mount target is available, add a mount target to the file system. For more information, see Create a mount target.</li> </ul>
3	Local directory	A local directory of the ECS instance on which you can mount the NAS file system. For example, you can enter <i>/mnt</i> for a Linux image or $Z$ for a Windows image.
4	Protocol	<ul> <li>You can specify NFSv3.0 or NFSv4.0 for an NFS file system. If you do not use file locks, we recommend that you specify NFSv3.0.</li> <li>For an SMB file system, select SMB.</li> </ul>

#### Mount multiple NAS file systems

If you want to mount multiple file systems on the ECS instance, click Add NAS File System. Before you proceed, take note of the following parameters:

- Mount target:
  - All mount targets must reside in the same VPC.
  - If no mount target is available in the VPC where the ECS instance resides, add a mount target that matches the VPC. For more information, see Create a mount target.
- Local directory:
  - Each local directory must be unique.
  - You can use multilevel directories for Linux images, for example, /mnt and /mnt/sub.

**Note** If you want to mount more than five NAS file systems, submit a ticket.

#### Limits

When you mount NAS file systems on an ECS instance, take note of the following limits:

- Image: You can mount NAS file systems only on official images. You cannot mount NAS file systems on a custom image that is created by using an ECS snapshot.
- Local directory: You can specify only the root directory of NAS file systems. You cannot specify the subdirectories of NAS file systems.

**Note** If the feature cannot meet your business requirements, mount NAS file systems after the ECS instance is created. For more information, see Mount a NAS file system on an ECS instance.

#### Check the mount result

After you purchase the ECS instance, the NAS file system that you specify is automatically mounted on the ECS instance. If you restart the ECS instance, the NAS file system is also automatically mounted on the ECS instance. You can use the following methods to check the mount result:

#### • ECS Linux instances

Connect to the ECS instance and run the df -h command to view the details of the mounted NAS file system.

[root@flacking]# df -h					
Filesystem					
/dev/vda1					
devtmpfs					
tmpfs					
tmpfs					
tmpfs			3.8G		/sys/fs/cgroup
	1.0P	1.1G	1.0P	1%	/mnt
tmpfs					

In the command output, **Used** indicates the used space. **Size** indicates the maximum size of the file system. The billing of the file system is not related to the maximum size.

#### ? Note

- The settings for an automatic mount are stored in the */etc/fstab* file. You can change a local directory based on your business requirements. For more information, see Automatically mount the NFS file system.
- For information about how to connect to an ECS instance, see Guidelines on instance connection.

#### • ECS Windows instances

Connect to the ECS instance and start File Explorer. The SMB file that you specify is mounted as a network drive.



Onte The settings for an automatic mount are stored in the c:\ProgramData\Microsoft\Wi ndows\Start Menu\Programs\StartUp\my\_mount.bat file. You can change these settings based on your business requirements.

# 2.2. Mount an NFS file system on a Linux ECS instance

This topic describes how to mount a Network File System (NFS) file system on a Linux Elastic Compute Service (ECS) instance. If an NFS client is installed on a Linux ECS instance, you can manually mount the NFS file system or automatically mount it.

#### Configure a Linux ECS instance

To mount an NFS file system in a Linux operating system, you must configure a Linux ECS instance. After you configure each Linux ECS instance for the first time, you no longer need to configure the instance each time you mount a file system.

- 1. Connect to the ECS instance. For more information, see Connection methods.
- 2. Install an NFS client.
  - If CentOS, Red Hat Enterprise Linux (RHEL), or Alibaba Cloud Linux is running on the ECS instance, run the following command to install the NFS client:

sudo yum install nfs-utils

• If Ubuntu or Debian is running on the ECS instance, run the following commands to install the NFS client:

sudo apt-get update

sudo apt-get install nfs-common

3. Increase the number of concurrent NFS requests.

Run the following command to set the maximum number of concurrent NFS requests to 128. For more information, see How do I modify the maximum number of concurrent NFS requests?

```
if (lsmod | grep sunrpc); then
(modinfo sunrpc | grep tcp_max_slot_table_entries) && sysctl -w sunrpc.tcp_max_slot_tab
le_entries=128
(modinfo sunrpc | grep tcp_slot_table_entries) && sysctl -w sunrpc.tcp_slot_table_entri
es=128
fi
(modinfo sunrpc | grep tcp_max_slot_table_entries) && echo "options sunrpc tcp_max_slot
_table_entries=128" >> /etc/modprobe.d/alinas.conf
(modinfo sunrpc | grep tcp_slot_table_entries) && echo "options sunrpc tcp_slot_table_e
ntries=128" >> /etc/modprobe.d/alinas.conf
```

#### Manually mount an NFS file system

You can use the mount target of an NFS file system to mount the file system on an ECS instance.

- 1. Mount the NFS file system.
  - To mount a Capacity or Performance NAS file system, run the following command.

? Note

- We recommend that you mount a file system by using the NFSv3 protocol to ensure optimal performance.
- If your application depends on file locks and you need to edit a file on multiple Linux ECS instances at the same time, use the NFSv4 protocol to mount the file system.

#### Use the NFSv3 protocol to mount a file system:

sudo mount -t nfs -o vers=3,nolock,proto=tcp,rsize=1048576,wsize=1048576,hard,timeo=6
00,retrans=2,noresvport file-system-id.region.nas.aliyuncs.com:/ /mnt

#### Use the NFSv4 protocol to mount a file system:

```
sudo mount -t nfs -o vers=4,minorversion=0,rsize=1048576,wsize=1048576,hard,timeo=600
,retrans=2,noresvport file-system-id.region.nas.aliyuncs.com:/ /mnt
```

#### • To mount an Extreme NAS file system, run the following command:

```
sudo mount -t nfs -o vers=3,nolock,noacl,proto=tcp,rsize=1048576,wsize=1048576,hard,t
imeo=600,retrans=2,noresvport file-system-id.region.extreme.nas.aliyuncs.com:/share /
mnt
```

#### The following table lists the parameters that are used in the mount command.

Parameter	Description
Capacity or Performance NAS: file-system- id.region.nas.aliyuncs. com: / /mnt Extreme NAS: file- system- id.region.extreme.nas. aliyuncs.com:/share /mnt	<ul> <li><i>Comain name of the mount target&gt;:<name directory="" of="" shared="" the=""> <pat directory="" h="" mount="" of="" the=""></pat></name></i>. Replace the domain name of the mount target, the name of the shared directory, and the path of the mount directory with the actual values.</li> <li><i>Domain name of the mount target:</i> Log on to the NAS console and click <b>File System List</b> in the left-side navigation pane. On the page that appears, click <b>Manage</b> in the Actions column on the right of the file system. Then, click the <b>Mount Targets</b> tab to view the domain name of the mount target. For more information, see Manage mount targets.</li> <li><i>Name of the shared directory:</i> specifies the root directory / or a subdirectory. If you specify a subdirectory such as <i>/share</i>, make sure that the subdirectory exists.</li> <li><i>Path of the mount directory</i>: the root directory (/) of the Linux ECS instance or a subdirectory such as <i>/mnt</i>. If the mount directory is a subdirectory is a subdirectory exists.</li> </ul>
vers	<ul> <li>The protocol version of the file system.</li> <li>vers=3: uses NFSv3 to mount the file system.</li> <li>vers=4: uses NFSv4 to mount the file system.</li> <li>? Note <ul> <li>Capacity NAS and Performance NAS file systems support NFSv3 and NFSv4.</li> <li>Extreme NAS file systems support only NFSv3.</li> </ul> </li> </ul>

Parameter	Description
	When you mount a file system, you can select multiple mount options. These mount options are separated by commas (,). For example, you can select the following options:
	<ul> <li><i>rsize</i>: specifies the size of data blocks that the client reads from the file system. Recommended value: 1048576.</li> </ul>
	<ul> <li>wsize: specifies the size of data blocks that the client writes to the file system. Recommended value: 1048576.</li> </ul>
	<b>Note</b> We recommend that you specify the maximum value (1048576) for both the rsize mount option and the wsize mount option to prevent performance degradation.
	<ul> <li><i>hard</i>: specifies that applications stop accessing a file system when the file system is unavailable, and wait until the file system is available. We recommend that you enable the hard parameter.</li> </ul>
	<ul> <li><i>timeo</i>: specifies the period in deciseconds (tenths of a second) for which the NFS client waits before it retries to send a request. Recommended value: 600 (60 seconds).</li> </ul>
Mount options	<b>Note</b> If you need to modify the timeo mount option, we recommend that you specify 150 or a greater value. The timeo mount option is measured in deciseconds (tenths of a second). For example, the value 150 indicates 15 seconds.
	<ul> <li>retrans: specifies the number of times the NFS retries to send a request Recommended value: 2.</li> </ul>
	<ul> <li>noresvport: specifies that a new TCP port is used to ensure network continuity between the file system and the ECS instance when the network recovers from failure. We recommend that you enable the noresvport parameter.</li> </ul>
	<ul> <li>Notice</li> <li>We recommend that you do not use the <i>soft</i> mount option to</li> </ul>
	prevent data inconsistency. Use of the soft option is at your own risk.
	We recommend that you use the default values for other mount options. Performance degradation may occur due to changes in several mount options. These mount options include the size of the read or write buffer or the use of attribute caching.

2. Run the mount -1 command to view the mount result.

The command output in the following figure indicates a successful mount.



After the file system is mounted, you can run the df -h command to view the capacity of the file system.

If the file system fails to be mounted, troubleshoot the issue. For more information, see Troubleshoot and fix mount issues.

3. After the NAS file system is mounted, you can read data from and write data to the NAS file system on the Linux ECS instance.

You can access the file system the same way you access a local directory. The following figure shows an example.

[root@i7=5=6==121d==2616=7	~]#	mkdir /mnt/dir1
[root@i	~]#	mkdir /mnt/dir2
[root@i	~]#	touch /mnt/file1
[root@i	~]#	<pre>echo 'some file conent' &gt; /mnt/file2</pre>
[root@i	~]#	ls /mnt
<pre>dir1 dir2 file1 file2 tmp</pre>		

#### (Optional)

#### Automatically mount the NFS file system

When you restart the ECS instance to which the file system is mounted, the information about all the mounted file systems may be lost. To prevent the loss of such information, you can edit the */etc/fstab* configuration file in the Linux ECS instance to enable automatic mounting of NFS file system at startup.

**?** Note Before you configure automatic mounting, make sure that the preceding manual mounting is successful.

1. To mount an Extreme NAS file system, run the following command.

To mount a Capacity or Performance NAS file system, skip this step and go to Step 2.

vi /etc/systemd/system/sockets.target.wants/rpcbind.socket

Open the */etc/systemd/system/sockets.target.wants/rpcbind.socket* configuration file, and comment out the rpcbind parameter that is related to IPv6 in the following figure. Otherwise, the rpcbind service fails to run at startup.

Description=RPCbind Server Activation Socket
[Socket]
ListenStream=/var/run/rpcbind.sock
# RPC netconfig can't handle ipv6/ipv4 dual sockets
#BindIPv60nly=ipv6-only
ListenStream=0.0.0.0:111
<u>ListenDatagram=0.0.0.0:111</u>
#ListenStream=[::]:111
#ListenDatagram=[::]:111
[Install]
WantedBy=sockets.target

If you want to enable an automatic mounting in CentOS 6.x, perform the following steps:

- i. Run the chkconfg netfs on command to enable the netfs service at startup.
- ii. Open the *netconfig* file in the /etc/ directory, and comment out inet6-related information.

#						
udp	tpi_clts		inet	udp		
tcp	tpi cots ord	v	inet	tcp	-	-
#udp6	tpi_clts		inet6	udp		-
#tcp6	tpi_cots_ord	v	inet6	tcp	-	-
rawip	tpi_raw	-	inet	-	-	-
local	tpi_cots_ord		loopback			
unix	tpi_cots_ord		loopback			

2. Open the /etc/fstab configuration file.

- If you mount a Capacity or Performance NAS file system, run the following command.
  - Use the NFSv3 protocol to mount a file system:

```
file-system-id.region.nas.aliyuncs.com:/ /mnt nfs vers=3,nolock,proto=tcp,rsize=104
8576,wsize=1048576,hard,timeo=600,retrans=2,_netdev,noresvport 0 0
```

Use the NFSv4 protocol to mount a file system:

file-system-id.region.nas.aliyuncs.com:/ /mnt nfs vers=4,minorversion=0,rsize=10485
76,wsize=1048576,hard,timeo=600,retrans=2,\_netdev,noresvport 0 0

• To mount an Extreme NAS file system, run the following command:

```
file-system-id.region.extreme.nas.aliyuncs.com:/share /mnt nfs vers=3,nolock,noacl,pr
oto=tcp,noresvport,_netdev 0 0
```

? Note

- If you want to enable an automatic mounting in CentOS 6.x, run the chkconfg netfs on command to enable the netfs service at startup.
- If you want to enable an automatic mounting in Ubuntu, run the following command:

```
[ ! -f /etc/rc.local ] && echo '#!/bin/bash' > /etc/rc.local; echo "mount -
a -t nfs" >> /etc/rc.local; chmod +x /etc/rc.local
```

Add ,x-systemd.automount next to the automatic mount parameter noresvport and retain "0 0" in the command.

If you want to enable an automatic mounting in Alibaba Cloud Linux, run the following command:

```
[ ! -f /etc/rc.local ] && echo '#!/bin/bash' > /etc/rc.local; echo "mount -
a -t nfs" >> /etc/rc.local; chmod +x /etc/rc.local
```

Add ,x-systemd.automount,x-systemd.requires=systemd-resolved.service,x-systemd.after=systemd-resolved.service next to the automatic mount parameter nore syport and retain "0 0" in the command.

## For more information, see Mount parameters. The following table describes the parameters that are not included in the preceding table.

Parameter	Description
_netdev	Prevents the automatic mounting before the network is connected.
0 (the first value after noresvport)	Specifies whether to back up the file system by running the dump command. If a value is not zero, it indicates that a file system is backed up. For a NAS file system, the default value is 0.
0 (the second value after noresvport)	Indicates the order in which the fsck command checks file systems at startup. For a NAS file system, the default value is 0. It indicates that the fsck command is not run at startup.

3. Run the reboot command to restart the ECS instance.

(?) Note Before you restart the ECS instance, make sure that the manual mounting is successful. Otherwise, the ECS instance may fail to restart. If the automatic mounting is enabled, you can view the mounted NAS file systems by running the df -h command after the ECS instance is restarted.

## 2.3. Mount an SMB file system on Windows

This topic describes how to mount a Server Message Block (SMB) file system of Apsara File Storage NAS on a Windows Elastic Compute Service (ECS) instance. After you configure a Windows ECS instance, you can manually or automatically mount an SMB file system.

#### Configure a Windows ECS instance

Perform the following steps when you mount an SMB file system on a Windows ECS instance:

- 1. Connect to the ECS instance. For more information, see Connect to ECS instances.
- 2. For Windows 2016 or later, run the following command to allow anonymous access from clients:

REG ADD HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\services\LanmanWorkstation\Parameters /f /v AllowInsecureGuestAuth /t REG DWORD /d 1

- 3. Start the Workstation service.
  - i. Open the Windows Start menu. Choose All Programs > Accessories > Run or press Win+R , and enter services.msc to open the Services window.
  - ii. Make sure that the Workstation service is in the **Started** state and the startup type is **Automatic**.

By default, the Workstation service is in the Started state.

§ Services						_0
File Action View	Help					
> 🔿   🗖   🗊	o 🗈 🛛 🖬 🕨 🔲 🖬 🕨					
🔆 Services (Local)	🔕 Services (Local)	,				
	Workstation	Name 🔺	Description	Status	Startup Type	Log On As
		User Profile Service	This servic	Started	Automatic	Local System
	Stop the service	🎑 Virtual Disk	Provides m		Manual	Local System
	Pause the service Restart the service	🔍 vminit service	<failed th="" to<=""><th></th><th></th><th>Local System</th></failed>			Local System
	reading the service	🔍 Volume Shadow Copy	Manages a		Manual	Local System
		Windows Audio	Manages a		Manual	Local Service
	Description: Creates and maintains client network	🔍 Windows Audio End	Manages a		Manual	Local System
	connections to remote servers using the	🔍 Windows Color Sys	The WcsPl		Manual	Local Service
	SMB protocol. If this service is stopped,	🔍 Windows Driver Fo	Creates an		Manual	Local System
	these connections will be unavailable. If this service is disabled, any services that	Windows Error Rep	Allows erro		Manual	Local System
	explicitly depend on it will fail to start.	🔍 Windows Event Coll	This servic		Manual	Network S
		🎑 Windows Event Log	This servic	Started	Automatic	Local Service
		🔍 Windows Firewall	Windows Fi	Started	Automatic	Local Service
		🔍 Windows Font Cac	Optimizes	Started	Automatic	Local Service
		🔅 Windows Installer	Adds, modi		Manual	Local System
		🔍 Windows Managem	Provides a	Started	Automatic	Local System
		🔍 Windows Modules I	Enables ins		Manual	Local System
		🔍 Windows Remote M	Windows R	Started	Automatic	Network S
		🎑 Windows Time	Maintains d	Started	Automatic (D	Local Service
		🔍 Windows Update	Enables th	Started	Automatic (D	Local System
		🔍 WinHTTP Web Prox	WinHTTP i		Manual	Local Service
		Wired AutoConfig	The Wired		Manual	Local System
		🔍 WMI Performance	Provides p		Manual	Local System
		Workstation	Creates an	Started	Automatic	Network S
	Extended Standard					

- 4. Start the TCP/IP NetBIOS Helper service.
  - i. Open the **Network and Sharing Center** window, and click the active network connection.
  - ii. Select Properties. In the Connection Properties dialog box, double-click Internet Protocol Version 4 (TCP/IPv4). In the Internet Protocol Version 4 (TCP/IPv4) Properties dialog box, click Advanced.

iii. In the Advanced TCP/IP Settings dialog box, choose WINS > Enable NetBIOS over TCP/IP.



- iv. Open the Windows Start menu, choose All Programs > Accessories > Run or press Win+R , and enter services.msc to open the Services window.
- v. Make sure that the TCP/IP NetBIOS Helper service is in the **Started** state and the startup type is **Automatic**.

By default, the TCP/IP NetBIOS Helper service is in the Started state.

File Action View	Help						
Þ 🔿   🖬   🛅	Q 🗟 🚺 📷 🕨 🔳 🕪 👘						
Services (Local)	🔅 Services (Local)						
	TCP/IP NetBIOS Helper	Name 🔺	Description	Status	Startup Type	Log On As	
		Smart Card Remov	Allows the		Manual	Local System	
	Stop the service	SNMP Trap	Receives tr		Manual	Local Service	
	Restart the service	Software Protection	Enables th		Automatic (D	Network S	
		🔍 Special Administrati	Allows adm		Manual	Local System	
	Description:	SPP Notification Ser	Provides S		Manual	Local Service	
	Provides support for the NetBIOS over TCP/IP (NetBT) service and NetBIOS	SSDP Discovery	Discovers		Disabled	Local Service	
	name resolution for dients on the	🔍 System Event Notifi	Monitors s	Started	Automatic	Local System	
	network, therefore enabling users to	Cask Scheduler	Enables a	Started	Automatic	Local System	_
	share files, print, and log on to the network. If this service is stopped, these	TCP/IP NetBIOS He	Provides s	Started	Automatic	Local Service	
	functions might be unavailable. If this	California Telephony	Provides T		Manual	Network S	
	service is disabled, any services that	🥋 Thread Ordering Se	Provides or		Manual	Local Service	
	explicitly depend on it will fail to start.	🔍 UPnP Device Host	Allows UPn		Disabled	Local Service	
		🔍 User Profile Service	This servic	Started	Automatic	Local System	
		🔍 Virtual Disk	Provides m		Manual	Local System	
		🔍 Volume Shadow Copy	Manages a		Manual	Local System	
		Windows Audio	Manages a		Manual	Local Service	
		🔍 Windows Audio End	Manages a		Manual	Local System	
		Windows Color Sys	The WcsPl		Manual	Local Service	
		Windows Driver Fo	Creates an		Manual	Local System	
		Windows Error Rep	Allows erro		Manual	Local System	
		Windows Event Coll	This servic		Manual	Network S	
		🔍 Windows Event Log	This servic	Started	Automatic	Local Service	
	1	Windows Firewall	Windows Fi	Started	Automatic	Local Service	
	Extended Standard						

#### Manually mount an SMB file system

1. Run the mount command to mount an SMB file system.

Open the command prompt and run the following command to mount an SMB file system:

net use Z: \\file-system-id.region.nas.aliyuncs.com\myshare

Command syntax: net use <Letter of the destination drive> \\<Domain name of the mount target>\m yshare .

• Letter of the destination drive: the letter of the drive on which you mount an SMB file system. Replace this parameter with the letter of the actual destination drive.

**?** Note The letter of the destination drive must be different from the existing drive letters.

- Domain name of the mount target: The domain name of the mount target is automatically generated when the mount target is created. Replace this parameter with the actual domain name of the mount target. To obtain the domain name of the mount target, perform the following steps: Log on to the NAS console, find the file system, and then click Manage. On the details page that appears, the domain name of the mount target is displayed.
- myshare: the name of an SMB share. You cannot change the name.

(?) Note For Windows 2019 or later, we recommend that you use the Powershell command New-SmbGlobalMapping to mount the SMB file system. Run the following command to mount the SMB file system:

New-SmbGlobalMapping -LocalPath z: -RemotePath \\file-system-id.region.nas.aliyuncs.c om\myshare -Persistent \$true

If an AD domain controller is installed on the ECS instance, you must pass the identity verification when you run the command. You can enter a valid identity in the AD domain.

#### 2. Confirm that the SMB file system is mounted.

Run the net use command to check the mount result.

The command output in the following figure indicates a successful mount.

	Administrato ctions will	or>net use be remembere	d.
Status	Local	Remote	Network
ок	D:	<u>\\6</u>	.nas.aliyuncs.com\myshare Microsoft Windows Network
The comma	nd completed	l successfull	y.

After the SMB file system is mounted, you can read data from and write data to the NAS file systems on the ECS instance.

#### Automatically mount an SMB file system

Before you configure automatic mounting, make sure that the preceding manual mounting is successful.

(?) Note In Windows operating systems, most system calls including Services and Scheduled Tasks are performed by using the SYSTEM account. However, the following mount process is performed by using a logon user account. If you need to mount an SMB file system by using the SYSTEM account, see .

1. Open the command prompt and run the following command to configure the auto\_mount.bat

#### script:

echo %HOMEPATH%\mount.bat > auto\_mount.bat

2. Run the following three commands to enable the auto\_mount.bat script to automatically run after user logon and grant the read and execute permissions to other users:

MOVE auto\_mount.bat "c:\ProgramData\Microsoft\Windows\Start Menu\Programs\StartUp\auto\_mo unt.bat"

icacls "c:\ProgramData\Microsoft\Windows\Start Menu\Programs\StartUp\auto\_mount.bat" /gra
nt everyone:rx

REG ADD HKEY\_LOCAL\_MACHINE\Software\Microsoft\Windows\CurrentVersion\Run /f /v MyMount /t REG SZ /d "c:\ProgramData\Microsoft\Windows\Start Menu\Programs\StartUp\auto mount.bat"

#### 3. Run the following command to configure the mount.bat script:

echo net use z: \\file-system-id.region.nas.aliyuncs.com\myshare > "%HOMEPATH%\mount.bat"

Replace file-system-id.region.nas.aliyuncs.com with the actual domain name of the mount target.

**?** Note If the SMB file system supports ADs and ACLs, you can run the following command to configure the script. Then, you can mount the SMB file system as a domain user rather than a Windows logon user:

echo net use z: \\file-system-id.region.nas.aliyuncs.com\myshare /user:user@domain pa
ssword > "%HOMEPATH%\mount.bat"

Replace file-system-id.region.nas.aliyuncs.com with the actual domain name of the

mount target, replace <u>user@domain</u> with the actual domain username, and replace <u>passwor</u> a with the actual domain user password.

4. Restart the ECS instance.

Run the net use command to check the mount result.

# 2.4. Mount an NFS file system on a Windows ECS instance

If you need to share data among instances that run different operating systems, you can mount an NFS file system on a Windows ECS instance. This way, you can upload data to and download data from the NFS file system. The topic describes how to mount an NFS file system on an ECS instance that resides in a virtual private cloud (VPC). Windows Server 2012 R2 is used in this example.

#### Install the NFS client

- 1. Connect to the ECS instance. For more information, see Connection methods.
- 2. Start Server Manager.
- 3. Choose Manage > Add Roles and Features.
- 4. Follow the Add Roles and Features Wizard to install the NFS client.
  - i. In the Server Roles step, choose File and Storage Services > File and iSCSI Services and select Server for NFS.
  - ii. In the Features step, select Client for NFS.
- 5. Restart the ECS instance.

6. Start the **command prompt** and run the mount command.

The following command output indicates that the NFS client is installed.

C:\User	s\Administrator	Rount
Local	Remote	Properties
C:\User	s\Administrator	>_

#### Manually mount the NFS file system

1. On a Windows client, run the following command to mount the NFS file system:

```
mount -o nolock -o mtype=hard -o timeout=60 \\file-system-id.region.nas.aliyuncs.com\!
Z:
```

In the preceding command, file-system-id.region.nas.aliyuncs.com indicates the domain name of the mount target. You need to replace the domain name with the actual value.

(?) Note If you mount a subdirectory of a NAS file system, the mount may fail. We recommend that you do not mount a subdirectory of a NAS file system. For more information, see How do I resolve the invalid device error that is returned when I try to rename a file on the Windows NFS client?

2. Run the mount command to check whether the NFS file system is mounted.

If the NFS file system is mounted, the command output shows that the value of mount is hard, the value of locking is no, and the value of timeout is greater than or equal to 10. Otherwise, the NFS file system is not mounted.

:·Users Administrator>mount						
_ocal	Remote	Properties				
1:	Ntert-and & 72-ayeff, en-shanghai, nas"	UID=0, GID=0 rsize=1048576, wsize=1048576 mount=hard, timeout=10.0 retry=1, locking=no fileaccess=755, lang=GB2312-80 casesensitive=no sec=sys				

3. Double-click the This PC icon to view the shared file system.

Create files and folders in the shared file system to check whether you can manage the shared file system.



#### Automatically mount the NFS file system

1. Enter the following content in the nas\_auto.bat script file and save the file to the corresponding disk path.

```
Example: mount -o nolock -o mtype=hard -o timeout=60 \\file-system-id.region.nas.aliyuncs
.com\! Z:
```

You need to replace the drive letter (Z:) and the domain name (file-systemid.region.nas.aliyuncs.com) with their actual values. For more information, see Automatically mount the NFS file system.

**?** Note If a success message is displayed but the NAS disk does not appear, move the nas\_auto.bat file to the *C:\ProgramData\Microsoft\Windows\StartMenu\Programs\StartUp* directory. In this case, you do not need to create a task in the Task Scheduler window.

2. Create a scheduled task.

i. Open the **Control Panel** and choose **Administrative Tools > Task Scheduler**.

ii. In the Task Scheduler window, choose Actions > Create Task.



iii. On the General tab, enter the name of the task, and select Run whether user is logged on or not and Run with highest privileges.

General Trig	gers Actions Conditions Settings
Na <u>m</u> e:	nas
Location:	l l
Author:	1 - A - A - THE
<u>D</u> escription:	
-Security opt	ions
When runn	ing the task, use the following user account:
10.00	Change User or Group
Run only	y when user is logged on
Run whee	ther user is logged on or not
📃 Do n	ot store <u>p</u> assword. The task will only have access to local computer resources.
🔽 Run w <u>i</u> th	n highest privileges
Hidd <u>e</u> n	Configure for: Windows Vista™, Windows Server™ 2008

iv. Click the Triggers tab. Click New Trigger. Select At startup from the Begin the task dropdown list. In the Advanced settings section, select Enabled. Click OK.

New Trigger	
Begin the task: Settings	At log on 🔹
Any user	
Specific us	sen: HZ\wb-yjp354868 Change User
Advanced setti	ngs
Delay task t	for: 15 minutes 👻
🔲 Repeat task	k every: 1 hour v for a duration of: 1 day v
Stop	all running tasks at end of repetition duration
📄 Stop task if	f it runs longer than: 3 days 👻
Activate:	2018/11/21 🔍 10:14:56 📄 Synchronize across time zones
Expire:	2019/11/21 🔍 🔹 🖾 Synchronize across time zones
Enabled	
	OK Cancel

v. Click the Actions tab. Click New Action. Select Start a program from the Action dropdown list, select the *nas\_auto.bat* file in the Program/script field, and then click OK.

New Action	x
You must specify what action this task will perform.	
Act <u>i</u> on: Start a program	]
Settings	
Program/script:	
B <u>r</u> owse	
Add arguments (optional):	
Start in (optional):	
OK Cancel	

vi. Click the Conditions tab. Select Start only if the following network connection is available. Select Any connection from the drop-down list under Start only if the following network connection is available.

Create Task		
ieneral Triggers Actions Conditions Settings Specify the conditions that, along with the trigger, deter run if any condition specified here is not true.	mine whether the task sho	uld run. The task will not
Idle Start the task only if the <u>c</u> omputer is idle for:	10 minutes	-
W <u>a</u> it for idle for:	1 hour	-
Stop if the comput <u>er</u> ceases to be idle		
Restart if the idle state resumes		
Stop if the computer switches to battery power           Wake the computer to run this task           Network           Z Start only if the following network connection is avail	lable:	
Any connection		
		OK Cancel

vii. Click the Settings tab. Select If the running task does not end when requested, force it to stop. Select Do not start a new instance from the drop-down list under If the task is already running, then the following rule applies.

(b) Create Task	
General Triggers Actions Conditions Settings	
Specify additional settings that affect the behavior of the task.	
Allow task to be run on demand	
Run task as soon as possible after a scheduled start is missed	
If the <u>t</u> ask fails, restart every:	1 minute 👻
Attempt to <u>r</u> estart up to:	3 times
Stop the tas <u>k</u> if it runs longer than:	3 days 👻
If the running task does not end when requested, force it to s	top
If the task is not scheduled to run again, <u>d</u> elete it after:	30 days 👻
If the task is already running, then the following rule applies:	
Do not start a new instance 🔹	
	OK Cancel

- viii. Click OK.
- ix. Restart the ECS instance to check whether the scheduled task is created.

The following example shows that the scheduled task is created.

🕑 Task Scheduler		
File Action View Help		
🗢 🔿 🙍 🖬 👔		
<ul> <li>Task Scheduler (Local)</li> <li>Task Scheduler Library</li> <li>Lenovo</li> <li>Microsoft</li> <li>OfficeSoftwareProtect</li> <li>Symantec Endpoint Pi</li> <li>TVT</li> <li>WPD</li> </ul>	Name       Status       Triggers         Image: Provide the status       At log on of any user         Image: Provide the status       At log on of any user         Image: Provide the status       At log on of any user         Image: Provide the status       At log on of any user         Image: Provide the status       At log on of any user         Image: Provide the status       At log on of any user         Image: Provide the status       At log on of any user         Image: Provide the status       At system startup         Image: Provide the status       Conditions       Settings       History         Name:       Name:       Name       Provide the status       Provide the status         Image: Provide the status       Image: Provide the status       Provide the status       Provide the status       Provide the status         Image: Provide the status       Image: Provide the status       Image: Provide the status       Provide the status       Provide the status       Provide the status         Image: Provide the status       Image: Provide the status       Image: Provide the status       Provide the status	Actions         Task Sc       ▲         Image: Cr       ▲         Image: Cr       ▲         Image: Cr       ▲         Image: Dis       ■         Image: Dis <td< td=""></td<>
		• • • •
L		

#### FAQ

If the file handle error message appears, check the following registry keys.

**Note** If you cannot find the Locking, AnonymousGID, and AnonymousUID registry keys, you need to follow the format required by the Windows registry to create these keys.

Choose HKEY\_LOCAL\_MACHINE > SOFT WARE > Microsoft > Client ForNFS > Current Version > Users > Default > Mount, create a key named Locking, and set the value of this key to 1.

19 A A A A A A A A A A A A A A A A A A A		megion,	Luitor	
File Edit View Favorites Help				
Microsoft      M	Name (Default) Contains Defaults MountType ReadBuffer Retransmissions WriteBuffer	Type REG_SZ REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD	Data           (value not set)         0x0000000 (0)           0x0000000 (0)         0x0000000 (1)           0x00000000 (1)         0x0000000 (1048576)           0x00000000 (1)         0x00000000 (8)           0x00100000 (1048576)         0x00100000 (1048576)	

Create the following registry keys to configure the GID and UID.

- Navigate to the Default path: HKEY\_LOCAL\_MACHINE > SOFT WARE > Microsoft > ClientForNFS > CurrentVersion > Default.
- Right-click a blank area on the right side of the Registry Editor window, choose New > DWORD (32-bit) Value, and create the following registry keys.
  - AnonymousGID: Set the value of the key to 0.
  - AnonymousUID: Set the value of the key to 0.

⊿ - 🌽 Microsoft	^	Name	Туре	Data
	=	ab (Default) AnonymousGid AnonymousUid CacheBlocks DeleteSymLinks FirstContact MaxNfsUser MountType Protocols Retransmissions Timeout UseReservedPorts	REG_SZ REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD REG_DWORD	(value not set) 0x00000000 (0) 0x00000000 (0) 0x00000000 (0) 0x00000001 (1) 0x00000000 (32) 0x00000001 (1) 0x00000001 (1) 0x00000000 (8) 0x00000001 (1)

3. Restart the ECS instance.

4. Run the following command to mount the NAS file system:

```
mount -o nolock -o mtype=hard -o timeout=60 \\file-system-id.region.nas.aliyuncs.com\!
Z:
```

In the preceding command, file-system-id.region.nas.aliyuncs.com indicates the domain name of the mount target. You need to replace the domain name with the actual value.

5. Run the mount command to check the UID and GID.

If the NAS file system is mounted, the command output shows that the value of mount is hard, the value of locking is no, and the value of timeout is greater than or equal to 10. Otherwise, the NAS file system is not mounted.



# 2.5. Mount an SMB file system on a Linux ECS instance

This topic describes how to mount a Server Message Block (SMB) file system on a Linux Elastic Compute Service (ECS) instance. This topic also describes how to read data from and write data to the SMB file system.

#### Prerequisites

**Notice** We recommend that you mount an NFS file system on a Linux ECS instance. Linux is not highly compatible with the SMB protocol. Therefore, we recommend that you mount an SMB file system on a Linux ECS instance only if you want to share data across operating systems.

• An ECS instance is available in the region where you want to create a file system. For more information, see Create an ECS instance.

One of the following Linux distributions is run on the ECS instance. Unless otherwise specified, this topic applies only to the following Linux distributions:

- Cent OS 7.6 64-bit (3.10.0-957.5.1.el7.x86\_64)
- Ubunt u 18.04 64-bit (4.15.0-48-generic)
- Debian 9.9 64-bit (4.9.0-9-amd64)
- SUSE Enterprise Server 12 SP2 64-bit (4.4.74-92.35-default)
- openSUSE 42.3 64-bit (4.4.90-28-default)
- Alibaba Cloud Linux (4.19.34-11.al7.x86\_64)
- CoreOS (4.19.43-coreos VersionID=2079.4.0)

(?) Note If you use a version of Linux that is not in this list, errors may occur on the SMB client. If you use an unsupported version, Alibaba Cloud does not guarantee the reliability of SMB file systems.

- An SMB file system is created. For more information, see Create a General-purpose NAS file system in the NAS console.
- A mount target is created. For more information, see Create a mount target.
- The network is available.
  - The Linux ECS instance and the SMB file system reside in the same virtual private cloud (VPC).
  - The IP address of the Linux ECS instance is in the whitelist of the SMB file system and the ECS instance can access the SMB file system.
  - TCP port 445 is enabled for the Linux ECS instance to access the SMB file system.
     If port 445 is disabled, you must add a rule to a security group of the ECS instance for the port. For more information, see Add security group rules.

#### Install the cifs-utils package

Install the cifs-utils package on a Linux ECS instance.

If you are using Ubuntu or Debian, use the apt-get package manager to install the cifs-utils package.

```
sudo apt-get update
```

sudo apt-get install cifs-utils

• If you are using Red Hat Enterprise Linux (RHEL), CentOS, or Alibaba Cloud Linux, use the Yellowdog Updater, Modified (YUM) package manager to install the cifs-utils package.

sudo yum install cifs-utils

• If you are using openSUSE or SUSE Linux Enterprise Server 12 Service Pack 2 (SLES 12 SP2), use the Zypper or Yet another Setup Tool (YaST) package manager to install the cifs-utils package.

sudo zypper install cifs-utils

Run the sudo yast2 command, choose Software > Software Management, and then install the c ifs-utils package.

- If you are using CoreOS, install the cifs-utils package by performing the following steps:
  - i. Configure Security-Enhanced Linux (SELinux).

sed -i 's/SELINUXTYPE=mcs/SELINUXTYPE=targeted/' /etc/selinux/config

ii. Compile the cifs-utils package on a CoreOS ECS instance. You can run the following command to create a Fedora container and compile the cifs-utils package. You can also download the cifs-utils package that Alibaba Cloud provides for CoreOS, and then copy the package to the /tmp or /bin directory.

```
$ docker run -t -i -v /tmp:/cifs fedora /bin/bash
fedora # yum groupinstall -y "Development Tools" "Development Libraries"
fedora # yum install -y bzip2
fedora # curl https://download.samba.org/pub/linux-cifs/cifs-utils/cifs-utils-
6.9.tar.bz2 --output cifs-utils-6.9.tar.bz2;
fedora # bunzip cifs-utils-6.9.tar.bz2; && tar xvf cifs-utils-6.9.tar
fedora # cd cifs-utils-6.9; ./configure && make
fedora # cp mount.cifs /cifs/
fedora # exit
```

#### Mount a file system

- 1. Log on to the Linux ECS instance as a root user or a sudo-enabled user.
- 2. Run the following command to mount the file system:

```
mount -t cifs //xxx-crf23.eu-west-1.nas.aliyuncs.com/myshare /mnt -o vers=2.0,guest,uid
=0,gid=0,dir_mode=0755,file_mode=0755,mfsymlinks,cache=strict,rsize=1048576,wsize=10485
76
```

**Command syntax:** mount -t cifs //<Domain name of the mount target>/myshare <Path of the s hared directory> -o <Mount options>

Parameter	Description	
File system type	Specify the -t cifs parameter in the command for the SMB file system.	
Domain name of the mount target	Specify the domain name of the mount target in the command. The domain name is automatically generated when you create the mount target. For more information, see Manage mount targets.	
myshare	The name of the shared directory for the SMB file system. You cannot change the name after you specify this parameter.	
Path of the shared directory	The path of the directory on which the SMB file system is mounted. For example, you can specify <i>/mnt/sharepath</i> .	

Parameter	Description
	Specify the required mount options by adding the $-\circ$ argument in the mount command:
	• vers: specifies the version of the SMB protocol. Specify 2.0 or 3.0 for the option.
	• guest: specifies the identity that you want to use to mount the file system. You must use the guest identity that is authenticated based on the New Technology Lan Manager (NTLM) protocol.
	• rsize: specifies the maximum size of a data packet that the SMB client can read from the SMB file system. In most cases, set this option to 1048576 (1 MB).
	<ul> <li>wsize: specifies the maximum size of a data packet that the SMB client can write to the SMB file system. In most cases, set this option to 1048576 (1 MB).</li> </ul>
	Specify the following mount options by adding the> parameter:
	<ul> <li>uid: specifies the user to which the files in the file system belong after a successful mount. The default value of uid is 0.</li> </ul>
	<ul> <li>gid: specifies the user group to which the files in the file system belong after a successful mount. The default value of gid is 0.</li> </ul>
	<ul> <li>dir_mode: specifies the read, write, and execute permissions that you want to grant to the user on the specified directories. The value must start with 0, for example, 0755 and 0644. The default value of dir_mode is 0755.</li> </ul>
	<ul> <li>file_mode: specifies the read, write, and execute permissions that you want to grant to the user on files. The value must start with 0, for example, 0755 and 0644. The default value of file_mode is 0755.</li> </ul>
Mount options	<ul> <li>mfsymlinks: specifies whether symbol links are supported.</li> </ul>
	° cache:
	If you set this option to strict, caching is enabled for the SMB client. The default value of cache is strict.
	If you set this option to none, caching is disabled for the SMB client.
	<ul> <li>atime relatime: If file access time does not affect your business, we recommend that you do not set this option to atime. The default value of this option is relatime.</li> </ul>
	⑦ Note
	<ul> <li>An authorized administrator of the Linux ECS instance has full control over the SMB file system.</li> </ul>
	<ul> <li>If you want to view the details about a mount target, you can run the mount   grep cifs command.</li> </ul>
	<ul> <li>If you are using an inapplicable Linux distribution, we recommend that you use a Linux distribution whose kernel version is later than 3.10.0-514. If the kernel version of the distribution is 3.7 or earlier, you must set the cache option to strict. You can run the uname -a command to view the kernel version.</li> </ul>

3. Run the mount -1 command to view the mount result.

The output in the following figure indicates that the file system is mounted.

debugfs on /sys/kernel/debug type debugfs (rw,relatime)		
hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime)		
mqueue on /dev/mqueue type mqueue (rw,relatime)		
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw,relatime)		
tmpfs on /run/user/0 type tmpfs (rw,nosuid,nodev,relatime,size=800920k,mode=700)		
// cn-hangzhou.nas.aliuuncs.com/wushare on /mnt type cifs (rw,relatime,vers=2.1,sec=none,cache=strict,domain=,ui		
d=0,noforceuid,gid=0,noforcegid,addr=,file_mode=0755,dir_mode=0755,soft,nounix,serverino,mapposix,rsize=1048576,ωs		
ze=1048576.echo_interval=60.actimeo=1)		

4. Read data from and write data to the file system.

You can access the file system the same way you access a local directory. The following figure shows an example.

[root@i7=5c6omig=121d=2616a2	~]# mkdir /mnt/dir1
[root@i	~]# mkdir /mnt/dir2
[root@i	~]# touch /mnt/file1
[root@i	~]# echo 'some file conent' > /mnt/file2
[root@i	~]# ls /mnt
dir1 dir2 file1 file2 tmp	_

#### Scenarios

To ensure optimal performance of the file system, you can specify mount options based on specific scenarios. This section lists scenarios and the mount options that are suitable for each scenario:

• Shared access to a file system from multiple Linux ECS instances

Multiple Linux ECS instances share access to a file system and no access control is required. In this scenario, you can use an authorized administrator of each ECS instance to mount the file system on the ECS instances. The following command shows an example:

mount -t cifs //smbfs.hangzhou-g.aliyun.com/myshare /mnt/sharepath -o vers=2.1,guest,mfs
ymlinks

- Shared access from multiple Linux ECS instances to a home directory Multiple Linux ECS instances share access to a home directory and you need to control access to the home directory. You can set the uid, gid, dir\_mode, and file\_mode options in the mount command to manage permissions.
- Shared access to a file system from multiple Linux ECS instances that function as web servers You can install web server applications such as Apache HTTP Server on multiple Linux ECS instances and use an SMB file system as shared file storage.
  - ? Note
    - The SMB file system provides shared access, horizontal scalability, and high availability. When users access small files in the SMB file system, the performance of the SMB file system may be compromised. This occurs because the mechanism of SMB file systems is different from the mechanism of local disks. In this scenario, we recommend that you store shared files in the SMB file system and other files in local disks to ensure optimal performance.
    - In most cases, web server applications have heavy workloads. You can enable the acceleration feature for the applications to process heavy workloads. You can contact NAS technical support to enable this feature.
- Shared access from both a Windows ECS instance and a Linux ECS instance to a file system A Windows ECS instance and a Linux ECS instance share access to an SMB file system. In this scenario, you must set the cache option to strict or use the default value of this option when you mount the file system on the Linux ECS instance.

For information about how to resolve issues, see Troubleshoot issues that may occur when you access an SMB file system from a Linux ECS instance.

### 2.6. Troubleshoot and fix mount issues

This topic describes how to troubleshoot and fix issues that may occur when you mount Apsara File Storage NAS file systems.

#### Mount an NFS file system on an ECS instance running Linux

• Enable automatic troubleshooting by using a script

You may fail to mount an NFS file system on an ECS instance running Linux due to several different reasons. You can use the following script to troubleshoot a mount issue and identify the root cause.

- i. Log on to a Linux ECS instance on which you fail to mount a file system.
- ii. Use the following commands to download and run the check\_alinas\_nfs\_mount.py script. Then, you can follow instructions provided by the script to fix mount issues.

wget -N https://code.aliyun.com/nas\_team/nas-client-tools/raw/master/linux\_client/che ck\_alinas\_nfs\_mount.py -P /tmp/

```
python2.7 /tmp/check_alinas_nfs_mount.py file-system-id.region.nas.aliyuncs.com:/ /mn
t
```

In the preceding command, file-system-id.region.nas.aliyuncs.com specifies the mount target of the NAS file system, the forward slash (/) following the mount target specifies the root directory of the NAS file system, and /mnt specifies a local mount point that resides on the Linux ECS instance. You need to replace these example parameters based on your business requirements. After all issues are fixed, a specific mount command is displayed. A prompt also appears indicating that troubleshooting for the issue is complete.

(?) Note If several questions appear when the script is running, we recommend that you log to the Alibaba Cloud console and confirm the answers to the questions. After the answer to each question is confirmed, you can click **Yes** or **No** to continue running the script and find more issues.

iii. Copy and run the mount command to enable the mount of a file system.

• Troubleshoot and fix other issues

Several errors prompted by the mount command cannot be fixed by using the script. You can use the following methods to fix these errors.

Failed to mount the subdirectory of a file system
 Error message: mount.nfs: access denied by server while mounting xxxx.nas.aliyuncs.com:/<dir>

**?** Note If an error message showing "Permission denied" appears, you can use the script to troubleshoot the issue.

An error occurs when you attempt to mount a subdirectory of a NAS file system on an ECS instance but the subdirectory does not exist. You can first mount the root directory of the NAS file system. After the root directory is mounted, you can create a subdirectory on the NAS file system and mount the subdirectory again. Failed to mount a file system on two instances with duplicate names
 Error message: mount.nfs: Operation not permitted. This error occurs when you mount an NFSv4
 file system. However, the mount is successful if the file system complies with NFSv3.
 For several kernel versions of Linux, an error may occur in the following scenario: You attempt to
 mount an NFSv4 file system on an ECS instance with the same name as that of another ECS
 instance on which the file system is mounted. You can perform the following steps to fix the error:

a. Use the following command on the ECS instance on which you fail to mount a file system.

echo 'install nfs /sbin/modprobe --ignore-install nfs nfs4\_unique\_id=`cat /sys/clas s/dmi/id/product\_uuid`' >> /etc/modprobe.d/nfs.conf

- b. Restart the ECS instance during off-peak hours.
   You can also unmount all available NFS file systems and use the rmmod command to uninstall the nfsv4 and nfs kernel modules.
- c. Re-mount the NFS file system.

#### Mount an SMB file system on an ECS instance running Windows

• Enable automatic troubleshooting by using a script

You may fail to mount an SMB file system on an ECS instance running Windows due to several different reasons. You can use the following script to troubleshoot a mount issue and find the root cause.

- i. Log on to a Windows ECS instance on which you fail to mount a file system.
- ii. Use the following commands to download and run the check\_alinas\_nfs\_mount.py script. Then, you can follow instructions provided by the script to fix mount issues.

wget https://code.aliyun.com/nas\_team/nas-client-tools/raw/master/windows\_client/alin
as\_smb\_windows\_inspection.ps1 -OutFile alinas\_smb\_windows\_inspection.ps1

```
.\alinas_smb_windows_inspection.ps1 -MountAddress abcde-123.region-id.nas.aliyuncs.co m -Locale zh-CN
```

In the preceding command, abcde-123.region-id.nas.aliyuncs.com specifies the domain name of a mount target. You need to replace the domain name with a domain name that is specific to your environment.

• Troubleshoot and fix other issues

For more information about how to troubleshoot and fix issues that occur when you mount a file system on a Windows ECS instance, see Troubleshoot SMB mount failures. You can find the corresponding solution to each error code.

#### Mount an SMB file system on an ECS instance running Linux

Apsara File Storage NAS allows you to mount an SMB file system on a Linux ECS instance. If you fail to mount an SMB file system, see Troubleshoot issues that may occur when you access an SMB file system from a Linux ECS instance.

# 3.Mount a file system on a pod inKubernetes3.1. Recommended mount methods

This topic describes the recommended methods that you can use to mount NAS file systems on Container Service for Kubernetes (ACK). You can mount file systems by using storage plug-ins such as FlexVolume and Container Storage Interface (CSI). You can also mount file systems on Windows containers.

Notice We recommend that you mount a NAS file system by using one of the following methods. If you mount a NAS file system by using other methods, the file system may be exposed to stability risks. The NAS technical team cannot estimate the impact of these risks. You are solely responsible for all the losses and consequences that may arise from the risks.

#### Storage plug-ins

ACK supports FlexVolume and CSI. FlexVolume and CSI support multiple storage services such as Apsara File Storage NAS, Cloud Paralleled File System (CPFS), Object Storage Service (OSS), and Block Storage. FlexVolume and CSI provide flexible parameters to improve user experience and reduce O&M complexity.

We recommend that you use FlexVolume or CSI if you want to mount file systems on ACK clusters or self-managed Kubernetes clusters. Take note of the following items:

- If you want to mount a file system on a Kubernetes cluster that is newly created, we recommend that you use CSI.
- If you want to mount a file system on an existing Kubernetes cluster, we recommend that you use a storage plug-in that is already installed.
- You cannot use both plug-ins on the same ACK cluster.
- You cannot change the plug-in from FlexVolume to CSI for an ACK cluster.

For information about the differences between FlexVolume and CSI, see Differences between the CSI and FlexVolume plug-ins.

#### Use CSI to mount file systems

For information about how to mount file systems by using CSI, see Overview.

CSI supports the following methods:

- Mount file systems as static persistent volumes (PVs). For more information, see Mount a statically provisioned NAS volume.
- Mount file systems as dynamic PVs. For more information, see Mount a dynamically provisioned NAS volume.

For more information, see Install and upgrade the CSI plug-in, Set quotas on the subdirectories of NAS volumes, and FAQ about NAS volumes.

#### Use FlexVolume to mount file systems

For information about how to mount file systems by using FlexVolume, see Use NAS volumes. To ensure flexibility and reduce O&M complexity, we recommend that you use PVs or persistent volume claims (PVCs) instead of volumes when you use FlexVolume to mount file systems. You can mount file systems by using the FlexVolume storage driver provided by Alibaba Cloud. You cannot mount file systems by using the NFS driver provided by Kubernetes.

FlexVolume supports the following methods:

- Mount file systems as static PVs. For more information, see Mount a statically provisioned NAS volume.
- Mount file systems as dynamic PVs. For more information, see Mount a dynamically provisioned NAS volume.

For more information, see FAQ about NAS volumes.

#### Mount file systems on Windows containers

For more information, see Mount SMB file systems to Windows containers.

# 3.2. Use the CSI storage plug-in to mount a NAS instance

### 3.2.1. Overview

Container Service for Kubernetes (ACK) allows you to mount Apsara File Storage NAS (NAS) file systems as persistent volumes (PVs) in ACK clusters. This topic describes the limits of NAS volumes and provides usage notes for NAS volumes.

NAS file systems can be mounted to an ACK cluster by using the Container Storage Interface (CSI) plugin in two forms:

- Mount as statically provisioned PVs
- Mount as dynamically provisioned PVs

#### Prerequisites

A NAS file system is created and a mount target is added to the file system. To create a NAS file system and add a mount target, log on to the NAS console. The mount target of the NAS file system and your cluster are deployed in the same virtual private cloud (VPC).

The mount target is in the following format: 055f84ad83-ixxxx.cn-hangzhou.nas.aliyuncs.com .

#### Usage notes

- Apsara File Storage NAS is a shared storage service. A persistent volume claim (PVC) that is used to mount a NAS file system can be shared among pods.
- Do not delete the mount target before you unmount the NAS file system. Otherwise, the operating system hang may occur.
- After a mount target is created, wait until the mount target is Available for use.
- We recommend that you use NFSv3.
- We also recommend that you upgrade the CSI plug-in to the latest version before you mount NAS file systems as PVs.
- General-purpose and Extreme NAS file systems have different limits on mounting scenarios, the number of file systems, and file sharing protocols. For more information, see Limits of Apsara File

#### Storage NAS.

### 3.2.2. Install and upgrade the CSI plug-in

The CSI plug-in consists of CSI-Plugin and CSI-Provisioner. This topic describes how to install and upgrade the CSI plug-in in a Container Service for Kubernetes (ACK) cluster.

#### Prerequisites

- A cluster of ACK later than 1.14 is created, and the CSI plug-in is specified as the volume plug-in of the cluster. For more information, see Create an ACK managed cluster.
- You are connected to the cluster by using kubectl. For more information, see Connect to ACK clusters by using kubectl.

#### Install CSI-Plugin and CSI-Provisioner

If you do not specify FlexVolume as the volume plug-in when you create a managed Kubernetes cluster or a dedicated Kubernetes cluster, the system installs CSI-Plugin and CSI-Provisioner by default.

#### Verify the installation

Check whet her CSI-Plugin and CSI-Provisioner are installed in the cluster.

• Run the following command to check whether CSI-Plugin is installed in the cluster:

kubectl get pod -n kube-system | grep csi-plugin

• Run the following command to check whether CSI-Provisioner is installed in the cluster:

kubectl get pod -n kube-system | grep csi-provisioner

#### Upgrade CSI-Plugin and CSI-Provisioner

You can upgrade CSI-Plugin and CSI-Provisioner in the ACK console.

- 1. Log on to the ACK console.
- 2. In the left-side navigation pane of the ACK console, click **Clusters**.
- 3. On the **Clusters** page, find the cluster that you want to manage and click the name of the cluster or click **Details** in the **Actions** column. The details page of the cluster appears.
- 4. In the left-side navigation pane of the details page, choose **Operations > Add-ons**.
- 5. Click the **Storage** tab, find **csi-plugin** and **csi-provisioner**, and click **Upgrade**.
- 6. In the Note message, confirm the versions of the plug-ins and click OK. After the plug-ins are upgraded, the system prompts that the upgrades are completed and the current versions of the plug-ins are displayed.

If the plug-ins fail to be upgraded in the console or the plug-ins fail to pass the precheck, you can perform the following operations accordingly:

- CSI-Plugin fails to pass the precheck.
  - If volumes that use disks, Apsara File Storage NAS (NAS) file systems, or Object Storage Service (OSS) buckets are not provisioned in the cluster, you must manually upgrade CSI-Plugin. For more information, see Upgrade CSI-Plugin.
  - If volumes that use disks, NAS file systems, or OSS buckets are provisioned in the cluster, and the cluster is created in a staging environment, you must manually upgrade CSI-Plugin. For more information, see Upgrade CSI-Plugin.
- If volumes that use disks, NAS file systems, or OSS buckets are provisioned in the cluster, and business critical data is stored in the volumes, Submit a ticket to request technical support.
- CSI-Plugin passes the precheck but fails to be upgraded. Check whether the nodes in the cluster are in the Ready state. If CSI-Plugin is installed on a node that is in the NotReady state, you must fix the state of the node. If you cannot identify the cause of the failure, Submit a ticket to request technical support.
- CSI-Plugin is displayed in the console but CSI-Provisioner is not displayed. CSI-Provisioner is deployed by using a StatefulSet. In this case, Submit a ticket to request technical support.
- CSI-Provisioner fails to pass the precheck.
  - If no volumes that use disks or NAS file systems are dynamically provisioned by using StorageClasses in the cluster, you must manually upgrade CSI-Provisioner. For more information, see Upgrade CSI-Provisioner.
  - If volumes that use disks or NAS file systems are dynamically provisioned by using StorageClasses in the cluster, and the cluster is created in a staging environment, you must manually upgrade CSI-Provisioner. For more information, see Upgrade CSI-Provisioner.
  - If volumes that use disks or NAS file systems are dynamically provisioned by using StorageClasses in the cluster, and business critical data is stored in the volumes, Submit a ticket to request technical support.
- CSI-Provisioner passes the precheck but fails to be upgraded. In this case, Submit a ticket to request technical support.

# 3.2.3. Mount a statically provisioned NAS

### volume

Apsara File Storage NAS (NAS) is a distributed file system that supports shared access, elastic scaling, high reliability, and high performance. This topic describes how to mount a statically provisioned NAS volume, and how to enable persistent storage and shared storage by using a statically provisioned NAS volume.

#### Prerequisites

- A Container Service for Kubernetes (ACK) cluster is created. For more information, see Create an ACK managed cluster.
- A NAS file system is created. For more information, see Create a NAS file system. If you want to encrypt data in a NAS volume, configure the encryption settings when you create the NAS file system.
- A mount target is created for the NAS file system. For more information, see Manage mount targets. The mount target and the cluster node to which you want to mount the NAS file system must belong to the same virtual private cloud (VPC).
- A kubectl client is connected to the cluster. For more information, see Connect to ACK clusters by using kubectl.

#### Scenarios:

- Your application requires high disk I/O.
- You need a storage service that offers higher read and write throughput than Object Storage Service (OSS).

• You want to share files across hosts. For example, you want to use a NAS file system as a file server.

#### Precautions

- To mount an Extreme NAS file system, set the path parameter of the NAS volume to a subdirectory of */share*. For example, you can specify the */share/path1* subdirectory when you mount an Extreme NAS file system to a pod.
- If a NAS file system is mounted to multiple pods, the data in the file system is shared by the pods. In this case, the application must be able to synchronize data across these pods when data modifications are made by multiple pods.

**?** Note You cannot grant permissions to access the / directory (root directory) of the NAS file system. The user account and user group to which the directory belongs cannot be modified.

• If the securityContext.fsgroup parameter is set in the application template, kubelet performs the ch mod or chown operation after the volume is mounted, which increases the time consumption.

(?) Note For more information about how to speed up the mounting process when the securityContext.fsgroup parameter is set, see Why does it require a long time to mount a NAS volume?.

#### Mount a statically provisioned NAS volume in the console

#### Step 1: Create a PV

- 1. Log on to the ACK console.
- 2. In the left-side navigation pane of the ACK console, click **Clusters**.
- 3. On the **Clusters** page, find the cluster that you want to manage and click the name of the cluster or click **Details** in the **Actions** column. The details page of the cluster appears.
- 4. In the left-side navigation pane of the cluster details page, choose **Volumes > Persistent Volumes**.
- 5. In the upper-right corner of the Persistent Volumes page, click Create.
- 6. In the Create PV dialog box, set the following parameters.

Parameter	Description
РV Туре	You can select Cloud Disk, NAS, or OSS. In this example, NAS is selected.
Volume Name	The name of the persistent volume (PV) that you want to create. The name must be unique in the cluster. In this example, pv-nas is used.
Volume Plug-in	You can select FlexVolume or CSI. In this example, CSI is selected.
Capacity	The capacity of the PV. A NAS file system provides unlimited capacity. This parameter does not limit the storage usage of the NAS file system but defines the capacity of the PV.
Access Mode	You can select ReadWriteMany or ReadWriteOnce. Default value: ReadWriteMany.

Parameter	Description
Mount Target Domain Name	You can Select Mount Target or enter a Custom mount target.
Show Advanced Options	• <b>Subdirectory</b> : the subdirectory of the NAS file system that you want to mount. The subdirectory must start with a forward slash (/). After you set this parameter, the PV is mounted to the subdirectory.
	If the specified subdirectory does not exist, the system automatically creates the subdirectory in the NAS file system and mounts the subdirectory to the cluster.
	If you do not set this parameter, the root directory of the NAS file system is mounted.
	If you want to mount an Extreme NAS file system, the subdirectory must be under the <i>/share</i> directory.
	• Version: the version of the PV.
Label	Add labels to the PV.

#### 7. Click Create.

#### Step 2: Create a PVC

- 1. In the left-side navigation pane of the details page, choose Volumes > Persistent Volume Claims.
- 2. In the upper-right corner of the **Persistent Volume Claims** page, click **Create**.
- 3. In the **Create PVC** dialog box, set the following parameters.

Parameter	Description					
РVС Туре	You can select Cloud Disk, NAS, or OSS. In this example, NAS is selected.					
Name	The name of the persistent volume claim (PVC). The name must be unique in the cluster.					
Allocation Mode	In this example, Existing Volumes is selected.					
	<b>Note</b> If no PV is created, you can set <b>Allocation Mode</b> to <b>Create Volume</b> , and set the required parameters to create a PV. For more information, see <b>Create a PV</b> .					
Existing Volumes	Click <b>Select PV</b> . Find the PV that you want to use and click <b>Select</b> in the Actions column.					

Parameter	Description					
Capacity	The capacity claimed by the PVC.					
	<b>Note</b> The capacity claimed by the PVC cannot exceed the capacity of the PV that is bound to the PVC.					

#### 4. Click Create.

After the PVC is created, you can view the PVC in the PVCs list. The PVC is bound to the corresponding PV.

#### Step 3: Create an application

- 1. In the left-side navigation pane of the details page, choose Workloads > Deployments.
- 2. In the upper-right corner of the **Deployments** page, click **Create from Image**.
- 3. Set the application parameters.

This example shows how to set the volume parameters. For more information about other parameters, see Create a stateless application by using a Deployment. You can add local volumes and cloud volumes.

- Add Local Storage: You can select HostPath, ConfigMap, Secret, or EmptyDir from the PV Type drop-down list. Then, set the Mount Source and Container Path parameters to mount the volume to a container path. For more information, see Volumes.
- Add PVC: You can add cloud volumes.

In this example, a NAS volume is specified as the mount source and mounted to the /tmp path in the container.

	Volume: 🞯	Add Local Storage					
		PV Type	Mount Source	Container Path			
Volume		• Add PVC					
No.		PV Type	Mount Source	Container Path			
		Cloud Storage	✓ csi-disk-pvc +	Лттр	•		

4. Set the other parameters and click Create.

After the application is created, you can use the NAS volume to store application data.

#### Mount a statically provisioned NAS volume by using kubectl

1. Run the following command to create a statically provisioned PV:

kubectl create -f pv-nas.yaml

The following YAML template provides an example on how to create a statically provisioned PV:

```
apiVersion: v1
kind: PersistentVolume
metadata:
 name: pv-nas
 labels:
  alicloud-pvname: pv-nas
spec:
 capacity:
   storage: 5Gi
 accessModes:
   - ReadWriteMany
 csi:
   driver: nasplugin.csi.alibabacloud.com
   volumeHandle: pv-nas
   volumeAttributes:
     server: "2564f4****-ysu87.cn-shenzhen.nas.aliyuncs.com"
     path: "/csi"
 mountOptions:
 - nolock,tcp,noresvport
  - vers=3
```

Parameter	Description
name	The name of the PV.
labels	The labels that you want to add to the PV.
storage	The capacity of the NAS volume.
accessModes	The access mode of the PV.
driver	The type of driver. In this example, the parameter is set to nasplugin.csi.alibabacloud.com. This indicates that the NAS Container Storage Interface (CSI) plug-in provided by Alibaba Cloud is used.
volumeHandle	The unique identifier of the PV. If multiple PVs are used, the identifier of each PV must be unique.
server	The mount target of the NAS file system.
path	The subdirectory of the NAS file system that you want to mount. If you want to mount an Extreme NAS file system, the subdirectory must be under the <i>/share</i> directory.
vers	The version of the Network File System (NFS) protocol. We recommend that you use NFSv3. Extreme NAS file systems support only NFSv3.

#### 2. Run the following command to create a PVC used for static provisioning:

When you create a PVC of the NAS type, set the selector parameter to configure how to select a

#### PV and bind it to the PVC.

kubectl create -f pvc-nas.yaml

The following YAML template provides an example on how to create a PVC used for static provisioning:

kind: PersistentVolumeClaim
apiVersion: v1
metadata:
name: pvc-nas
spec:
accessModes:
- ReadWriteMany
resources:
requests:
storage: 5Gi
selector:
matchLabels:
alicloud-pvname: pv-nas

Parameter	Description
name	The name of the PVC.
accessModes	The access mode of the PVC.
storage	The capacity claimed by the PVC. The claimed capacity cannot exceed the capacity of the PV bound to the PVC.
mathLabels	The labels are used to select a PV and bind it to the PVC.

3. Run the following command to create an application named **nas-static** and mount the created PVC to the application:

kubectl create -f nas.yaml

The following YAML template provides an example of the *nas.yaml* file that is used to create the **nas-static** application:

apiVersion: apps/vl	
kind: Deployment	
metadata:	
name: nas-static	
labels:	
app: nginx	
spec:	
replicas: 2	
selector:	
matchLabels:	
app: nginx	
template:	
metadata:	
labels:	
app: nginx	
spec:	
containers:	
- name: nginx	
image: nginx	
ports:	
- containerPort: 80	
volumeMounts:	
- name: pvc-nas	
mountPath: "/data"	
volumes:	
- name: pvc-nas	
persistentVolumeClaim:	
claimName: pvc-nas	

Parameter	Description
mountPath	The path of the container where the NAS volume is mounted.
claimName	The name of the PVC mounted to the application.

#### 4. Run the following command to query the pods that run the application:

kubectl get	pod				
Expected outp	out:				
NAME		READY	STATUS	RESTARTS	AGE
nas-static-5	b5cdb85f6-n****	1/1	Running	0	32s
nas-static-c	:5bb4746c-4****	1/1	Running	0	32s

#### Verify that the NAS file system can be used to persist data

1. Query the pods that run the application and the files in the mounted NAS file system.

i. Run the following command to query the pods that run the application:

kubectl get pod

Expected output:

```
NAMEREADYSTATUSRESTARTSAGEnas-static-5b5cdb85f6-n***1/1Running032snas-static-c5bb4746c-4***1/1Running032s
```

ii. Run the following command to query files in the /data path of a pod. The pod nas-static-5 b5cdb85f6-n\*\*\*\* is used as an example:

```
kubectl exec nas-static-5b5cdb85f6-n**** ls /data
```

No output is returned. This indicates that no file exists in the /data path.

2. Run the following command to create a file named *nas* in the */data* path of the pod nas-static-5b5cdb85f6-n\*\*\*\* :

kubectl exec nas-static-5b5cdb85f6-n\*\*\*\* touch /data/nas

3. Run the following command to query files in the /data path of the pod nas-static-5b5cdb85f6-n

```
kubectl exec nas-static-5b5cdb85f6-n**** ls /data
```

Expected output:

nas

4. Run the following command to delete the pod:

```
kubectl delete pod nas-static-5b5cdb85f6-n****
```

5. Open another command-line interface (CLI) and run the following command to view how the pod is deleted and recreated:

kubectl get pod -w -l app=nginx

- 6. Verify that the file still exists after the pod is deleted.
  - i. Run the following command to query the name of the recreated pod:

kubectl get pod

Expected output:

NAME		READY	STATUS	RESTARTS	AGE
nas-static-5b5cdb85f6-n****	1/1	Runni	ng 0	32s	
nas-static-c5bb4746c-4****	1/1	Runni	ng 0	32s	

ii. Run the following command to query files in the /data path of the pod nas-static-5b5cdb85

kubectl exec nas-static-5b5cdb85f6-n\*\*\*\* ls /data

Expected output:

nas

The *nas* file still exists in the /data path. This indicates that data is persisted to the NAS file system.

#### Verify that data in the NAS file system can be shared across pods

- 1. Query the pods that run the application and the files in the mounted NAS file system.
  - i. Run the following command to query the pods that run the application:

kubectl get pod

Expected output:

NAME	READY	STATUS	RESTARTS	AGE
nas-static-5b5cdb85f6-n****	1/1	Running	0	32s
nas-static-c5bb4746c-4****	1/1	Running	0	32s

ii. Run the following command to query files in the /data path of each pod:

```
kubectl exec nas-static-5b5cdb85f6-n**** ls /data
kubectl exec nas-static-c5bb4746c-4**** ls /data
```

2. Run the following command to create a file named *nas* in the */data* path of a pod:

kubectl exec nas-static-5b5cdb85f6-n\*\*\*\* touch /data/nas

- 3. Query files in the /data path of each pod.
  - i. Run the following command to query files in the /data path of the pod nas-static-5b5cdb85 f6-n\*\*\*\* :

kubectl exec nas-static-5b5cdb85f6-n\*\*\*\* ls /data

Expected output:

nas

ii. Run the following command to query files in the /data path of the pod nas-static-c5bb4746

kubectl exec nas-static-c5bb4746c-4\*\*\*\* ls /data

Expected output:

nas

When you create a file in the /data path of one pod, you can also find the file in the /data path of the other pod. This indicates that data in the NAS file system is shared by the two pods.

# 3.2.4. Mount a dynamically provisioned NAS

# volume

You can use the Container Storage Interface (CSI) driver to mount a dynamically provisioned Apsara File Storage NAS (NAS) volume to a Container Service for Kubernetes (ACK) cluster in subpath and filesystem modes. This topic describes how to mount a dynamically provisioned NAS volume and how to test whether the NAS volume can persist and share data as expected.

#### Prerequisites

- An ACK cluster is created. For more information, see Create an ACK managed cluster.
- A NAS file system is created. For more information, see Create a NAS file system. If you want to encrypt data in a NAS volume, configure the encryption settings when you create the NAS file system.
- A mount target is created for the NAS file system. For more information, see Manage mount targets. The mount target and the cluster node to which you want to mount the NAS file system must belong to the same virtual private cloud (VPC).

#### Scenarios:

- Your application requires high disk I/O.
- You need a storage service that offers higher read and write throughput than Object Storage Service (OSS).
- You want to share files across hosts. For example, you want to use a NAS file system as a file server.

#### Precautions

- To mount an Extreme NAS file system, set the path parameter of the NAS volume to a subdirectory of /share. For example, a value of ocd8b4a576-g\*\*\*\*.cn-hangzhou.nas.aliyuncs.com:/share/subpath indicates that the mounted subdirectory of the NAS file system is /share/subpath .
- If a NAS file system is mounted to multiple pods, the data in the file system is shared by the pods. In this case, the application must be able to synchronize data across these pods when data modifications are made by multiple pods.

**Note** You cannot grant permissions to access the / directory (root directory) of the NAS file system. The user account and user group to which the directory belongs cannot be modified.

• If the securityContext.fsgroup parameter is set in the application template, kubelet performs the ch mod or chown operation after the volume is mounted, which increases the time consumption.

(?) Note For more information about how to speed up the mounting process when the securityContext.fsgroup parameter is set, see Why does it require a long time to mount a NAS volume?.

#### Mount a dynamically provisioned NAS volume in the console

You can mount a dynamically provisioned NAS volume only in subpath mode if you use the console. To mount a dynamically provisioned NAS volume in filesystem mode, you must use the kubectl command-line tool.

#### Step 1: Create a StorageClass

- 1. Log on to the ACK console.
- 2. In the left-side navigation pane of the ACK console, click **Clusters**.
- 3. Log on to the ACK console.
- 4. In the left-side navigation pane of the details page, choose **Volumes > StorageClasses**.
- 5. In the upper-right corner of the **StorageClasses** page, click **Create**.
- 6. In the **Create** dialog box, set the parameters.

The following table describes some of the parameters.

Parameter	Description
Name	The name of the StorageClass. The name must start with a lowercase letter and can contain only lowercase letters, digits, periods (.), and hyphens (-).
РV Туре	You can select <b>Cloud Disk</b> or <b>NAS</b> . In this example, <b>NAS</b> is selected.
Volume Plug-in	By default, CSI is selected.
Reclaim Policy	The reclaim policy. By default, this parameter is set to Delete. You can also set this parameter to Retain.
	<ul> <li>Delete mode: When a persistent volume claim (PVC) is deleted, the related PV and NAS file system are also deleted.</li> </ul>
	• Retain mode: When a PVC is deleted, the related PV and NAS file system are retained and can only be manually deleted.
	If you require higher data security, we recommend that you use the Retain mode to prevent data loss caused by user errors.
Mount Options	The mount options, such as the Network File System (NFS) version.
Mount Target Domain Name	The mount target of the NAS file system. If no mount target is available, you must create a NAS file system first. For more information, see Use CNFS to manage NAS file systems.
Path	The mount path of the NAS file system.

#### 7. Click Create.

You can find the created StorageClass in the StorageClasses list.

#### Step 2: Create a PVC

- 1. In the left-side navigation pane of the details page, choose Volumes > Persistent Volume Claims.
- 2. In the upper-right corner of the **Persistent Volume Claims** page, click **Create**.
- 3. In the Create PVC dialog box, set the following parameters.

Parameter	Description
РVС Туре	You can select Cloud Disk, NAS, or OSS. In this example, NAS is selected.

Parameter	Description
Name	The name of the PVC. The name must be unique in the cluster.
Allocation Mode	In this example, Use StorageClass is selected.
Existing Storage Class	Click <b>Select</b> . In the <b>Select Storage Class</b> dialog box, find the StorageClass that you want to use and click <b>Select</b> in the <b>Actions</b> column.
Capacity	The capacity claimed by the PVC.
Access Mode	Default value: ReadWriteMany. You can also select ReadWriteOnce.

#### 4. Click Create.

After the PVC is created, you can find the PVC in the PVCs list. The PVC is bound to the corresponding PV.

#### Step 3: Create an application

- 1. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
- 2. In the upper-right corner of the **Deployments** page, click **Create from Image**.
- 3. Set the application parameters.

This example shows how to set the volume parameters. For more information about other parameters, see Create a stateless application by using a Deployment. You can add local volumes and cloud volumes.

- Add Local Storage: You can select HostPath, ConfigMap, Secret, or EmptyDir from the PV Type drop-down list. Then, set the Mount Source and Container Path parameters to mount the volume to a container path. For more information, see Volumes.
- Add PVC: You can add cloud volumes.

In this example, a NAS volume is specified as the mount source and mounted to the /tmp path in the container.

Volume: 🚱	<ul> <li>Add Local Storage</li> </ul>			
	PV Type	Mount Source	Container Path	
Volume	Add PVC			
No.	PV Type	Mount Source	Container Path	
	Cloud Storage	✓ csi-disk-pvc + //tm	np	•
	Cloud Storage	✓ csi-disk-pvc + /tm	np	

4. Set the other parameters and click Create.

After the application is created, you can use the NAS volume to store application data.

# Mount a dynamically provisioned NAS volume in subpath mode by using kubectl

The subpath mode is applicable to scenarios where you want to share a NAS volume among different applications or pods. You can also use this mode to mount different subdirectories of the same NAS file system to different pods.

To mount a dynamically provisioned NAS volume in subpath mode, you must manually create a NAS file system and a mount target.

- 1. Create a NAS file system and a mount target.
  - i. Log on to the NAS console.
  - ii. Create a NAS file system. For more information, see Create a NAS file system.
  - iii. Create a mount target. For more information, see Manage mount targets.
- 2. Create a StorageClass.
  - i. Create an *alicloud-nas-subpath.yaml* file and copy the following content into the file:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
   name: alicloud-nas-subpath
mountOptions:
    nolock,tcp,noresvport
    vers=3
parameters:
    volumeAs: subpath
    server: "0cd8b4a576-g****.cn-hangzhou.nas.aliyuncs.com:/k8s/"
provisioner: nasplugin.csi.alibabacloud.com
reclaimPolicy: Retain
```

Parameter	Description
mountOptions	Set the options parameter and specify the NFS version in the mountOptions field.
volumeAs	You can select subpath or filesystem. subpath indicates that a subdirectory is mounted to the cluster. filesystem indicates that a file system is mounted to the cluster.
server	When you mount a subdirectory of the NAS file system as a PV, this parameter specifies the mount target of the NAS file system.
provisioner	The type of driver. In this example, the parameter is set to nasplugin. csi.alibabacloud.com . This indicates that the NAS Container Storage Interface (CSI) plug-in provided by Alibaba Cloud is used.
	The reclaim policy of the PV. By default, this parameter is set to Delete. You can also set this parameter to Retain.
an ala in Dalian	<ul> <li>Delete mode: When a persistent volume claim (PVC) is deleted, the related PV and NAS file system are also deleted.</li> </ul>
reclaimPolicy	<ul> <li>Retain mode: When a PVC is deleted, the related PV and NAS file system are retained and can only be manually deleted.</li> </ul>
	If you require higher data security, we recommend that you use the Retain mode to prevent data loss caused by user errors.

Parameter	Description	
archiveOnDelete	This parameter specifies the reclaim policy of backend storage when reclaimPolicy is set to Delete. NAS is a shared storage service. You must set both reclaimPolicy and archiveOnDelete to ensure data security. Configure the policy in the parameters section. The default value is true This value indicates that the subdirectory or files are not deleted when the PV is deleted. Instead, the subdirectory or files are renamed in the format of archived-{pvName}.{timestamp}}. If the value is set to false, it indicates that the backend storage is deleted when the PV is deleted.	
	<b>Note</b> We recommend that you do not set the value to false when the service receives a large amount of network traffic. For more information, see What do I do if the task queue of alicloud-nas-controller is full and PVs cannot be created when I use a dynamically provisioned NAS volume?.	

#### ii. Run the following command to create a StorageClass:

kubectl create -f alicloud-nas-subpath.yaml

#### 3. Create a PVC.

i. Create a *pvc.yaml* file and copy the following content into the file:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
   name: nas-csi-pvc
spec:
   accessModes:
        ReadWriteMany
   storageClassName: alicloud-nas-subpath
   resources:
        requests:
        storage: 20Gi
```

Parameter	Description
name	The name of the PVC.
accessModes	The access mode of the PVC.
storageClassName	The name of the StorageClass that you want to associate with the PVC.
storage	The storage that is requested by the application.

#### ii. Run the following command to create a PVC:

kubectl create -f pvc.yaml

#### 4. Create applications.

Deploy two applications named **nginx-1** and **nginx-2** to share the same subdirectory of the NAS file system.

i. Create an *nginx-1.yml* file and copy the following content into the file:

```
apiVersion: apps/v1
kind: Deployment
metadata:
 name: deployment-nas-1
 labels:
   app: nginx-1
spec:
  selector:
   matchLabels:
     app: nginx-1
  template:
   metadata:
     labels:
       app: nginx-1
   spec:
     containers:
      - name: nginx
       image: nginx:1.7.9
       ports:
       - containerPort: 80
       volumeMounts:
         - name: nas-pvc
          mountPath: "/data"
     volumes:
        - name: nas-pvc
         persistentVolumeClaim:
           claimName: nas-csi-pvc
```

- mountPath : the path where the NAS file system is mounted in the container.
- claimName : the name of the PVC that you want to mount to the application. In this example, the value is set to nas-csi-pvc.

ii. Create an *nginx-2.yml* file and copy the following content into the file:

```
apiVersion: apps/v1
kind: Deployment
metadata:
 name: deployment-nas-2
 labels:
   app: nginx-2
spec:
  selector:
   matchLabels:
     app: nginx-2
  template:
   metadata:
     labels:
       app: nginx-2
   spec:
     containers:
     - name: nginx
       image: nginx:1.7.9
       ports:
        - containerPort: 80
       volumeMounts:
         - name: nas-pvc
           mountPath: "/data"
      volumes:
        - name: nas-pvc
         persistentVolumeClaim:
           claimName: nas-csi-pvc
```

- mountPath : the path where the NAS file system is mounted in the container. In this example, the value is set to /data.
- claimName : Enter the name of the PVC that is mounted to **nginx-1**. In this example, the value is set to **nas-csi-pvc**.
- iii. Run the following command to deploy applications nginx-1 and nginx-2:

```
kubectl create -f nginx-1.yaml -f nginx-2.yaml
```

5. Run the following command to query the pods that run the applications:

kubectl get pod

Expected output:

NAME	READY	STATUS	RESTARTS	AGE
deployment-nas-1-5b5cdb85f6-n****	1/1	Running	0	32s
deployment-nas-2-c5bb4746c-4****	1/1	Running	0	32s

Note The subdirectory 0cd8b4a576-g\*\*\*\*.cn-hangzhou.nas.aliyuncs.com:/share/nas-7 9438493-f3e0-11e9-bbe5-00163e09\*\*\*\* of the NAS volume is mounted to the /data directory of pods deployment-nas-1-5b5cdb85f6-n\*\*\*\* and deployment-nas-2-c5bb4746c-4\*\*\*\*.
The following information is displayed.

- /share : the subdirectory is mounted in subpath mode as specified in the StorageClass configurations.
- nas-79438493-f3e0-11e9-bbe5-00163e09\*\*\*\* : the name of the PV.

To mount different subdirectories of a NAS file system to different pods, you must create a separate PVC for each pod. To do this, you can create pvc-1 for nginx-1 and pvc-2 for nginx-2.

# Mount a dynamically provisioned NAS volume in filesystem mode by using kubectl

**Notice** By default, if you delete a PV that is mounted in filesystem mode, the system retains the related NAS file system and mount target. To delete the NAS file system and mount target together with the PV, set reclaimPolicy to Delete and set deleteVolume to true in the StorageClass configurations.

The filesystem mode is applicable to scenarios where you want to dynamically create and delete NAS file systems and mount targets.

When you mount a NAS volume in filesystem mode, you can create only one NAS file system and one mount target for each pod. You cannot share a NAS volume among multiple pods. The following procedure shows how to mount a dynamically provisioned NAS volume in filesystem mode.

1. Configure a Resource Access Management (RAM) permission policy and attach it to a RAM role.

The filesystem mode allows you to dynamically create and delete NAS file systems and mount targets. To enable this feature, you must grant the required permissions to csi-nasprovisioner. The following code block shows a permission policy that contains the required permissions:

```
{
    "Action": [
        "nas:DescribeMountTargets",
        "nas:CreateMountTarget",
        "nas:DeleteFileSystem",
        "nas:CreateFileSystem"
    ],
    "Resource": [
        "*"
    ],
        "Effect": "Allow"
}
```

You can grant the permissions by using the following methods:

• Attach the preceding permission policy to the master RAM role of your ACK cluster. For more information, see ACK default roles.

Overview Basic Information Connection Informat	on Cluster Resources Cluster Logs
The following resources are managed by the current Kubern applications that are deployed in the cluster	tes cluster. We recommend that you do not modify or delete these resources. Otherwise, errors may occur in the cluster and affect the running of the
Resource Orchestration Service (ROS)	k8s-for-cs-
VPC	vpc-
NAT Gateway	ngw-
SLB EIP for Access to API Server	eip-
Master Instance	
Master RAM Role	KubernetesMasterRole-
Security Group	sg-
Worker RAM Role	KubernetesWorkerRole-
Scaling Group	asg-
APIServerSLB	lb-

Note The master RAM role is automatically assigned to a managed Kubernetes cluster. However, for a dedicated Kubernetes cluster, you must manually assign the master RAM role.

Create a RAM user and attach the preceding permission policy to the RAM user. Then, generate an AccessKey pair and specify the AccessKey pair in the env variable in the configurations of the csi-provisioner StatefulSet. For more information, see ACK default roles.

```
env:
    - name: CSI_ENDPOINT
        value: unix://socketDir/csi.sock
    - name: ACCESS_KEY_ID
        value: ""
    - name: ACCESS_KEY_SECRET
        value: ""
```

#### 2. Create a StorageClass.

i. Create an *alicloud-nas-fs.yaml* file and copy the following content into the file:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
 name: alicloud-nas-fs
mountOptions:
- nolock, tcp, noresvport
- vers=3
parameters:
 volumeAs: filesystem
 storageType: Performance
 zoneId: cn-hangzhou-a
 vpcId: "vpc-2ze9c51qb5kp1nfqu****"
 vSwitchId: "vsw-gw8tk6gecif0eu9ky****"
 accessGroupName: DEFAULT VPC GROUP NAME
 deleteVolume: "false"
provisioner: nasplugin.csi.alibabacloud.com
reclaimPolicy: Retain
```

Parameter	Description
volumeAs	<ul> <li>The mode in which the NAS file system is mounted. Supported modes are:</li> <li>filesystem: csi-nasprovisioner automatically creates a NAS file system. Each PV corresponds to a NAS file system.</li> <li>subpath: csi-nasprovisioner automatically creates a subdirectory in a NAS file system. Each PV corresponds to a subdirectory of the NAS file system.</li> </ul>
storageType	The type of NAS file system. You can select <b>Performance</b> or <b>Capacity</b> . Default value: Performance.
zoneld	The ID of the zone to which the NAS file system belongs.
vpcld	The ID of the VPC to which the mount target of the NAS file system belongs.
vSwitchld	The ID of the vSwitch to which the mount target of the NAS file system belongs.
accessGroupName	The permission group to which the mount target of the NAS file system belongs. Default value: DEFAULT_VPC_GROUP_NAME.
deleteVolume	The reclaim policy of the NAS file system when the related PV is deleted. NAS is a shared storage service. Therefore, you must specify both deleteVolume and reclaimPolicy to ensure data security.
provisioner	The type of driver. In this example, the parameter is set to nas plugin.csi.alibabacloud.com . This indicates that the NAS CSI plug-in provided by Alibaba Cloud is used.
reclaimPolicy	The reclaim policy of the PV. When you delete a PVC, the related NAS file system is automatically deleted only if you set deleteVolume to true and reclaimPolicy to Delete.

#### ii. Run the following command to create a StorageClass:

kubectl create -f alicloud-nas-fs.yaml

3. Create a PVC and pods to mount a NAS volume.

i. Create a *nas.yaml* file and copy the following content into the file:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
   name: nas-csi-pvc-fs
spec:
   accessModes:
        - ReadWriteMany
   storageClassName: alicloud-nas-fs
   resources:
        requests:
        storage: 20Gi
```

ii. Create an *nginx.yaml* file and copy the following content into the file:

```
apiVersion: apps/v1
kind: Deployment
metadata:
 name: deployment-nas-fs
 labels:
   app: nginx
spec:
  selector:
   matchLabels:
     app: nginx
  template:
   metadata:
     labels:
       app: nginx
   spec:
     containers:
      - name: nginx
       image: nginx:1.7.9
       ports:
        - containerPort: 80
       volumeMounts:
         - name: nas-pvc
           mountPath: "/data"
      volumes:
        - name: nas-pvc
         persistentVolumeClaim:
           claimName: nas-csi-pvc-fs
```

iii. Run the following command to create the PVC and pods:

kubectl create -f pvc.yaml -f nginx.yaml

In filesystem mode, the CSI driver automatically creates a NAS file system and a mount target when you create the PVC. When the PVC is deleted, the file system and the mount target are retained or deleted based on the settings of the deleteVolume and reclaimPolicy parameters.

#### Verify that the NAS file system can be used to persist data

NAS provides persistent storage. When a pod is deleted, the recreated pod automatically synchronizes the data of the deleted pod.

You can use the following example to verify persistent storage.

- 1. Query the pods that run the application and the files in the mounted NAS file system.
  - i. Run the following command to query the pods that run the application:

kubectl get pod

Expected output:

NAME	READY	STATUS	RESTARTS	AGE
deployment-nas-1-5b5cdb85f6-n****	1/1	Running	0	32s
deployment-nas-2-c5bb4746c-4****	1/1	Running	0	32s

ii. Run the following command to query files in the /data path of a pod. The pod named deplo yment-nas-1-5b5cdb85f6-n\*\*\*\* is used as an example:

kubectl exec deployment-nas-1-5b5cdb85f6-n\*\*\*\* -- ls /data

No output is returned. This indicates that no file exists in the /data path.

2. Run the following command to create a file named *nas* in the */data* path of the pod deploymentnas-1-5b5cdb85f6-n\*\*\*\* :

kubectl exec deployment-nas-1-5b5cdb85f6-n\*\*\*\* -- touch /data/nas

3. Run the following command to query files in the /data path of the pod deployment-nas-1-5b5cdb 85f6-n\*\*\*\* :

kubectl exec deployment-nas-1-5b5cdb85f6-n\*\*\*\* -- ls /data

Expected output:

nas

4. Run the following command to delete a pod:

kubectl delete pod deployment-nas-1-5b5cdb85f6-n\*\*\*\*

5. Open another command-line interface (CLI) and run the following command to view how the pod is deleted and recreated:

kubectl get pod -w -l app=nginx

- 6. Verify that the file still exists after the pod is deleted.
  - i. Run the following command to query the name of the recreated pod:

kubectl get pod

Expected output:

NAME	READY	STATUS	RESTARTS	AGE
deployment-nas-1-5b5cdm2g5-m****	1/1	Running	0	32s
deployment-nas-2-c5bb4746c-4****	1/1	Running	0	32s

ii. Run the following command to query files in the /data path of the pod deployment-nas-1-5b 5cdm2q5-m\*\*\*\* :

```
kubectl exec deployment-nas-1-5b5cdm2g5-m**** -- ls /data
```

Expected output:

nas

The *nas* file still exists in the /data path. This indicates that data is persisted to the NAS file system.

#### Verify that data in the NAS file system can be shared across pods

You can mount a NAS volume to multiple pods. When the data is modified in one pod, the modifications are automatically synchronized to other pods.

You can use the following example to verify shared storage.

- 1. Query the pods that run the application and the files in the mounted NAS file system.
  - i. Run the following command to query the pods that run the application:

kubectl get pod

Expected output:

NAME	READY	STATUS	RESTARTS	AGE
deployment-nas-1-5b5cdb85f6-n****	1/1	Running	0	32s
deployment-nas-2-c5bb4746c-4****	1/1	Running	0	32s

ii. Run the following command to query files in the /data path of each pod:

kubectl exec deployment-nas-1-5b5cdb85f6-n\*\*\*\* -- ls /data
kubectl exec deployment-nas-2-c5bb4746c-4\*\*\*\* -- ls /data

2. Run the following command to create a file named *nas* in the */data* path of a pod:

kubectl exec deployment-nas-1-5b5cdb85f6-n\*\*\*\* -- touch /data/nas

- 3. Query files in the */data* path of each pod.
  - i. Run the following command to query files in the /data path of the pod deployment-nas-1-5b 5cdb85f6-n\*\*\*\* :

kubectl exec deployment-nas-1-5b5cdb85f6-n\*\*\*\* -- ls /data

Expected output:

nas

ii. Run the following command to query files in the /*data* path of the pod deployment-nas-2-c5 bb4746c-4\*\*\*\* :

kubectl exec deployment-nas-2-c5bb4746c-4\*\*\*\* -- ls /data

Expected output:

nas

When you create a file in the /data path of one pod, you can also find the file in the /data path of the other pod. This indicates that data in the NAS file system is shared by the two pods.

# 3.2.5. Set quotas on the subdirectories of NAS

## volumes

You can set quotas to manage resource allocation and improve the overall resource utilization. Container Service for Kubernetes (ACK) allows you to use the CSI plug-in to set quotas on the subdirectories of Apsara File Storage NAS volumes. This topic describes how to set quotas on the subdirectories of NAS volumes.

#### Prerequisites

- The image version of csi-plugin is V1.18.8.45 or later. For more information about csi-plugin versions, see csi-plugin.
- The NAS volume is mounted by using a subdirectory.

#### Limits

- Only NAS Capacity file systems support quota limits. For more information about the types of NAS file systems, see General-purpose NAS file systems.
- Quotas can be set only for volumes that are mounted by using subdirectories.
- Quot a limits are supported in all regions except the China (Hohhot) and China (Ulanqab) regions.
- For each file system, you can configure quot as only on a maximum of 10 directories.
  - You can set an enforcement quota on a directory. If the quota is exceeded, you cannot write data to the directory. The write operations include the operations that are used to increase the length of files, create files, subdirectories, and special files, and move files to another directory. An IOError error occurs at the frontend.
  - To avoid unexpected errors, use caution when you set enforcement quotas on critical directories.
  - A specific period of time is required before an enforcement quota is enabled or disabled due to asynchronous execution at the backend. In most cases, the time period ranges from 5 to 15 minutes.

#### Examples

1. Create a StorageClass that uses a subdirectory of a NAS file system to provision volumes.

In this example, the following template is used:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: alicloud-nas-sp8
mountOptions:
 - nolock,tcp,noresvport
 - vers=3
parameters:
 volumeAs: subpath
 server: "xxx.cn-hangzhou.nas.aliyuncs.com:/"
archiveOnDelete: "false"
path: "/abc"
 volumeCapacity: "true"
provisioner: nasplugin.csi.alibabacloud.com
reclaimPolicy: Delete
allowVolumeExpansion: "true"
```

Parameter	Description
mountOptions	Set the options parameter and Network File System (NFS) version in the mountOptions field.
volumeAs	You can select subpath or filesystem. subpath specifies that a subdirectory is mounted on the cluster while filesystem specifies that a file system is mounted on the cluster.
server	When you mount a subdirectory of the NAS file system as a persistent volume (PV), this parameter specifies the mount target of the NAS file system.
archiveOnDelete	This parameter specifies whether to delete the backend storage when reclaimPolicy is set to Delete. NAS is a shared storage service. You must set both reclaimPolicy and deleteVolume to ensure data security. Default value: true.
path	The subdirectory of the NAS file system that is mounted. If you mount an Extreme NAS file system, the path must start with <i>/share</i> .
volumeCapacity	This parameter specifies whether to set a quota. Valid values: true and false.
provisioner	The provisioner of the PV that is provided by ACK.
reclaimPolicy	<ul> <li>The policy that is used to reclaim the PV. Valid values:</li> <li>Retain: retains the backend storage when the PV and PVC are deleted. The backend storage may be cloud disks.</li> <li>Delete: automatically deletes the backend storage and PV when the PVC is deleted.</li> </ul>
allowVolumeExpansion	This parameter specifies whether NAS storage volume expansion is supported.

**Note** To create a StorageClass that sets quotas on the subdirectory of a NAS file system, the volumeCapacity parameter must be set to true.

2. Create a PVC that claims a storage capacity of 20 GiB.

In this example, the following template is used:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
   name: pvc-nas-dynamic-create-subpath8
spec:
   accessModes:
        - ReadWriteMany
   storageClassName: alicloud-nas-sp8
   resources:
        requests:
        storage: 20Gi
```

3. Create a Deployment that uses the PVC that is created in Step 2.

In this example, the following template is used:

```
apiVersion: apps/v1
kind: Deployment
metadata:
 name: deployment-nas-dynamic-create8
 labels:
   app: nginx
spec:
 selector:
   matchLabels:
     app: nginx
  template:
   metadata:
     labels:
       app: nginx
    spec:
     containers:
       - name: nginx
         image: nginx:1.14.2
         ports:
            - containerPort: 80
         volumeMounts:
            - name: pvc-nas-dynamic-create-subpath8
             mountPath: "/data"
      volumes:
        - name: pvc-nas-dynamic-create-subpath8
         persistentVolumeClaim:
            claimName: pvc-nas-dynamic-create-subpath8
```

#### Verification

1. Run the following command to write 10 GiB of data to the */data* directory that is mounted on the Deployment that is created in Step 3:

```
dd if=/dev/zero of=10G.txt bs=1M count=10000
```

- 2. Wait 5 to 15 minutes and then check the quota details of the subdirectory.
  - i. Log on to the NAS console.
  - ii. In the left-side navigation pane, choose File System > File System List.
  - iii. Select the NAS file system that you have used and choose **More > Quota Management** in the **Operations** column.
  - iv. On the Quota Management page, click Manage Quotas in the Operations column.

The following figure shows that the subdirectory has a quota limit of 20 GiB. The used storage is 9 GiB.

NAS File System	NAS File System / File System	a / 1c9a648c10	/nas-	CTU25e 0	1713 - AND AN A	PET whether	10/11/06				×
Overview	← lefuiddel	0									
File System	Basic Information	Quota Management	Add U	lser Quotas							c
File System List	Mounting Use	Quota Management can monitor and		Type of	ID	Quota Type	Capacity	Current	File Limit	Current Number of	Operations
Access Group	Access Control	New Directory Quota Directory Pati		User	10	Quota type	Limit (GiB)	Capacity(	The chine	Files	operations
Resource Plans	Quota Management	Directory Path		All Users		Restricted	20	9	Unlimited	1	Edit   Delete
Data Service	Performance Monito	/nas-lig the company and the									
Snapshot											
File Migration 🖸											
File Backup											
Life Cycle Management											
Monitoring Audit											
Performance Monitoring											

When the 20 GiB of storage is used up, the Disk quota exceeded error appears if you try to write more data to the subdirectory.



## 3.2.6. FAQ about NAS volumes

This topic provides answers to some frequently asked questions about Apsara File Storage NAS (NAS) volumes.

- The system prompts chown: Operation not permitted
- What do I do if the task queue of alicloud-nas-controller is full and PVs cannot be created when I use a dynamically provisioned NAS volume?
- Why does it require a long time to mount a NAS volume?

#### The system prompts chown: Operation not permitted

#### Symptom:

The system prompts chown: Operation not permitted when I mount a NAS file system.

Cause:

Your container does not have permissions to use the specified NAS file system.

#### Solution:

Launch the container with root privileges.

# What do I do if the task queue of alicloud-nas-controller is full and PVs cannot be created when I use a dynamically provisioned NAS volume?

#### Symptom:

When you use a dynamically provisioned NAS volume, if the speed of subdirectory creation is faster than the speed of subdirectory deletion, the task queue of alicloud-nas-controller may be full and therefore PVs cannot be created.

#### Cause:

The reclaimPolicy parameter is set to Delete and the archiveOnDelete parameter is set to false in the configuration of the StorageClass that mounts the dynamically provisioned NAS volume.

#### Solution:

Set archiveOnDelete to true. This way, when a PV is deleted, only the name of the mounted subdirectory in the NAS file system is modified. The files in the subdirectory are not deleted.

You must delete these files yourself. For example, you can configure a node to automatically delete files in the root directory by schedule, or start multiple pods to concurrently delete files of specific formats in subdirectories.

#### Why does it require a long time to mount a NAS volume?

#### Symptom:

It requires a long time to mount a NAS volume.

Cause:

If the following conditions are met, the chmod or chown operation is performed when volumes are mounted, which increases the time consumption.

- The AccessModes parameter is set to ReadWriteOnce in the PV and PVC templates.
- The securityContext.fsgroup parameter is set in the application template.

#### Solution:

- If the securityContext.fsgroup parameter is set in the application template, delete the fsgroup parameter in the securityContext section.
- If you want to configure the user ID (UID) and mode of the files in the mounted directory, you can manually mount the directory to an Elastic Compute Service (ECS) instance. You can then perform ch own and chmod operations through a CLI and provision the NAS volume through the CSI plug-in.
   For more information about how to provision NAS volumes through the CSI plug-in, see Mount a statically provisioned NAS volume or Mount a dynamically provisioned NAS volume.
- For clusters of Kubernetes 1.20 or later, you can set the fsGroupChangePolicy parameter to OnRootMismatch. This way, the chmod or chown operation is performed only when the pod that uses the volume is first started. For more information, see Set the security context for a container.

# 3.3. Use the Flexvolume storage plug-in to mount NAS

# 3.3.1. Use NAS volumes

This topic describes how to mount Network Attached Storage (NAS) file systems to clusters of Container Service for Kubernetes (ACK) as volumes and how to use NAS volumes.

#### Prerequisites

A NAS file system is created and a mount target is added to the file system. To create a NAS file system and add a mount target, log on to the NAS console. The mount target of the NAS file system and your cluster are deployed in the same virtual private cloud (VPC).

The mount target is in the following format: 055f84ad83-ixxxx.cn-hangzhou.nas.aliyuncs.com .

#### **Background information**

You can mount file systems of Apsara File Storage NAS to ACK clusters in the following ways:

- Mount NAS file systems as static volumes
  - Directly mount NAS file systems as volumes.
  - Use a pair of persistent volume (PV) and persistent volume claim (PVC) to mount NAS file systems.
- Mount NAS file systems as dynamic volumes

#### Scenarios

• Static volumes

NAS provides shared storage services. You can mount NAS file systems as static volumes to meet the requirements of diverse scenarios.

• Dynamic volumes

You can mount NAS file systems as dynamic volumes when you need to use multiple NAS subdirectories for different applications.

You can also mount NAS file systems as dynamic volumes when you use the StatefulSet controller to deploy applications and want each pod to use a separate NAS volume.

#### How to mount NAS file systems

We recommend that you read the following information before you mount NAS file systems to ACK clusters:

• Recommended volume plug-in

We recommend that you use the Flexvolume driver to mount NAS file systems. The Flexvolume driver is installed by default when you create an ACK cluster in the console. You must make sure that the Flexvolume driver is upgraded to the latest version. For more information, see Upgrade the Flexvolume driver.

- Recommended mounting method
   We recommend that you mount a NAS file system by using a pair of PV and PVC. This makes the NAS file system easier to manage and maintain.
  - For more information about static volumes, see Mount a statically provisioned NAS volume.
  - For more information about dynamic volumes, see Mount a dynamically provisioned NAS volume.
- Not recommended mounting method

We recommend that you do not directly mount NAS file systems as volumes. You can use only the Flexvolume driver to mount volumes to ACK clusters. The Network File System (NFS) driver provided by Kubernetes is not supported.

#### Considerations

- NAS is a shared storage system that provides storage services for multiple pods at a time. A PVC can be shared among multiple pods.
- Do not delete a mount target if the related NAS file system is still mounted. Otherwise, the operating system may become unresponsive.
- After a mount target is created, wait until the mount target is Available for use.
- We recommend that you use NFS v3.
- We recommend that you upgrade Flexvolume to the latest version before you use NAS volumes.
- NAS file systems of Extreme type support only NFS v3. You must specify the nolock parameter when you mount these file systems.

# 3.3.2. Install and upgrade FlexVolume

If you specify FlexVolume as the volume plug-in for a Container Service for Kubernetes (ACK) cluster that runs Kubernetes earlier than 1.16, the system automatically installs FlexVolume and Disk Controller in the cluster. However, the system does not automatically install alicloud-nas-controller. This topic describes how to install and upgrade FlexVolume, and how to install alicloud-nas-controller.

#### Prerequisites

- An ACK cluster is created. For more information, see Create an ACK managed cluster.
- FlexVolume is specified as the volume plug-in of the ACK cluster.
- A kubectl client is connected to the cluster. For more information, see Step 2: Select a type of cluster credentials.

#### Precautions

If alicloud-nas-controller is deployed in the cluster, you must upgrade the image version of alicloudnas-controller to v1.14.8.17-7b898e5-aliyun or later before you can upgrade the Kubernetes version of the cluster to 1.20.

(?) **Note** If you use an open source version, such as nfs-provisioner, to replace alicloud-nascontroller provided by Alibaba Cloud, you may need to find a solution in the open source community to avoid self Link issues.

#### Limits

Only the CentOS 7 and Alibaba Cloud Linux 2 operating systems are supported.

#### Install the components

#### Install FlexVolume

- Clusters that run Kubernetes 1.16 and later do not support FlexVolume. You must install CSI-Plugin in these clusters. For more information, see Differences between the CSI and FlexVolume plug-ins.
- If you specify FlexVolume as the volume plug-in for an ACK cluster that runs Kubernetes earlier than 1.16, the system automatically installs FlexVolume in the cluster. For more information, see Component configurations.

#### Install Disk Controller

• Clusters of ACK 1.16 and later do not support Disk Controller. You must install CSI-Provisioner in these clusters. For more information, see Differences between the CSI and FlexVolume plug-ins.

• If you specify FlexVolume as the volume plug-in for an ACK cluster that runs Kubernetes earlier than 1.16, the system automatically installs Disk Controller in the cluster. For more information, see Component configurations.

#### Install alicloud-nas-controller

If FlexVolume is installed in your cluster, you can manually install alicloud-nas-controller, and then dynamically provision volumes that use Apsara File Storage NAS (NAS) file systems.

You can use the following YAML template to manually install alicloud-nas-controller:

File system mounting Mount a file s ystem on a pod in Kubernetes

```
kind: Deployment
apiVersion: apps/v1
metadata:
 name: alicloud-nas-controller
 namespace: kube-system
spec:
  selector:
   matchLabels:
     app: alicloud-nas-controller
  strategy:
   type: Recreate
  template:
    metadata:
     labels:
       app: alicloud-nas-controller
    spec:
     tolerations:
      - operator: Exists
      affinity:
        nodeAffinity:
          preferredDuringSchedulingIgnoredDuringExecution:
          - weight: 1
           preference:
              matchExpressions:
              - key: node-role.kubernetes.io/master
               operator: Exists
      priorityClassName: system-node-critical
      serviceAccount: admin
      hostNetwork: true
      containers:
        - name: nfs-provisioner
          image: registry.cn-hangzhou.aliyuncs.com/acs/alicloud-nas-controller:v1.14.8.17-7
b898e5-aliyun
         env:
          - name: PROVISIONER NAME
           value: alicloud/nas
          securityContext:
           privileged: true
          volumeMounts:
          - mountPath: /var/log
            name: log
      volumes:
      - hostPath:
          path: /var/log
        name: log
```

#### Verify the installation

Check whet her FlexVolume, Disk Controller, and alicloud-nas-controller are installed in the cluster.

• Run the following command to check whether FlexVolume is installed in the cluster:

```
kubectl get pod -nkube-system | grep flexvolume
```

• Run the following command to check whether Disk Controller is installed in the cluster:

kubectl get pod -nkube-system | grep alicloud-disk-controller

• Run the following command to check whether alicloud-nas-controller is installed in the cluster:

kubectl get pod -nkube-system | grep alicloud-nas-controller

#### Upgrade the components

You can upgrade FlexVolume and Disk Controller in the ACK console. You cannot upgrade alicloud-nascontroller in the ACK console.

If the Kubernetes version of your ACK cluster is upgraded to 1.16 or later, the cluster still supports FlexVolume. You can upgrade FlexVolume in the ACK console.

- 1. Log on to the ACK console.
- 2. In the left-side navigation pane of the ACK console, click Clusters.
- 3. On the **Clusters** page, find the cluster that you want to manage and click the name of the cluster or click **Details** in the **Actions** column. The details page of the cluster appears.
- 4. In the left-side navigation pane of the details page, choose **Operations > Add-ons**.
- 5. Click the Storage tab, find flexvolume and alicloud-disk-controller, and then click Upgrade.
- In the Note message, confirm the versions of the plug-ins and click OK. After the plug-ins are upgraded, the system prompts that the upgrades are completed and the current versions of the plug-ins are displayed.
- When you upgrade FlexVolume in the following scenarios, Submit a ticket to request technical support.
  - The system fails to update FlexVolume in the ACK console.
  - The version of FlexVolume is 1.12 or earlier, and volumes that use disks and Object Storage Service (OSS) buckets are provisioned in the cluster.
  - You want to ensure a successful upgrade because sensitive business data is stored in the cluster and a large number of volumes are used.
- The system fails to upgrade Disk Controller. In this case, Submit a ticket to request technical support.

# 3.3.3. Mount a statically provisioned NAS

### volume

You can use the FlexVolume plug-in provided by Alibaba Cloud to mount Apsara File Storage NAS (NAS) file systems to Container Service for Kubernetes (ACK) clusters. This topic describes how to mount a statically provisioned NAS volume.

#### Prerequisites

- FlexVolume is upgraded to the latest version. For more information, see Install and upgrade FlexVolume.
- A kubectl client is connected to the cluster. For more information, see Connect to ACK clusters by using kubectl.

#### Context

After FlexVolume installed in the cluster, you can mount NAS file systems by using persistent volumes

(PVs) and persistent volume claims (PVCs).

#### Precautions

If the securityContext.fsgroup parameter is set in the application template, kubelet performs the chmod or chown operation after the volume is mounted, which increases the time consumption.

**Note** For more information about how to speed up the mounting process when the securityContext.fsgroup parameter is set, see Why does it require a long time to mount a NAS volume?.

#### Procedure

You can mount a NAS file system by using a PV and a PVC.

1. Create a PV.

You can create a PV in the ACK console or by using a YAML file.

Create a PV by using a YAML file.
 Use the following *nas-pv.yaml* file to create a PV:

```
apiVersion: v1
kind: PersistentVolume
metadata:
 name: pv-nas
spec:
 capacity:
   storage: 5Gi
  storageClassName: nas
 accessModes:
    - ReadWriteMany
  flexVolume:
   driver: "alicloud/nas"
   options:
     server: "0cd8b4a576-u****.cn-hangzhou.nas.aliyuncs.com"
     path: "/k8s"
     vers: "3"
     options: "nolock, tcp, noresvport"
```

- Create a PV in the ACK console.
  - a. Log on to the ACK console.
  - b. In the left-side navigation pane of the ACK console, click Clusters.
  - c. On the **Clusters** page, find the cluster that you want to manage. Then, click the name of the cluster or click **Details** in the **Actions** column.
  - d. In the left-side navigation pane of the cluster details page, choose **Volumes > Persistent Volumes**.
  - e. On the Volumes page, click Create in the upper-right corner of the page.
  - f. In the Create PV dialog box, set the parameters.

Parameter	Description
РV Туре	In this example, NAS is selected.

Parameter	Description
Volume Name	The name of the PV that you want to create. The name must be unique in the cluster. In this example, <b>pv-nas</b> is used.
Volume Plug-in	In this example, Flexvolume is selected. For more information about volume plug-ins, see Differences between the CSI and FlexVolume plug-ins.
Capacity	The capacity of the PV that you want to create. The capacity of the PV cannot exceed that of the NAS file system to be mounted.
Access Mode	Default value: ReadWriteMany.
Mount Target Domain Name	The domain name of the mount target that is used to mount the NAS file system to the cluster. For more information about how to manage the mount targets of a NAS file system, see Manage mount targets.
	Enter a subdirectory in the NAS file system. The subdirectory must start with a forward slash (/). If this parameter is set, the PV will be mounted to the specified subdirectory.
Subdirectory	<ul> <li>If the specified subdirectory does not exist, the system automatically creates the subdirectory in the NAS file system and mounts the subdirectory to the cluster.</li> </ul>
	<ul> <li>If you do not set this parameter, the root directory of the NAS file system is mounted.</li> </ul>
	<ul> <li>If you specify a subdirectory of an Extreme NAS file system, the subdirectory must start with <i>/share</i>.</li> </ul>

Parameter	Description	
	The access permissions on the mounted directory. For example, you can set this parameter to 755, 644, or 777.	
	<ul> <li>Note</li> <li>You can set access permissions only on subdirectories.</li> <li>If the mounted directory stores a large number of files, we recommend that you do not set this parameter. Otherwise, the process of running the chmod command may require an excessive amount of time.</li> </ul>	
Permissions	<ul> <li>If the mounted directory is a subdirectory, this parameter is optional.</li> <li>If you do not set this parameter, the original permissions on the mounted directory are used.</li> <li>Take note of the following items when you set the permissions:</li> <li>For FlexVolume versions earlier than V1.14.6.15-8d3b7e7-aliyun, use the recursive mode when you configure permission settings. The permissions on all files and directories under the mounted directory will be modified.</li> <li>For FlexVolume V1.14.6.15-8d3b7e7-aliyun and later, set the chmod (Change Mode) parameter to configure permission settings.</li> </ul>	

Parameter	Description		
	<ul><li>The change mode of access permissions. Valid values: Non-recursive and Recursive.</li><li>Non-recursive: The permission changes apply only to the</li></ul>		
	mounted directory. The subdirectories and files in the mounted directory are not affected.		
	<ul> <li>Recursive mode: The permission changes apply to the mounted directory, and the subdirectories and files in the mounted directory.</li> </ul>		
chmod (Change Mode)	<b>Note</b> If you select the recursive mode for a mounted directory that contains a large number of files, the process of running the chmod command may require an excessive amount of time. The mount or unmount operation may fail. Exercise caution when you set this parameter.		
Version	The version of the NFS protocol. We recommend that you use NFSv3. Extreme NAS file systems support only NFSv3.		
Labels	Add labels to the PV.		

- g. Click Create.
- 2. Create a PVC.

Use the following *nas-pvc.yaml* file to create a PVC:

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
   name: pvc-nas
spec:
   accessModes:
    - ReadWriteMany
   storageClassName: nas
   resources:
    requests:
      storage: 5Gi
```

3. Create a pod.

Use the following *nas-pod.yaml* file to create a pod:
```
apiVersion: apps/v1
kind: Deployment
metadata:
 name: nas-static
 labels:
   app: nginx
spec:
 replicas: 1
  selector:
   matchLabels:
     app: nginx
  template:
   metadata:
     labels:
       app: nginx
    spec:
     containers:
      - name: nginx
       image: nginx
       ports:
        - containerPort: 80
       volumeMounts:
         - name: pvc-nas
           mountPath: /data
     volumes:
        - name: pvc-nas
         persistentVolumeClaim:
           claimName: pvc-nas
```

### 3.3.4. Mount a dynamically provisioned NAS

#### volume

You can create a subdirectory in an Apsara File Storage NAS (NAS) file system and map the subdirectory to a dynamically provisioned persistent volume (PV) for applications. This topic describes how to mount a dynamically provisioned NAS volume.

#### Prerequisites

- A Container Service for Kubernetes (ACK) cluster is created and the FlexVolume plug-in is installed in the cluster. For more information, see Create an ACK managed cluster.
- The alicloud-nas-controller component is deployed in the cluster. For more information, see Inst all and upgrade FlexVolume.

#### Precautions

If the securityContext.fsgroup parameter is set in the application template, kubelet performs the chmod or chown operation after the volume is mounted, which increases the time consumption.

**Note** For more information about how to speed up the mounting process when the securityContext.fsgroup parameter is set, see Why does it require a long time to mount a NAS volume?.

#### Create a dynamically provisioned NAS volume

1. Configure a StorageClass.

#### Sample code:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
    name: alicloud-nas
mountOptions:
    - nolock,tcp,noresvport
    vers=3
parameters:
    server: "23a9649583-i****.cn-shenzhen.nas.aliyuncs.com:/nasroot1/"
    driver: flexvolume
provisioner: alicloud/nas
reclaimPolicy: Delete
```

Parameter	Description
mountOptions	The mount options of the PV. The NAS volume is mounted based on the specified mount options.
server	The list of mount targets that are used to provision the PV. The format is <i>nfsurl1:/path1,nfsurl2:/path2</i> . When multiple servers are specified, the PV provisioned by this StorageClass uses the servers in a round robin manner. For Extreme NAS file systems, the path must start with <i>/share</i> .
driver	FlexVolume and NFS are supported. The default driver is NFS.
reclaimPolicy	<ul> <li>The reclaim policy of the PV. We recommend that you set the value to Retain.</li> <li>If you set the value to Delete, the name of the subdirectory mapped to the PV is automatically changed after you delete the PV. For example, <i>path-name</i> is changed to <i>archived-path-name</i>.</li> <li>If you want to delete the subdirectory in the NAS file system, set archiveOnDelete to <i>false</i> in the StorageClass configurations.</li> </ul>

2. Use the dynamically provisioned NAS volume in a StatefulSet.

Create a Service and a StatefulSet by using the following sample code:

apiVersion: v1 kind: Service metadata: name: nginx labels: app: nginx spec: ports: - port: 80 name: web clusterIP: None selector: app: nginx \_\_\_ apiVersion: apps/v1 kind: StatefulSet metadata: name: web spec: selector: matchLabels: app: nginx serviceName: "nginx" replicas: 5 volumeClaimTemplates: - metadata: name: html spec: accessModes: - ReadWriteOnce storageClassName: alicloud-nas resources: requests: storage: 2Gi template: metadata: labels: app: nginx spec: containers: - name: nginx image: nginx:alpine volumeMounts: - mountPath: "/data" name: html

## 3.3.5. Use NAS volumes for shared persistent

#### storage

You can use an Apsara File Storage NAS (NAS) volume to persist data and share the data among multiple pods. This topic describes how to use a NAS file system to persist and share data.

#### Prerequisites

- Create an ACK managed cluster.
- Connect to ACK clusters by using kubectl.
- A NAS file system is created in the NAS File System console. For more information, see Mount an NFS file system on a Linux ECS instance. The NAS file system and the cluster are deployed in the same zone.
- A mount target is added to the NAS file system. For more information, see . The NAS file system and the cluster are deployed in the same virtual private cloud (VPC).

#### Context

If a NAS file system is mounted on multiple pods, the data in the file system is shared among the pods. The application must be able to synchronize data across all pods when data modifications are made by multiple pods.

Scenarios:

- Your application requires high disk I/O.
- You need a storage service that offers higher read and write throughput than Object Storage Service (OSS).
- You want to share files across hosts. For example, you want to use a NAS file system as a file server.

Procedure

- 1. Create a NAS file system and create a mount target.
- 2. Create a persistent volume (PV) and a persistent volume claim (PVC).

The following section describes how to create a PV and a PVC by using the *FlexVolume* plug-in provided by Alibaba Cloud and then mount a NAS file system.

#### Create a PV

1. Create a file named *pv-nas.yaml*.

```
apiVersion: v1
kind: PersistentVolume
metadata:
 name: pv-nas
 labels:
   alicloud-pvname: pv-nas
spec:
 capacity:
   storage: 5Gi
 accessModes:
    - ReadWriteMany
 flexVolume:
   driver: "alicloud/nas"
   options:
      server: "***-**.cn-hangzhou.nas.aliyuncs.com" #Replace the value with the mount
target.
     path: "/k8s1"
     vers: "4.0"
```

Parameter	Description
alicloud-pvname	The name of the PV.
server	The mount target of the NAS file system. To obtain the mount target, log on to the NAS File System console. In the left-side navigation pane, click <b>File System List</b> , find the created file system, and then click <b>Management</b> in the <b>Operations</b> column. Click the <b>Mounting Use</b> tab and copy the mount address in the <b>Mount</b> <b>Target</b> column.
path	The mounted directory of the NAS file system. You can specify a subdirectory of a NAS file system. If no subdirectories exist, the system automatically creates a subdirectory.
vers	The version number of the Network File System (NFS) protocol. This parameter is optional. Valid values: 3 and 4.0. Default value: 3.
	The access permissions on the mounted directory. This parameter is optional. By default, this parameter is left empty.
mode	<ul> <li>Note</li> <li>You are not allowed to grant permissions to access the root directory of a NAS file system.</li> <li>If you set the mode parameter for a NAS file system that stores a large amount of data, the process of mounting the NAS file system may be time-consuming or even fail. We recommend that you leave this parameter empty.</li> </ul>

2. Run the following command to create a PV:

kubectl create -f pv-nas.yaml

#### Expected result

- 1. Log on to the ACK console.
- 2. In the left-side navigation pane, click **Clusters**.
- 3. On the **Clusters** page, find the cluster that you want to manage, and click the name of the cluster or click **Details** in the **Actions** column.

4. In the left-side navigation pane of the details page, choose **Volumes > Persistent Volumes**. Verify that the newly created PV is displayed.

#### Create a PVC

When you create a PVC of the NAS type, set the selector parameter to configure how to select the PV to which the PVC is bound.

1. Create a file named *pvc-nas.yaml*.

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
   name: pvc-nas
spec:
   accessModes:
    - ReadWriteMany
   resources:
      requests:
      storage: 5Gi
   selector:
      matchLabels:
      alicloud-pvname: pv-nas
```

2. Run the following command to create a PVC:

kubectl create -f pvc-nas.yaml

#### **Expected** result

- 1. Log on to the ACK console.
- 2. In the left-side navigation pane, click **Clusters**.
- 3. On the **Clusters** page, find the cluster that you want to manage, and click the name of the cluster or click **Details** in the **Actions** column.
- 4. In the left-side navigation pane of the details page, choose **Volumes > Persistent Volume Claims**. Verify that the newly created PVC is displayed.

#### Create an application

1. Create a file named *nas.yaml*.

```
apiVersion: apps/v1
kind: Deployment
metadata:
 name: nas-static
 labels:
   app: nginx
spec:
 replicas: 2
 selector:
   matchLabels:
    app: nginx
  template:
   metadata:
     labels:
       app: nginx
    spec:
     containers:
      - name: nginx
       image: nginx
       ports:
       - containerPort: 80
       volumeMounts:
         - name: pvc-nas
           mountPath: "/data"
     volumes:
       - name: pvc-nas
         persistentVolumeClaim:
           claimName: pvc-nas
```

2. Run the following command to deploy an application:

kubectl create -f nas.yaml

#### **Expected** result

- 1. Log on to the ACK console.
- 2. In the left-side navigation pane, click **Clusters**.
- 3. On the **Clusters** page, find the cluster that you want to manage, and click the name of the cluster or click **Applications** in the **Actions** column.
- 4. In the left-side navigation pane of the **cluster** details page, choose **Workloads > Deployments**. Verify that the newly created application is displayed.

#### Verify data sharing

1. Run the following command to query the pods that run the application.

kubectl get pod

#### Expected output:

```
        NAME
        READY
        STATUS
        RESTARTS
        AGE

        nas-static-f96b6b5d7-r***
        1/1
        Running
        0
        9m

        nas-static-f96b6b5d7-r***
        1/1
        Running
        0
        9m
```

2. Run the following commands to query the files in the /data path:

kubectl exec nas-static-f96b6b5d7-r\*\*\*\* ls /data

#### Expected output:

```
kubectl exec nas-static-f96b6b5d7-w**** ls /data
```

**?** Note The output indicates that no file exists in the /data path.

3. Run the following command to create a file named *nas* in the */data* path of a pod:

```
kubectl exec nas-static-f96b6b5d7-r**** touch /data/nas
```

4. Query files in the pods.

Run the following command to query files in the /data path of one pod:

kubectl exec nas-static-f96b6b5d7-r\*\*\*\* ls /data

Expected output:

nas

Run the following command to query files in the /data path of the other pod:

kubectl exec nas-static-f96b6b5d7-w\*\*\*\* ls /data

#### Expected output:

nas

**?** Note The file was created in the */data* path of one of the pods. However, you can find the file in the */data* path of both pods. This indicates that the pods share the NAS volume.

#### Verify data persistence

1. Run the following commands to delete all pods of the application:

kubectl delete pod nas-static-f96b6b5d7-r\*\*\*\* nas-static-f96b6b5d7-wthmb

Expected output:

pod "nas-static-f96b6b5d7-r\*\*\*\*" deleted
pod "nas-static-f96b6b5d7-w\*\*\*\*" deleted

2. Run the following command to view how the pods are deleted and recreated:

kubectl get pod -w -l app=nginx

Expected output:

NAME	READY	STATUS	RESTARTS	AGE
nas-static-f96b6b5d7-r****	1/1	Running	0	27m
nas-static-f96b6b5d7-w****	1/1	Running	0	27m
nas-static-f96b6b5d7-r****	1/1	Terminating	0	28m
nas-static-f96b6b5d7-w****	0/1	Pending	0	0s
nas-static-f96b6b5d7-w****	0/1	Pending	0	0s
nas-static-f96b6b5d7-w****	0/1	ContainerCreating	0	0s
nas-static-f96b6b5d7-w****	1/1	Terminating	0	28m
nas-static-f96b6b5d7-n****	0/1	Pending	0	0s
nas-static-f96b6b5d7-n****	0/1	Pending	0	0s
nas-static-f96b6b5d7-n****	0/1	ContainerCreating	0	0s
nas-static-f96b6b5d7-r****	0/1	Terminating	0	28m
nas-static-f96b6b5d7-w****	0/1	Terminating	0	28m
nas-static-f96b6b5d7-r****	0/1	Terminating	0	28m
nas-static-f96b6b5d7-r****	0/1	Terminating	0	28m
nas-static-f96b6b5d7-w****	1/1	Running	0	10s
nas-static-f96b6b5d7-w****	0/1	Terminating	0	28m
nas-static-f96b6b5d7-w****	0/1	Terminating	0	28m
nas-static-f96b6b5d7-n****	1/1	Running	0	17s

3. Run the following command to query the newly created pods:

kubectl get pod

#### Expected output:

NAME	READY	STATUS	RESTARTS	AGE
nas-static-f96b6b5d7-n****	1/1	Running	0	21s
nas-static-f96b6b5d7-w****	1/1	Running	0	21s

#### 4. Query files in the pods.

Run the following command to query files in the /data path of one pod:

kubectl exec nas-static-f96b6b5d7-n\*\*\*\* ls /data

#### Expected output:

nas

Run the following command to query files in the /data path of the other pod:

```
kubectl exec nas-static-f96b6b5d7-w**** ls /data
```

#### Expected output:

nas

**?** Note The *nas* file still exists. This indicates that data is persisted to the NAS volume.

#### 3.3.6. FAQ about NAS volumes

This topic provides answers to some frequently asked questions about Apsara File Storage NAS (NAS) volumes.

• Why does it require a long time to mount a NAS volume?

- Why does a timeout error occur when I mount a NAS volume?
- Why does the system prompt chown: option not permitted when I mount a NAS volume?
- What do I do if I fail to mount a NAS volume?
- What do I do if the task queue of alicloud-nas-controller is full and PVs cannot be created when I use a dynamically provisioned NAS volume?

#### Why does it require a long time to mount a NAS volume?

Symptom:

It requires a long time to mount a NAS volume.

Cause:

If the securityContext.fsgroup parameter is set in the application template, kubelet performs the chmod or chown operation after the volume is mounted, which increases the time consumption.

Solution:

- If the securityContext.fsgroup parameter is set in the application template, delete the fsgroup parameter in the securityContext section.
- If you want to configure the user ID (UID) and mode of the files in the mounted directory, you can manually mount the directory to an Elastic Compute Service (ECS) instance. You can then perform ch own and chmod operations through a CLI and provision the NAS volume through the FlexVolume plug-in. For more information about how to provision NAS volumes through FlexVolume, see Mount a statically provisioned NAS volume and Mount a dynamically provisioned NAS volume.
- For clusters of Kubernetes 1.20 or later, you can set the fsGroupChangePolicy parameter to OnRootMismatch. This way, the chmod or chown operation is performed only during the firsttime launch of the pod that uses the volume. For more information, see Set the security context for a container.

#### Why does a timeout error occur when I mount a NAS volume? Symptom:

A timeout error occurred when you mount a NAS volume.

Cause:

The mount target of the NAS file system and the cluster are not in the same virtual private cloud (VPC). Solution:

Select a NAS file system whose mount target is in the same VPC as the cluster.

## Why does the system prompt chown: option not permitted when I mount a NAS volume?

Symptom:

The system prompts chown: option not permitted when you mount a NAS volume.

Cause:

Your container does not have permissions to use the specified NAS volume.

Solution:

Launch the container with root privileges.

#### What do I do if I fail to mount a NAS volume?

#### Symptom:

Your attempt to mount a NAS PV failed and the system prompts the following error:

Unable to mount volumes for pod "dp-earnings-pod\_default(906172c6-3d68-11e8-86e0-00163e00\*\* \*\*)": timeout expired waiting for volumes to attach/mount for pod "default"/"dp-earnings-po d". list of unattached/unmounted volumes=[vol1 vol2]

#### Cause:

The FlexVolume plug-in is not installed.

Solution:

Install the FlexVolume plug-in. For more information, see Install and upgrade FlexVolume.

# What do I do if the task queue of alicloud-nas-controller is full and PVs cannot be created when I use a dynamically provisioned NAS volume?

#### Symptom:

When you use a dynamically provisioned NAS volume, if the speed of subdirectory creation is faster than the speed of subdirectory deletion, the task queue of alicloud-nas-controller may be full and therefore PVs cannot be created.

#### Cause:

The reclaimPolicy parameter is set to Delete and the archiveOnDelete parameter is set to false in the configuration of the StorageClass that mounts the dynamically provisioned NAS volume.

#### Solution:

Set archiveOnDelete to true. This way, when a PV is deleted, only the name of the mounted subdirectory in the NAS file system is modified. The files in the subdirectory are not deleted.

You must delete these files yourself. For example, you can configure a node to automatically delete files in the root directory by schedule, or start multiple pods to concurrently delete files of specific formats in subdirectories.

# 3.4. Mount SMB file systems to Windows containers

You can mount Server Message Block (SMB) file systems of Apsara File Storage NAS (NAS) to Windows containers that run in a Container Service for Kubernetes (ACK) cluster. This topic describes how to mount SMB file systems to Windows containers.

#### Prerequisites

- Create a Windows node pool.
- Step 2: Select a type of cluster credentials.
- In the NAS console, create an SMB file system in the virtual private cloud (VPC) where the ACK cluster is deployed, and create a mount target for the SMB file system. For more information, see Mount an SMB file system on Windows.

#### Step 1: Create a PV and a PVC

1. Use the following YAML template to create a persistent volume (PV) and a persistent volume claim (PVC).

The following table describes the required parameters in the PV template.

Parameter	Description
driver	The driver that is used to mount the SMB file system. Set the value to alicloud/smb.exe.
server	The domain name of the mount target for the SMB file system. The mount target must be in the same VPC as the ACK cluster.
path	The path where the SMB file system is mounted. Set the value to \myshare or a subdirectory that starts with \myshare.
user	The username that is used to log on to a node. We recommend that you use workgroup\administrator.
password	The password that is used to log on to a node.

YAML template that is used to create a PV	YAMLtemplate

```
ML template that is used to create a PVC
```

```
apiVersion: v1
 kind: PersistentVolume
 metadata:
  labels:
    alicloud-pvname: pv-smb
  name: pv-smb
 spec:
  accessModes:
   - ReadWriteMany
   capacity:
    storage: 5Gi
   flexVolume:
    driver: alicloud/smb.exe
    options:
      path: \myshare\test
      server: 25f3f4819c-eak52.cn-shenzhen.nas.aliyuncs.com
      user: workgroup\administrator
       password: ***
   persistentVolumeReclaimPolicy: Retain
                                                                                       2.
Run the kubectl get pvc |grep pvc-smb command to view the newly created PVC.
```

The following output is returned:

pvc-smb	Bound	pv-smb	5Gi	RWX
24h				

#### Step 2: Deploy an application

File system mounting Mount a file s ystem on a pod in Kubernetes

1. Use the following YAML template to deploy an application.

```
apiVersion: apps/v1
kind: Deployment
metadata:
 name: nginx-smb
 namespace: default
spec:
 selector:
   matchLabels:
     app: nginx-smb
  template:
   metadata:
     labels:
       app: nginx-smb
    spec:
     replicas: 2
     tolerations:
     - effect: NoSchedule
       key: os
       operator: Equal
       value: windows
     containers:
      - args:
       - -Command
       - start-sleep 10000
       command:
       - pwsh.exe
       image: registry.cn-hangzhou.aliyuncs.com/acs/flexvolume:v1.16.9.7be0fa0-windows
1809
       imagePullPolicy: IfNotPresent
       name: nginx
       volumeMounts:
       - mountPath: /data
         name: pvc-nas
     restartPolicy: Always
     volumes:
      - name: pvc-nas
       persistentVolumeClaim:
         claimName: pvc-smb
```

2. Run the kubectl get pod command to view the state of the application.

The following output is returned:

NAME	READY	STATUS	RESTARTS	AGE
nginx-smb-965fb4597-jz6fv	1/1	Running	0	95s
nginx-smb-965fb4597-zvbhk	1/1	Running	0	42s

If the application is in the Running state, the application is created.

# 4.Mount a file system across VPCs or Alibaba Cloud accounts

# 4.1. Mount a file system across VPCs or regions

This topic describes how to use Cloud Enterprise Network (CEN) to mount a file system across virtual private clouds (VPCs) or regions.

#### Prerequisites

- A file system is created. For more information, see Manage file systems.
- A mount target is created. For more information, see Create a mount target.

#### Context

By default, a file system can be mounted on an Elastic Compute Service (ECS) instance only if the instance and mount target reside in the same VPC. If the mount target and the ECS instance reside in different VPCs, you can use CEN to establish a connection between the VPCs. You can then enable a cross-VPC mount for the file system.

The following procedure uses an example to describe how to establish a connection between two VPCs that belong to the same Alibaba Cloud account. In this example, VPC 1 and VPC 2 are used and attached to the same CEN instance. The procedure applies regardless of whether VPC 1 and VPC 2 reside in the same region.



#### Step 1: Create a CEN instance

1. Log on to the CEN console.

- 2. In the left-side navigation pane, choose Instances > Create CEN Instance.
- 3. In the Attach Network section of the **Create CEN Instance** dialog box, click the **Your Account** tab and set the parameters.

Parameter	Description
Network Type	Select VPC.
Region	Select the region where the VPC resides. In this example, select the region where VPC 1 resides.
Networks	Select the VPC that you want to attach to the CEN instance. In this example, select VPC 1.

#### Step 2: Attach a VPC to the CEN instance

- 1. Log on to the CEN console.
- 2. In the left-side navigation pane, click **Instances**. On the page that appears, find the CEN instance that you created and click **Manage** in the Actions column.
- 3. Choose Networks > Attach Network.
- 4. In the Attach Network dialog box, click the Your Account tab and set the parameters.

The following table lists the required parameters.

Parameter	Description
Network Type	Select VPC.
Region	Select the region where the VPC resides. In this example, select the region where VPC 2 resides.
Networks	Select the VPC that you want to attach to the CEN instance. In this example, select VPC 2.

5. Click OK.

#### Step 3: Mount the file system

After the preceding configurations are complete, mount a file system in one of the VPCs on an ECS instance in the other VPC.

- For more information about how to mount an NFS file system on a Linux ECS instance, see Mount an NFS file system on a Linux ECS instance.
- For more information about how to mount an SMB file system on a Windows ECS instance, see Mount an SMB file system on Windows.

# 4.2. Enable a cross-account mount for a file system

This topic describes how to use Cloud Enterprise Network (CEN) to enable a cross-account mount for a file system.

#### Prerenuisites

#### 

Before you enable a cross-account mount for a file system, the following requirements must be met:

- A file system is created. For more information, see Create a General-purpose NAS file system in the NAS console.
- A mount target is created. For more information, see Create a mount target.

#### Context

By default, you can mount a file system only on an Elastic Compute Service (ECS) instance that is owned by the same account as that of the file system. Assume that you have multiple Alibaba Cloud accounts and want to allow mutual access between a file system and an ECS instance from these different accounts. You must establish a connection between the VPCs that host the file system and the ECS instance.

You can use CEN to connect the VPCs that are owned by different accounts.

This topic describes how to attach VPC 1 of Account A and VPC 2 of Account B to the same CEN instance.



#### Step 1 Create a CEN instance

Use Account A to create a CEN instance.

- 1. Log on to the CEN console.
- 2. In the left-side navigation pane, click Instances, and then click Create CEN Instance.
- 3. In the **Create CEN Instance** dialog box, set the parameters.

The following table describes the required parameters.

Parameter	Description
Network Type	Select VPC from the drop-down list.
Region	The region where the network resides. Select the region where VPC 1 resides.
Networks	The network that you want to attach to the CEN instance. Select VPC 1.

4. Click OK.

Back up the CEN instance ID for subsequent operations.

## Step 2: Authorize an account to access the network of a different account

Use Account B to authorize Account A to attach VPC 2. For more information, see VPC authorization.

## Step 3: Use an account to attach a network that is owned by a different account

Use Account A to attach VPC 2.

- 1. Log on to the Cloud Enterprise Network console.
- 2. In the left-side navigation pane, click **Instances**. On the page that appears, find the CEN instance and click **Manage**.
- 3. On the Networks tab, click Attach Network.
- 4. In the Attach Network dialog box, click Different Account and set the following parameters.

Parameter	Description
Owner Account	The ID of the account that owns the target network. Enter the ID of Account B.
Network Type	Select VPC.
Region	The region where the network resides. Select the region where VPC 2 resides.
Networks	The network that you want to attach. Select VPC 2.

5. Click OK.

#### Mount a file system

After the configuration is complete, you can perform a cross-account mount on a file system.

- For more information about how to mount an NFS file system on a Linux ECS instance, see Mount an NFS file system on a Linux ECS instance.
- For more information about how to mount an SMB file system on a Windows ECS instance, see Mount an SMB file system on Windows.

# 5.Access file systems in on-premises data centers5.1. Access a NAS file system from a datacenter by using NAT Gateway

This topic describes how to access an Apsara File Storage NAS file system from a data center by using a NAT gateway.

#### Context

You can mount a file system only on an ECS instance that resides in the same region as the file system. For example, a Network File System (NFS) or Server Message Block (SMB) file system that you create in the China (Hangzhou) region can be mounted only on an ECS instance that resides in the China (Hangzhou) region. You cannot mount the file system on an ECS instance that resides in a different region such as the China (Qingdao) region or on a local server. To implement a file system mount across regions or in a data center, you must use Express Connect to establish a connection between Virtual Private Clouds (VPCs) or between a VPC and a data center. However, this connection significantly increases the cost of mounting the file system.

If a VPN gateway is deployed in your data center, we recommend that you use Alibaba Cloud VPN Gateway to connect your data center to NAS. For more information, see Access an Apsara File Storage NAS file system from a local data center by using VPN Gateway.

If you only need to upload a small amount of data from your data center to NAS, we recommend that you use NAT Gateway to establish a connection.

The following figure shows the network topology that is adopted when NAT Gateway is used to establish a connection between a data center and NAS.



- Advantage: easy to configure
- Disadvantage:
  - In terms of security, a user who has an Elastic IP address (EIP) can create a mount target that relates to the EIP because connections are established between EIPs and VPCs.
  - Each combination of an EIP and port can be specified only for one mount target. If you want to access multiple mount targets at the same time, you must create multiple EIPs.

#### Create a file system and a mount target

- 1. Log on to the NAS console.
- 2. Create a file system. For more information, see Create a General-purpose NAS file system in the NAS console.
- 3. Create a mount target in a VPC. For more information, see Create a mount target.

#### Configure a NAT gateway

You can perform the following steps to mount a NAS file system on a Windows or Linux host that is connected to the Internet. After the file system is mounted, you can upload files to or download files from the file system on the host.

- 1. Log on to the VPC console.
- 2. Create a NAT gateway. For more information, see Create an Internet NAT gateway.

(?) Note The VPC in which the NAT gateway resides must be same as the VPC in which the NAS file system resides.

- 3. Bind an EIP to the NAT gateway For more information, see Apply for an EIP.
- 4. Create a DNAT entry. For more information, see Manage a DNAT entry.

You must set the following parameters:

- Public IP Address: specifies the public IP address that is generated when you create an EIP.
- **Private IP Address**: specifies the IP address of the mount target for the file system. To obtain the IP address, you can ping the mount target from the ECS instance on which the file system is mounted.

ping file-system-id.region.nas.aliyuncs.com

- **Port** : We recommend that you select **All Ports**. You can also select a port for your NFS or SMB file system.
- 5. Mount the file system.

- To mount an NFS file system, you must first install an NFS client. For more information, see Install an NFS client.
- Before you mount an SMB file system, make sure that the Workstation and TCP/IP NetBIOS Helper services are started in the Windows system on which you want to mount the SMB file system. For more information, see Start the Workstation and TCP/IP NetBIOS Helper services.
- The default port 445 of the SMB protocol is a high-risk port. By default, the port is disabled by your internet service provider (ISP). If you want to access NAS from a data center by using NAT Gateway over the Internet, you must configure port forwarding in the data center. Procedure:
  - a. Configure a DNAT entry to map Port 445 of NAS to Port 4456 for the EIP of a NAT gateway.
  - b. Run the netsh tool on a local Windows client to forward network traffic from Port 445 to Port 4456.
  - c. Mount the file system.

<sup>?</sup> Note

• If you want to mount an NFS file system, run the following command:

mount -t nfs4 10.10.10.1:/ /mnt

- 10.10.10.1 is the public IP address that is generated when you create an EIP. Replace the IP address based on your business requirements.
- /mnt is the directory on which you want to mount the file system. Replace the directory based on your business requirements.
- If you want to mount an SMB file system, run the following command:

```
net use D: \\10.10.1\myshare
```

- D: is the letter of the destination drive on which you want to mount a file system. Replace the drive letter based on your business requirements.
- 10.10.10.1 is the public IP address that is generated when you create an EIP. Replace the IP address based on your business requirements.
- myshare is the name of the shared SMB directory. You cannot change the name.
- 6. Verify the mount result.
  - NFS file system

If the result that is similar to the following information appears after you run the mount command, the mount is successful. You can read data from and write data to the files of the NFS file system.



• SMB file system

If you can access the SMB file system from your local file manager, the mount is successful. You can read data from and write data to the files of the SMB file system.

Ny myshare									
🚱 🖓 v 🖳 🔹 🗰 🖬 🖬 🖬	- <b>6</b>	Search myshare							
Organize 🔻 New folder									
☆ Favorites	Name *	Date modified	Туре						
🧾 Desktop	🚮 fio-2.18-x64.zip	2017/3/22 17:17	Compressed (zippe						
鷆 Downloads	iotest.txt	2017/7/14 10:03	Text Document						

## Differences between the NAT Gateway solution and the VPN Gateway solution

The following table describes the differences between the two solutions.

ltem	NAT Gateway solution	VPN Gateway solution
Configuration	Easy: You can configure all settings in the Alibaba Cloud Management Console.	Complex: You must configure a VPN gateway in the Alibaba Cloud console and configure a client-side VPN gateway in a data center.

ltem	NAT Gateway solution	VPN Gateway solution
Data security	Low	High
Flexibility	Low. Each EIP can be mapped to only one mount target.	High. You can access all NAS mount targets at the same time. EIPs are not required in this solution.
Scenarios	Establish temporary connections to transfer a small amount of data.	Establish a long-term connection between a data center and NAS.

## 5.2. Access an Apsara File Storage NAS file system from a local data center by using VPN Gateway

This topic describes how to access an Apsara File Storage NAS file system from a local data center by configuring a VPN gateway.

#### Context

You can only mount a file system on an ECS instance that resides in the same region as the file system. For example, an NFS or SMB file system that you create in China (Hangzhou) can only be mounted on an ECS instance that resides in China (Hangzhou). You cannot mount a file system that resides in China (Hangzhou) on a local data center or on an ECS instance that resides in a different region such as China (Qingdao). To resolve these issues, you can establish a connection over an Express Connect circuit. To enable a file system mount on a local data center, you can establish the connection between the data center and the Virtual Private Cloud (VPC) where the file system resides. To enable a cross-region file system mount, you can establish the connection between the VPC where the ECS instance resides and the VPC where the file system resides. However, high costs incur for establishing the connection.

Instead, we recommend that you use VPN Gateway to enable communication between a local data center and a VPC or between VPCs that reside in different regions. With VPN Gateway, you can mount a file system on the following target instances:

- A server that resides in a local data center
- An ECS instance that resides in a different region different from the region of the file system
   If you have created a VPN gateway on an ECS instance in one VPC, you need to create another VPN
   gateway in the other VPC. Then, you need to establish a connection between the two VPN
   gateways. For more information about detailed operations, see Enable a cross-region mount (one
   VPN gateway available). If no VPN gateway exists in your environment, we recommend that you
   create VPN gateways in the two VPCs and connect the gateways. For more information about
   detailed operations, see Enable a cross-region mount (no VPN gateway available).

The following figure shows the topology that is adopted when VPN gateways are used.



The advantages and disadvantages are listed as follows:

- Advantages
  - Fixes all connectivity issues.
  - Provides secure access by using IPsec to encrypt data in transit.
  - Compared with Express Connect, VPN Gateway helps you reduce a large number of costs.
- Disadvantages

The Internet bandwidth and latency between a local data center and a VPC or between VPCs restrict I/O performance of a file system over a VPN connection.

#### Mount a file system on a server that resides in a local data center

- 1. Create a file system and mount target.
  - i. Log on to the Apsara File Storage NAS console.
  - ii. Create a file system. For more information, see Create a General-purpose NAS file system in the NAS console.
  - iii. Create a mount target of the VPC type. For more information, see Create a mount target.
- 2. Create a connection between the VPC and your local data center. For more information, see Connect a data center to a VPC.
- 3. Verify the connection between a server that resides in the local data center and an ECS instance or a mount target that resides in the VPC.

Log on to an ECS instance that does not have an Internet IP address. On the ECS instance, use the **ping** command to **ping** the internal IP address of a server that resides in the local data center and verify the connection.

4. After you confirm the connection by using the ping command, you can mount a file system that resides in the VPC on a server that resides in the local data center. For more information, see Mount a file system.

#### Enable a cross-region mount (one VPN gateway available)

The following example shows a practical scenario of two VPCs named VPC 1 and VPC 2 that reside in different regions.

- 1. Create a file system and mount target.
  - i. Log on to the Apsara File Storage NAS console.

- ii. Create a file system. For more information, see Create a General-purpose NAS file system in the NAS console.
- iii. Create a mount target of the VPC type. For more information, see Create a mount target.Create a mount target in VPC 1.
- 2. In VPC 2, create a VPN gateway on an ECS instance as a customer gateway.
  - ? Note
    - You must specify an Internet IP address for the ECS instance to connect to the VPN gateway that resides in VPC 1.
    - For more information about how to create a VPN gateway on an ECS instance, see tutorials such as Using StrongSwan for IPsec VPN on CentOS 7.
- 3. Establish a connection between VPN gateways that reside in VPC 1 and VPC 2, respectively.
  - i. Log on to the VPC console.
  - ii. Create a VPN connection to enable communication between VPN gateways that resides in VPC 1 and VPC 2, which you created in Step 2. For more information, see Create an IPsec connection.
- 4. Configure static routes on other ECS instances that reside in VPC 2. For more information, see Configure routes on a VPN gateway. The required settings are described as follows.

**Dest inat ion CIDR Block** specifies the private classless inter-domain routing (CIDR) Block of VPC 1.**Next Hop** specifies the customer gateway that resides in VPC 2.

5. Verify the connection between VPC 1 and an ECS instance (or mount target) that resides in VPC 2.

Log on to an ECS instance that resides in VPC 1, use the **ping** command to **ping** the IP address of an ECS instance that resides in VPC 2, and verify the connection.

6. After you confirm the connection by using the ping command, you can mount a file system that resides in VPC 1 on an ECS instance that resides in VPC 2. For more information, see Mount a file system.

#### Enable a cross-region mount (no VPN gateway available)

The following example shows a practical scenario of two VPCs named VPC 1 and VPC 2 that reside in different regions.

- 1. Create a file system and mount target.
  - i. Log on to the Apsara File Storage NAS console.
  - ii. Create a file system. For more information, see Create a General-purpose NAS file system in the NAS console.
  - iii. Create a mount target of the VPC type. For more information, see Create a mount target.

Create a mount target in VPC 1.

- 2. Establish a connection between VPN gateways that reside in VPC 1 and VPC 2, respectively.
  - i. Log on to the VPC console.
  - ii. Create VPN gateways in VPC 1 and VPC 2, respectively. For more information, see Create a VPN gateway.

iii. Create customer gateways in VPC 1 and VPC 2, respectively. For more information, see Create a customer gateway. The required settings are described as follows.

**IP Address** specifies an IP address for the VPN gateway that resides in VPC 1 and a different IP address for the VPN gateway that resides in VPC 2.

- iv. Configure routes for VPN gateways that reside in VPC 1 and VPC 2, respectively. For more information, see Configure routes for a VPN gateway.
  - The following information is important when you configure routes for the VPN gateway that resides in VPC 1. Destination CIDR Block specifies the private CIDR block for VPC 2. Next Hop specifies the name of the customer gateway that resides in VPC 1.
  - The following information is important when you configure routes for the VPN gateway that resides in VPC 2. **Destination CIDR Block** specifies the private CIDR block for VPC 1. **Next Hop** specifies the name of the customer gateway that resides in VPC 2.
- 3. Verify the connection between VPC 1 and an ECS instance (or mount target) that resides in VPC 2.

Log on to an ECS instance that resides in VPC 1, use the **ping** command to **ping** the IP address of an ECS instance that resides in VPC 2, and verify the connection.

4. After you confirm the connection by using the ping command, you can mount a file system that resides in VPC 1 on an ECS instance that resides in VPC 2. For more information, see Mount a file system.

## 5.3. Access an SMB file system from a macOS client by using VPN Gateway

This topic describes how to mount a Server Message Block (SMB) file system on a macOS client and access the SMB file system by using the Kerberos protocol.

#### Prerequisites

- An SMB file system is created. For more information, see Manage file systems.
- A mount target in a virtual private cloud (VPC) is created. For more information, see Create a mount target.

#### Mount the SMB file system on a macOS client

1. Connect the macOS client to the VPC by using a virtual private network (VPN) gateway. For more information, see Connect a macOS client to a VPC.

When you create a Secure Sockets Layer (SSL) server, the Classless Inter-Domain Routing (CIDR) blocks of Local Network and Client Subnet cannot overlap with each other. Local Network specifies the CIDR block of the VPC. For more information, see Remote access from a Mac client. You can view the CIDR block on the VPC details page in the VPC console.

2. Check whether the macOS client can access the mount target of the SMB file system.

After the VPN gateway is set up between the VPC and the macOS client, run the **ping** command on the macOS client to ping the domain name of the mount target.

#### IT-C02WW0JRG8WN:Volumes \$ ping .nas.aliyuncs.com (172.31.16.146): 56 data bytes 64 bytes from 172.31.16.146: icmp\_seq=0 ttl=101 time=211.562 ms 64 bytes from 172.31.16.146: icmp\_seq=1 ttl=101 time=199.784 ms 64 bytes from 172.31.16.146: icmp\_seq=2 ttl=101 time=173.240 ms ^C ---- 3 packets transmitted, 3 packets received, 0.0% packet loss round-trip min/avg/max/stddev = 173.240/194.862/211.562/16.027 ms

**?** Note If you cannot ping the domain name of the mount target, run the ping command on an Elastic Compute Service (ECS) instance that resides in the same VPC as the SMB file system. You can then obtain the IP address of the mount target and use the IP address to mount the SMB file system.

- 3. Mount the SMB file system on the macOS client.
  - Mount the SMB file system on the macOS client by using graphical user interface (GUI)
    - a. In the menu bar of the macOS client desktop, choose Go > Connect to Server.



b. In the **Connect to Server** dialog box, enter the domain name of the mount target and click **Connect**.

	Connect t	o Server	
smb://	00-1614.00-6001-1	.nas.aliyuncs.com/my	share 🔽
Favorite Servers:			
+ - *~	?	Browse	Connect

c. In the Connect As section, select Guest, and then click Connect.



d. In the menu bar of the macOS client desktop, choose Go > Computer. Click the myshare disk to view the SMB file system that is mounted on the macOS client.

(?) Note After a file system is mounted, the macOS client reads all files that are stored in the file system. The **myshare** disk may be empty when the macOS client is reading the files. Wait until the read process is complete.

•••		🚍 myshare			
			Q Search		
				myshare	
	Name		Date Modified		
O Downloads	▶ 🖿 hg				
Recents	New folde				
		Document (2).txt			
Applications	New Text	Document.txt			
Documents					
🛄 Desktop					
Dropbox					
😭 fatestudio					
Creative Cloud Files					
lCloud Drive					
BOOTCAMP					
200004400602 fieldus rest: 1.na 4					
Network					
• Yellow					
<ul> <li>Blue</li> </ul>					
Red					
Purple					
Green					
Important					
O Work					
All Tags					
- All Tays					

 Mount the SMB file system on the macOS client by using command line interface (CLI) Run the mount\_smbf command to mount the SMB file system. The following code provides an example of the mount\_smbf command:

mount\_smbfs '//guest@nas-mount-point.nas.aliyuncs.com/myshare' /Volumes/myshare/

nas-mount-point.nas.aliyuncs.com is the domain name of the mount target in the VPC. The following figure shows a successful mount.



#### Access the SMB file system by using the Kerberos protocol

After an SMB file system is mounted on a macOS client based on NT LAN Manager (NTLM), the macOS client has all permissions on the SMB file system by default. To grant different permissions to different users, NAS allows you to authenticate users and control access to SMB file systems based on an Active Directory (AD) domain. You can perform the following steps to control access to the SMB file system based on an AD domain:

1. Configure an ECS instance as an AD domain controller and set up an AD domain.

- 2. Join the mount target of the SMB file system to the AD domain. For more information, see Add the mount target of an SMB file system to an AD domain.
- 3. Add the CIDR block of the SSL VPN network to a security group of the ECS instance. For more information, see Add security group rules.

Add rules for the following ports to a security group of the ECS instance. This ensures that the SMB file system can be mounted on the macOS client based on the AD domain.

- Domain Name System (DNS) port: UDP 53
- Kerberos port: TCP 88
- LDAP port: TCP 389
- LDAP Global Catalog port: TCP 3268
- 4. Change the DNS server that the macOS client uses to the AD domain controller.
  - i. Run the **ipconfig** command on the ECS instance to query the internal IP address of the AD domain controller.
  - ii. In the menu bar of the macOS client desktop, choose **Go > Network**.
  - iii. In the **Network** dialog box, set the DNS server of the macOS client to the internal IP address of the AD domain controller.
- 5. Verify the connection between the macOS client and the AD domain.

On the macOS client, ping the AD domain controller. The following figure shows a successful mount.

```
IT-C02WW0JRG8WN:Volumes $ ping smb-hk.com
PING smb-hk.com (172.31.59.36): 56 data bytes
64 bytes from 172.31.59.36: icmp_seq=0 ttl=127 time=172.032 ms
64 bytes from 172.31.59.36: icmp_seq=1 ttl=127 time=173.217 ms
64 bytes from 172.31.59.36: icmp_seq=2 ttl=127 time=176.154 ms
^C
--- smb-hk.com ping statistics ---
3 packets transmitted, 3 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 172.032/173.801/176.154/1.733 ms
IT-C02WW0JRG8WN:Volumes
```

- 6. Use an AD domain identity to mount the SMB file system on the macOS client by using the Kerberos protocol.
  - i. Run the **kinit** command to verify the security of the AD domain identity. The following code provides an example of the kinit command:

kinit user@MYDOMAIN.COM

ii. Run the **klist** command to view the AD domain identity. The following code provides an example of the klist command:

klist

iii. Run the **kinit** command to use the AD domain identity to log on to the macOS client. The following code provides an example of the kinit command:

kinit

iv. Run the **mount\_smbfs** command to mount the SMB file system. The following code provides an example of the mount\_smbfs command:

mount\_smbfs //administrator@nas-mount-point.nas.aliyuncs.com/myshare /Volumes/mysha
re

**Note** If an error message mount\_smbfs: server rejected the connection: Authent ication error is returned, run the kinit command to verify the AD domain identity and remount the file system.

#### The following figure shows a successful mount.



After the successful mount, run the **klist** command. Two service principals are returned. The following figure gives an example of the klist command.

IT-C02WW0JRG8WN:Volum Credentials cache: AP Principal:		1
Issued	Expires	Principal
Apr 15 03:49:18 2020	Apr 15 13:49:12 2020	krbtgt/SMB-HK.COM@SMB-HK.COM
Apr 15 03:49:27 2020	Apr 15 13:49:12 2020	cifs/icn-hongkong.nas.aliyuncs.com@SMB-HK.COM

Onte SMB access control lists (ACLs) are not displayed on the macOS client. However, when you perform an operation on the SMB file system, the SMB server verifies the ACLs and then allows or denies the operation. You can configure the ACLs of the SMB file system when you mount the SMB file system on the AD domain controller.

## 6.Unmount a file system 6.1. Unmount a file system from a Linux ECS instance

If you no longer need to use a file system or the data in the file system, you can unmount the file system by using the NAS console. You can also unmount the file system by running the unmount command on the Linux ECS instance on which the file system is mounted.

#### Unmount a NAS file system by using the NAS console

- 1. Log on to the NAS console.
- 2. In the left-side navigation pane, choose File System > File System List.
- 3. On the File System List page, click the name of the file system that you want to unmount.
- 4. On the details page of the file system, click Mounting Use.
- 5. On the Mounting Use tab, find the mount target and click Unmount in the Actions column.
- 6. In the Unmount from ECS dialog box, confirm the information of the ECS instance and click OK.

## Unmount a NAS file system by running the unmount command on the associated ECS instance

- 1. Log on to the ECS console.
- 2. Connect to the ECS instance.
- 3. Run the umount /mnt command to unmount a Network File System (NFS) file system. Replace the /mnt directory with the actual value.

**?** Note We recommend that you do not specify other parameters in the unmount command or change the default values of these parameters.

When you unmount the file system, the error message device is busy may occur. In this case, you must perform the following steps to end the process that is accessing the file system:

- i. Install fuser.
  - For an ECS instance that runs CentOS, Red Hat Enterprise Linux (RHEL), or Alibaba Cloud Linux, fuser is preinstalled.
  - For an ECS instance that runs Ubuntu or Debian, run the apt install -y fuser command to install fuser.
- ii. Run the fuser -mv <Local directory of the mount target> command to view the ID of the process that is accessing the NAS file system.
- iii. Run the kill <pid> command to end the process.

(?) Note If the process is a kernel process, you can skip this step.

4. Run the mount -1 command to view the unmount result.

If the file system is not displayed in the unmount result, the file system is unmounted.

# 6.2. Unmount a file system from an ECS instance running Windows

This topic describes how to unmount an SMB file system from an ECS instance running Windows.

#### Procedure

- 1. Log on to the ECS console.
- 2. Open the command prompt and run the following command to unmount a file system.

net use D: /delete

In the preceding command, replace the drive letter D: with a drive letter specific to your environment. You can run the net use command to retrieve the drive letter of a mount point.

? Note

- You can run the **net use \* /delete** command to unmount each available file system one by one in Windows.
- You can run the **net use \* /delete /y** command to unmount all the available file systems without any confirmation in Windows.
- 3. You can run the net use command to view the unmount results.

If an SMB file system is not displayed in the results, it indicates that the file system is unmounted.

## 7.FAQ about mounting

- Mount an NFS file system on Linux
  - How do I modify the maximum number of concurrent NFS requests?
  - How do I create and mount a subdirectory of a NAS file system on Linux?
  - Laccidentally deleted a mount target on an ECS instance that runs Linux, and issues started to occur. How do I fix these issues?
  - How do I prevent the listening port of a file system that supports NFSv4.0 protocol from being mist aken for a Trojan horse?
  - Can I mount an NFS file system and SMB file system on the same ECS instance?
  - How do I prevent exceptions if multiple processes or clients concurrently write data to a log file?
  - Why does a file in a NAS file system belong to different owners when I query the file on two ECS instances?
  - Why am I unable to use a Linux client to read data from or write data to files that are named in Chinese characters in a file system?
  - Why is a 523 error returned when I run the ls command on a Linux client on which an NFS file system is mounted?
- Mount an SMB file system on Windows
  - How do I fix the error that occurs when I mount an SMB file system on Windows?
  - Why am I unable to mount an SMB file system on a Windows operating system that is later than Windows Server 2016?
  - Why am I unable to mount an SMB file system?
  - Why is a mounted SMB directory visible only to an administrator?
  - Why is Internet Information Services (IIS) unable to load the files of an SMB volume on Windows Server 2016?
  - How do I fix the error that occurs when I use IIS to access a NAS file system?
  - If the system fails to terminate processes that connect to a file system, how do I clear the handles that are exposed by a client?
- Mount an SMB file system on Linux
  - Why am I unable to mount an SMB file system on Linux?
  - How do I improve the performance of an SMB file system that is mounted on Linux?
  - Why does file migration and replication take a long time when I mount an SMB file system on Linux?
  - How do I fix the Permission denied error when I access an SMB file system on Linux?
  - How do I rename the files of an SMB file system by changing the letter case?
  - Why am I unable to change the owner of a file and access mode of a file or directory?
  - Why does the server not respond within 35 seconds when multiple clients concurrently access a file?
  - Why does the mount target of an SMB file system become unresponsive?
  - If the system fails to terminate processes that connect to a file system, how do I clear the handles that are exposed by a client?
- Mount an NFS file system on Windows

- How do I fix the errors that occur when I soft mount an NFS file system on Windows?
- How do I fix the error that occurs when I mount an NFS file system on Windows?
- How do I resolve the invalid device error that is returned when I try to rename a file on the Windows NFS client?
- FAQ about the noresvport parameter
  - Why do I need to mount a NAS file system by using the noresvport parameter?
  - How do I fix issues related to the noresvport parameter?
  - What happens when a network switchover or an HA switchover occurs on backend services?
  - Why do I need to remount a file system? Can I use an alternative solution?
- Access files
  - How are files prefixed with .nfs generated? How do I delete files prefixed with .nfs?
  - When Laccess a file in the directory of a NAS file system, the bind conn to session failed on NFSv4 server error message is returned. How do I fix this issue?

#### How do I modify the maximum number of concurrent NFS requests?

By default, the maximum number of concurrent requests from a Network File System (NFS) agent is 2. This reduces the performance of NFS file systems. We recommend that you set the maximum number to 128.

- Method 1
  - i. Install an NFS client. For more information, see Install an NFS client.
  - ii. Run the following commands to set the maximum number of concurrent NFS requests to 128.

echo "options sunrpc tcp\_slot\_table\_entries=128" >> /etc/modprobe.d/sunrpc.conf
echo "options sunrpc tcp max slot table entries=128" >> /etc/modprobe.d/sunrpc.conf

**?** Note The first time you install an NFS agent, run the preceding commands once with root permissions. You do not need to run the commands again.

iii. Use the following command to restart the ECS instance.

reboot

- iv. Mount the NAS file system to the ACK cluster. For more information, see Precautions.
- v. Use the following command to verify the results. If the value 128 is returned, the maximum number is modified.

cat /proc/sys/sunrpc/tcp\_slot\_table\_entries

- Method 2
  - i. Install an NFS client. For more information, see Install an NFS client.
  - ii. Run the following commands to set the maximum number of concurrent NFS requests to 128.

```
echo "options sunrpc tcp_slot_table_entries=128" >> /etc/modprobe.d/sunrpc.conf
echo "options sunrpc tcp max slot table entries=128" >> /etc/modprobe.d/sunrpc.conf
```

**?** Note The first time you install an NFS agent, run the preceding commands once with root permissions. You do not need to run the commands again.

- iii. Remount the file system. For more information, see Precautions.
- iv. Use the following command to verify the results.If the value 128 is returned, the maximum number is modified.

cat /proc/sys/sunrpc/tcp\_slot\_table\_entries

## How do I create and mount a subdirectory of a NAS file system on Linux?

Make sure that a file system is mounted. For more information, see Mount an NFS file system on a Linux ECS instance.

If you mount the */mnt* directory of the file system on a Linux ECS instance, the */mnt* directory is used as the root directory of the file system. You can create subdirectories in the */mnt* directory.

1. Create a subdirectory in the root directory of the NAS file system on the server, for example, a Linux ECS instance.

mkdir /mnt/subdir

2. Create a local directory on which you can mount the NAS file system.

mkdir /tmp/mnt

(?) Note After you create a local directory on a server, you can mount only one file system on the local directory. To mount multiple file systems, you must create multiple local directories.

3. Remount the file system.

```
sudo mount -t nfs -o vers=3,nolock,proto=tcp,rsize=1048576,wsize=1048576,hard,timeo=600
,retrans=2,noresvport file-system-id.region.nas.aliyuncs.com:/ /mnt
```

The following list describes the required fields. Replace the values of these fields with the actual values.

- *file-system-id.region.nas.aliyuncs.com*: specifies the endpoint of the mount target. To obtain the mount target, perform the following operations: Log on to the NAS console. On the File System List page, click the name of the file system. On the details page, click Mounting Use and copy the mount command.
- /subdir: specifies the subdirectory of the NAS file system.
- /tmp/mnt: specifies the local directory of the server.

#### I accidentally deleted a mount target on an ECS instance that runs Linux, and issues started to occur. How do I fix these issues?

Issue

A file system is mounted on a Linux ECS instance by using Mount Target A. However, the mount target is deleted from the NAS console before the file system is unmounted. As a result, issues occur on Linux. For example, the system gives slow responses or does not respond when you run commands.

- Solution
  - i. On the Linux ECS instance, press Ctrl+C to stop the commands that are being run.
  - ii. Run the mount command to view the mount information.

Obtain the mount directory from the mount information, for example, */mnt/data*, as shown in the following figure.

debugfs on /sys/kernel/debug type debugfs (rw,relatime)
mqueue on /dev/mqueue type mqueue (rw,relatime)
hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime)
sumpc on var/lib/nfs/rpc_pipefs type rpc_pipefs (rw,relatime)
tmpfs on /run/user/0 type tmpfs (rw,nosuid,nodev,relatime,size=773356k,mode=700)
8790746464 - Spy10.co. Assignment and a lignment count of an annual Spyne of all Sector and Annual Annual Systems (Annual System), and an annual system (Annual System), and an annual system).
and an analysis of the second s
ard.sore:repart, proto-top, t isso-100, retrans 2, rec-pp; c1 isstable-192, 100 0, 115, local, lock-sore, able-192, 100 0, 113)
1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
ard_normsympt.protoctop.time=100.retrans<2.rec-ups.elientable=192.501.0.115, host_host=101.001.0.115, host_host=102.501.0.115
8-11-11-1-tgp98.cn-hangzhou.nas.aliyuncs.com:/ on /mmt/data type nfs4 (rw,relatime,vers=4.0,rsize=1048576,wsize=1048576,namlen=
255, hard , noresyport , proto=tcp , timeo=600 , retrans=2 , sec=sys , clientaddr=191 🗰 🛽 115, local_lock=none , addr=192 . 💷 💷 117)

iii. Run the umount -f /mnt/data command to unmount the file system. Command syntax: unmount-f <Mount directory>

Note If the unmount -f <Mount directory> command fails to unmount the file system, run the umount -l <Mount directory> command.

After you unmount the file system, you can create a mount target to mount the file system.

## How do I prevent the listening port of a file system that supports NFSv4.0 protocol from being mistaken for a Trojan horse?

• Issue

After you mount a file system that supports the NFSv4.0 protocol on a compute node, the protocol listens to the random listening port 0.0.0.0. The netstat command cannot identify the process of the listening port.

This is because the listening port is a random port, and therefore the backend application of the listening port cannot be traced. As a result, the listening port is mistaken for a Trojan horse.

[root@iz	28. <b>1 - 11</b>	16:22.0. 12:0z /]#	netstat -tanpl   grep tcp	grep LISTEN	
tcp	0	0 0.0.0.0:111	0.0.0.°×	LISTEN	1/systemd
tcp	0	0 0.0.0.0:22	0.0.0:*	LISTEN	2159/sshd
					innng jænennas.aliyuncs.com:/ /mnt
[root@iz	z8vi ii		netstat -tanpl   grep tcp	grep LISTEN	
tcp	0	0 0.0.0.0:38862	0.0.0:*	LISTEN	-
tcp	0	0 0.0.0.0:111	0.0.0:*	LISTEN	1/systemd
tcp	0	0 0.0.0.0:22	0.0.0:*	LISTEN	2159/sshd
tcp6	0	0 :::41841	:::*	LISTEN	_

Cause

The NFSv4.0 protocol listens to this random port for *callback* operations. By default, the fs.nfs.nf s\_callback\_tcpport kernel parameter is set to 0. Therefore, the NFSv4.0 protocol listens to a random port. This random port does not cause security risks.

Solution

Before you mount the file system, you can specify a non-zero value for the <code>fs.nfs.nfs\_callback\_tc</code> pport parameter. The non-zero value is used as the port number of the callback port.

sudo sysctl fs.nfs.nfs\_callback\_tcpport=<port>

In the following example, the fs.nfs.nfs\_callback\_tcpport parameter is set to 45450. After you mount a file system that supports the NFSv4.0 protocol, the result that is returned by the netstat command indicates that the system listens to port 45450. The root user is used in this example, as shown in the following figure. Therefore, you do not need to use sudo to run the systel command.

[root@i:	z8thrhhri	16to2a8uox1toz ∕]# net	stat -tanpl   grep tcp	grep LISTEN	
tcp	0	0 0.0.0.0:111	0.0.0.÷	LISTEN	1/systemd
tcp	0	0 0.0.0.0:22	0.0.0:×	LISTEN	2159/sshd
[root@i:	z8( 🖬 🖬 I	🖆 î 💼 Latioz /]# sys	ctl fs.nfs.nfs_callback_	_tcpport=45450	
fs.nfs.n	nfs_call	back_tcpport = 45450			
[root@i:	z8(` <b>```</b>	🎬 î 🛍 Luitoz /]# mou	nt -t nfs -o vers=4.0 91	.56	🛛 🎦 🐂 nas.aliyuncs.com:/ /mnt
[root@i:	z8cilini	e.l. i bz /]# net	stat -tanpl   grep tcp	grep LISTEN	
tcp	0	0 0.0.0.0:45450	0.0.0:×	LISTEN	-
tcp	0	0 0.0.0.0:111	0.0.0:×	LISTEN	1/systemd
tcp	0	0 0.0.0:22	0.0.0:*	LISTEN	2159/sshd
tcp6	0	0 :::45450	:::*	LISTEN	-

## Why am I unable to use a Linux client to read data from or write data to files that are named in Chinese characters in a file system?

The files that are named in Chinese characters in a file system are created on a Windows client. The names of these files are in the GBK format. By default, Linux clients can recognize only file names that are in the UTF-8 format. Therefore, Linux clients cannot recognize files that are named in Chinese characters. We recommend that you use a Windows client to read data from and write data to the files that are named in Chinese characters.

## Why is a 523 error returned when I run the 1s command on a Linux client on which an NFS file system is mounted?

Issue

The following error message is returned when you run the ls command on a Linux client on which an NFS file system is mounted:

#### Ls: reading directory '.': Unknown error 523

Cause

If you run the 1s command on a directory of a file system while multiple rename operations are concurrently called, a 523 error occurs.

Solution

Try again later. If the error persists, submit a ticket.

## How do I fix the error that occurs when I mount an SMB file system on Windows?

- 1. System error 53
  - Error message

The network path is not found.

- Cause
  - The network connection fails.
  - The TCP/IP Net BIOS Helper service is not started.
  - LanmanWorkstation is not specified in the registry.
- Solution

- a. Run the **ping** <**Endpoint** of a mount target > command to check whether the endpoint of a mount target is accessible and whether the latency is within the expected range.
  - If you can ping the endpoint of the mount target, go to Step 2.
  - Otherwise, perform the following steps:
    - Make sure that the mount command is valid. For example, check whether the command does not include redundant forward slashes ( / ), backslashes ( \ ), spaces, or myshare.

To run a valid command to mount an SMB file system, use the following syntax:

```
net use <Letter of the destination drive> \\ amount target>\mysh are
```

Example:

```
net use z: \\xxxx.cn-hangzhou.nas.aliyuncs.com\myshare
```

• Make sure that the protocol type of the file system is SMB.

File System ID/Name	Storage Type	Protocol Type	Storage Capacity	Zone	Bound Storage Package	Number of Mount Points		Action
CODHNE2 (seed	Capacity-type	SMB	0 B	China North 1 Zone C	No	1	Add Mount Point   Manage	Delete

- Make sure that the endpoint of the mount target is valid.
- Make sure that the Elastic Compute Service (ECS) instance and the mount target reside in the same virtual private cloud (VPC).
- If the ECS instance and the mount target do not reside in the same VPC, make sure that the VPC or virtual private network (VPN) configurations of the ECS instance are valid.
- b. Run the **telnet** <**Endpoint of a mount target**> **445** command to check whether the SMB protocol is enabled.
- c. Check whether the TCP/IP NetBIOS Helper service is started. For more information, see Mount an SMB file system on Windows.
- d. Open the Registry Editor, choose HKEY\_LOCAL\_MACHINE > System > Current ControlSet > Control > NetworkProvider > Order. In the dialog box that appears, check whether the ProviderOrder key contains the LanmanWorkstation value. If the ProviderOrder key does not contain the LanmanWorkstation value, add the value to the ProviderOrder key.
- 2. System error 58
  - Error message

The specified server cannot perform the requested operation.

Cause

The Windows operating system that the ECS instance runs is incompatible with the SMB protocol that is used by the file system.

• Solution

Make sure that the ECS instance runs Windows Server 2008 R2 or a later version, excluding Windows Server 2008.

- 3. System error 64
  - Error message The specified network name is unavailable.
  - Cause
- The IP address of the ECS instance is not included in the permission groups of the NAS file system.
- The internal IP address or VPC IP address of the ECS instance is not included in the permission groups of the NAS file system.
- Your Alibaba Cloud account has overdue payments.
- The ECS instance and the NAS file system reside in the classic network. However, the ECS instance and the NAS file system belong to different Alibaba Cloud accounts.
- The protocol type of the file system is not SMB.

#### • Solution

This error occurs because the NAS file system is inaccessible. Perform the following steps to fix the error:

- a. Make sure that the internal IP address or VPC IP address of the ECS instance is included in the permission groups of the NAS file system.
- b. Make sure that your Alibaba Cloud account does not have overdue payments.
- c. If the ECS instance and the NAS file system reside in the classic network, make sure that they belong to the same Alibaba Cloud account.
- d. Make sure that the protocol type of the file system is SMB.

File System ID/Name	Storage Type	Protocol Type	Storage Capacity	Zone	Bound Storage Package	Number of Mount Points	Action
an bernaue Is and	Capacity-type	SMB	0 B	China North 1 Zone C	No	1	Add Mount Point   Manage   Delete

#### 4. System error 67

• Error message

The network name cannot be found.

• Cause

The required network services are not started.

• Solution

Start the following services. For more information, see Mount an SMB file system on Windows.

- a. The Workstation service.
- b. The TCP/IP Net BIOS Helper service.
- 5. System error 85
  - Error message

The local device name is already in use.

• Cause

The specified drive letter is already in use.

• Solution

Change the drive letter and remount the file system.

- 6. System error 1231
  - Error message The network location is unavailable.
  - Cause
    - The Client for Microsoft Networks component is uninstalled or disabled.
    - The File and Printer Sharing for Microsoft Networks component is uninstalled or disabled.

• Solution

Install and enable the Client for Microsoft Networks component or the File and Printer Sharing for Microsoft Networks component.

If the Client for Microsoft Networks component or the File and Printer Sharing for Microsoft Networks component is installed but not enabled, select the Client for Microsoft Networks option or the File and Printer Sharing for Microsoft Networks option. Perform the following steps to install and enable the component:

- a. On the Network and Sharing Center page, click the active network connection.
- b. Click **Properties**.
- c. In the WLAN Properties dialog box, click Install.
  - Install the Client for Microsoft Networks component.
    - a. In the Select Network Feature Type dialog box, select Client, and click Add.
    - b. Select Client for Microsoft Networks and click OK.
  - Install the File and Printer Sharing for Microsoft Networks component.
    - a. In the Select Network Feature Type dialog box, select Service, and click Add.
    - b. Choose Microsoft > File and Printer Sharing for Microsoft and click OK.
- 7. System error 1272
  - Error message

You cannot access this shared folder because the security policies of your organization block access from unauthorized guests. These policies protect your PC from unsafe or malicious devices on the network.

• Cause

The security policies of the Windows operating system block access from guest users to the SMB file system.

• Solution

If the ECS instance runs a Windows operating system that is later than Windows Server 2016, excluding Windows Server 2016, configure the following registry to allow access from guest users:

[HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters] "AllowInsecureGuestAuth"=dword:1

For more information, see Guest access in SMB2 disabled by default in Windows.

#### Why does a file in a NAS file system belong to different owners when I query the file on two ECS instances?

In NAS file systems, users are identified by User Identifiers (UIDs) or Group Identifiers (GIDs) instead of user names. The owner name that you query on the ECS instance is converted from a UID. If a UID is converted into different user names on different ECS instances, the UID is identified as a different owner on each ECS instance.

For example, create a file named *admin\_on\_machine1* on ECS Instance 1 and a file named *admin\_on\_machine2* on ECS Instance 2 as the admin user. Run the ls -l command on ECS Instance 1 to view the created file, as shown in the following figure.

\$11								
total O								
-rw-rw-r	1	admin	admin	0	Apr	6	17:10	admin_on_machine1
-rw-rw-r	1	terminal	19062	0	Apr	6	17:12	admin_on_machine2

Run the ls -l command on ECS Instance 2 to view the created file, as shown in the following figure.

\$11								
total 0								
-rw-rw-r	1	505	505	0	Apr	6	17:10	admin_on_machine1
-rw-rw-r	1	admin	admin	0	Apr	6	17:12	admin_on_machine2

The query results on the two ECS instances indicate that the same file has different owner names.

Run the **id** command on the two ECS instances to query the admin user. The UID of the admin user on ECS Instance 1 is 505, as shown in the following figure.

#### \$id uid=505(admin) gid=505(admin) groups=505(admin)

The UID of the admin user on ECS Instance 2 is 2915, as shown in the following figure.

#### \$id uid=2915(admin) gid=19062(admin) groups=19062(admin)

Run the stat admin\_on\_machine1 admin\_on\_machine2 command, as shown in the following figure. The results indicate that the two files belong to two different UIDs.



Why am I unable to mount an SMB file system on a Windows operating system that is later than Windows Server 2016?

• Issue

#### An error occurs when you run the following command:

C:\Users\Administrator>net use z: \\xxxxx.nas.aliyuncs.com\myshare System error 1272 has occurred.

You can't access this shared folder because your organization's security policies block u nauthenticated guest access. These policies help protect your PC from unsafe or malicious devices on the network.

#### Solution

This error occurs because the security policies of a Windows operating systems that are later than Windows Server 2016 do not allow guest users to access remote shared directories. Perform the following steps to fix the error:

#### • Locate the following registry key:

[HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters] "AllowInsecureGuestAuth"=dword:0

#### Modify the key.

[HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters] "AllowInsecureGuestAuth"=dword:1

• Open PowerShell and run the following command:

New-ItemProperty -Path \$registryPath -Name \$name -Value \$value -PropertyType DWORD -For ce

For more information, see Guest access in SMB2 disabled by default in Windows 10, Windows Server 2016 version 1709, and Windows Server 2019.

#### Why am I unable to mount an SMB file system?

Issue

The **net use** command is used to mount SMB file systems. If you accidentally use the command to mount NFS file systems, you can no longer use the command to mount an SMB file system.

Solution

Make sure that the protocol of the file system is SMB. Then, stop the mount operation and try to mount the file system again after 5 minutes. If the issue persists, submit a ticket.

#### Why is a mounted SMB directory visible only to an administrator?

This issue occurs because Windows user accounts are isolated from each other. For example, if you log on to Windows as User A, you cannot view the directory that you mounted as User B.

To enable access from multiple users, create a shared directory. For example, you can run the following command to create a shared directory named myshare on drive C:

mklink /D C:\myshare \\xxxxxx.cn-beijing.nas.aliyuncs.com\myshare\

# Why is Internet Information Services (IIS) unable to load the files of an SMB volume on Windows Server 2016?

For more information about how to fix this issue, see Install and configure Active Directory domains.

# How do I fix the error that occurs when I use IIS to access a NAS file system?

When you mount an SMB file system on Windows Server 2016, an HTTP error500.19 (error code0x8007003a) occurs. For information about how to fix this error, see Best practices for using IIS toaccess a NAS file system.

### Can I mount an NFS file system and SMB file system on the same ECS instance?

No, you cannot mount an NFS file system and SMB file system on the same ECS instance.

To prevent compatibility issues, we recommend that you do not access an SMB file system by using a Linux client. For example, the supported character sets and the length of a file name for Windows and Linux are different. In Windows, a maximum length of 255 Unicode wide characters is supported. In Linux, a maximum length of 255 UTF-8 characters is supported.

To mount an NFS file system and an SMB file system on the same ECS instance, you can mount the SMB file system on a Linux ECS instance. The kernel of the Linux ECS instance must support SMBv2 or a later version.

Run one of the following mount commands: mount -t cifs -o vers=2.0 \\<Mount target>\myshare /mnt Or mount -t cifs -o vers=2.0 //<Mount target>/myshare /mnt .

**Note** If you are prompted to enter a password after you run the commands, press Enter.

To check whether your Linux kernel supports the CIFS protocol, view the value of CONFIG\_CIFS in the */boot* directory. A value of y or m indicates that the protocol is supported and a CIFS file system can be mounted.

#### \$grep -i cifs /boot/config-2.6.18-274.alios5.1 CONFIG\_CIFS=m

#### ? Note

• Before you run the preceding command, you must install the cifs-utils tool. For example, you can run the following command to install the cifs-utils tool on CentOS:

yum install samba-client samba-common cifs-utils

• If the CIFS protocol is not supported on your version of Linux, we recommend that you upgrade the Linux kernel to version 3.10.0-514 or later.

#### Why am I unable to mount an SMB file system on Linux?

Cause

- You are using an early or incompatible version of Linux distributions. SMB file systems support the following Linux distributions:
  - Cent OS 7.6 64-bit (3.10.0-957.5.1.el7.x86\_64)
  - Ubunt u 18.04 64-bit (4.15.0-48-generic)
  - Debian 9.9 64-bit (4.9.0-9-amd64)
  - SUSE Enterprise Server 12 SP2 64-bit (4.4.74-92.35-default)
  - o openSUSE 42.3 64-bit (4.4.90-28-default)

- Alibaba Cloud Linux (4.19.34-11.al7.x86\_64)
- CoreOS (4.19.43-coreos VersionID=2079.4.0)
- The cifs-utils tool is not installed on the client, or the executable file of the mount.cifs command is stored in a directory that is different from the directory specified by the PATH environment variable.
- No network connection is established between the Linux ECS instance and the SMB file system.
  - The Linux ECS instance and the SMB file system belong to different Alibaba Cloud accounts.
  - The Linux ECS instance and the SMB file system reside in different regions.
  - The Linux ECS instance and the SMB file system reside in different networks. For example, the Linux ECS instance and the SMB file system reside in different VPCs or one of them resides in a VPC and the other resides in the classic network.

(?) Note You can mount a NAS file system on an on-premises Linux client. If you cannot access the file system from the Linux client, a probable cause is that no network connection is established between the Linux client and the file system. To establish a network connection, use Express Connect.

- The IP address of the Linux ECS instance is not included in the permission groups of the SMB file system.
- The firewall of the Linux ECS instance denies access to the IP address or port 445 of the SMB file system.
- The Linux ECS instance attempts to connect with the SMB file system by using an unsupported Transmission Control Protocol (TCP) port. SMB file systems support only port 445.

#### ? Note

You can run the ping <VolumeDomainName> and telnet <VolumeDomainName>445 commands to check the network connectivity.

If port 445 is disabled, you must add rules for port 445 to a security group of the ECS instance. For more information, see Add security group rules.

- The account that you use to log on to the Linux ECS instance does not have the root permissions or is not authorized to run the **mount** command. You can run the sudo command to authorize the account to run the mount command.
- The file system type is not set to Common Internet File System (CIFS).
- The value of the vers parameter in the mount command is not 2.0.
- The identity of the guest user is not specified when the NAS file system is mounted.
- The specified value for uid, gid, dir\_mode, or file\_mode is invalid.
- The Security-Enhanced Linux (SELinux) settings for the mount directory are invalid.
- The file system is mounted on more than 1,000 ECS instances that run Linux. This issue often occurs when you mount a file system on containers.

#### Solution

- 1. Fix the issue based on the preceding causes. For more information, see Mount an SMB file system on a Linux ECS instance.
- 2. Fix the issue based on the */var/log/messages* file and the output of the dmesg command.
- 3. Submit a ticket to contact Alibaba Cloud for technical support.

When you submit a ticket, you must provide the version of your Linux distribution, mount commands, */var/log/messages* file, and output of the dmesg command.

# How do I improve the performance of an SMB file system that is mounted on Linux?

If the performance of your SMB file system cannot meet your requirements, you can use the following solutions based on specific causes:

- Cause 1: The maximum read/write throughput of the SMB file system has a linear relationship with the capacity of the file system. Solution: Use the fio tool to test the performance of the SMB file system. For more information, see Performance testing for Apsara File Storage NAS.
- Cause 2: The bandwidth of the Linux ECS instance is low. Solution: Use multiple Linux ECS instances to ensure that the file system can provide expected performance.
- Cause 3: Caching is disabled for the SMB client.
   Solution: If the cache parameter is set to strict, caching is enabled. If the cache parameter is set to none, caching is disabled. By default, caching for an SMB client is enabled. You can run the sudo mount | grep cifs command to check the value of the cache parameter.
- Cause 4: The I/O size of the SMB client does not meet your business requirements. Solution: Specify the rsize and wsize parameters based on your business requirements. The default value of the two parameters is 1048576.
- Cause 5: The Linux ECS instance uses low-specification CPU or memory, or most CPU or memory resources are occupied by other processes.
   Solution: Specify the required specifications for the Linux ECS instance based on the usage of CPU and memory resources. This ensures that the file system can function as expected. You can run the top command to check the usage of CPU and memory resources.
- Cause 6: The atime parameter is specified when you mount the file system. Solution: Do not specify the atime parameter if your business does not require fast file access.
- Cause 7: The web server such as Apache HTTP Server on the Linux ECS instance processes a few write requests. These requests require notifications and frequent read operations on a large number of small files.

Solution: Configure the caching mechanism of the web server on the Linux ECS instance. You can also contact Alibaba Cloud to enable the acceleration feature for the web server.

# Why does file migration and replication take a long time when I mount an SMB file system on Linux?

Check whether the file system provides poor performance. If the performance of the file system can meet your requirements, a probable cause is that the files are not concurrently migrated or replicated. You can use the following open source tools to migrate or replicate files.

• GNU Parallel

```
Specify a number of threads based on your system resources. Example: find * -type | parallel -
-will-cite -j 10 cp {} /mnt/smb/ &
```

- Fpart
- Fpsync
- multi

### How do I fix the Permission denied error when I access an SMB file system on Linux?

Cause: You specified an invalid value for the uid, gid, file\_mode, or dir\_mode parameter when you mounted the file system.

Solution: Check whether the values that are specified for the uid, gid, file\_mode, and dir\_mode parameters are valid. For more information, see Mount an SMB file system on a Linux ECS instance.

### How do I rename the files of an SMB file system by changing the letter case?

The file names of an SMB file system are case-insensitive. This also applies to Windows systems. A file in an SMB file system cannot be renamed by changing only the letter case.

However, you can change a file name to a different name that includes different letters. Then, you can change the file name to the original name with a different letter case.

### Why am I unable to change the owner of a file and access mode of a file or directory?

You can specify the owner of a file and access mode of a file or directory in a file system only when you mount the file system. For more information, see Mount an SMB file system on a Linux ECS instance.

# Why does the server not respond within 35 seconds when multiple clients concurrently access a file?

Cause: The kernel driver of the current SMB protocol fails to work as expected. If the SMB protocol version is 2.1 or 3.0, the server does not respond within 35 seconds. In this case, the clients cannot send SMB break acknowledgment packets to the server.

Solution 1: Set the vers parameter to 2.0 if you mount the file system on the Linux ECS instance.

Solution 2: Perform the following operations:

- 1. If the CIFS module is being loaded, run the following command to disable the oplock feature:
   # modprobe cifs enable\_oplocks=0
- 2. If the CIFS module is already loaded, run the following command to disable the oplock feature: # echo 0 > /sys/module/cifs/parameters/enable\_oplocks
- 3. Run the following command to check the status of the oplock feature:
   # cat /sys/module/cifs/parameters/enable\_oplocks

In the result, Y indicates that the feature is enabled. N indicates that the feature is disabled.

? Note

- To apply the preceding changes, unmount and remount the SMB file system.
- To permanently apply the preceding changes, create the /etc/modprobe.d/cifs.conf file and add the options cifs enable oplocks=0 statement to the file.

# Why does the mount target of an SMB file system become unresponsive?

Cause: If the kernel version of your Linux distribution is 3.10.0-514 or earlier, the kernel driver of the SMB protocol may fail to respond when multiple clients concurrently access the file system. In this case, the mount target is inaccessible. The following record is included in the kernel log:

```
...
[<fffffffc03c9bc1>] cifs_oplock_break+0x1f1/0x270 [cifs]
[<fffffff810a881a>] process_one_work+0x17a/0x440
[<fffffff810a8d74>] rescuer_thread+0x294/0x3c0
...
```

Solution

- Set the cache parameter to none to remount the file system. This may affect the performance of the file system.
- Upgrade the operating system of the Linux ECS instance.

### How do I fix the error that occurs when I mount an NFS file system on Windows?

- Error message: Invalid file handler Solution: Perform the required steps and set the parameters to remount the file system. For more information, see Mount an NFS file system on a Windows ECS instance.
- Error message: Network error 53 Solution: Perform required steps and set the parameters to remount the file system. For more information, see Mounting NFS on a Windows Client.
- Error message: Network error 1222 Solution: After you install an NFS client, remount the file system. For more information, see Install the NFS client.

# How do I fix the errors that occur when I soft mount an NFS file system on Windows?

• Issue

By default, the soft mode is used when you mount an NFS file system on Windows. However, in some scenarios, soft mounts result in data inconsistency or unexpected application exits.

- Data inconsistency: An application sends an ECS instance a request to write data to a softmounted file system. If the application does not receive a response from the instance before the request times out, an error is returned. This applies even if data is written to the file system. In this case, the application determines that the request failed. However, the ECS instance determines that the request is successful. As a result, data inconsistency occurs.
- Unexpected application exits: An application sends an ECS instance a request to access a softmounted file system. If the application does not receive a response from the instance before the request times out, an error is returned. In this case, an exception may be thrown based on the programming language that is used to write the application. If the exception is not handled, the application exits.
- Solution 2

To prevent these issues, hard mount the NFS file system on the ECS Windows instance.

- i. Run the mount command to view the mount mode.
  - If the command output includes mount=soft , perform the following steps.

 If the command output includes mount=hard, you do not need to perform the following steps.

Local	Remote	Properties
X:		<pre>UID=-2, GID=-2 rsize=32768, wsize=32768 mount=soft, timeout=1.6 retry=1, locking=yes fileaccess=755, lang=ANSI casesensitive=no sec=sys</pre>

- ii. Stop the application that is using the NFS file system.
- iii. Run the following command to unmount the NFS file system:

umount H:

Replace the drive letter H: based on your business requirements.

iv. Run the following command to remount the NFS file system:

```
mount -o nolock -o mtype=hard -o timeout=60 \\xxxxx.cn-hangzhou.nas.aliyuncs.com\! h
:
```

Replace the mount target address xxxxx.cn-hangzhou.nas.aliyuncs.com and the drive letter h: based on your business requirements.

v. Run the mount command to verify the mount result:

The mount is successful if the command output includes mount=hard, locking=no, and timeout=10 or a number greater than 10.

C: Wser	s\Administrator>mount	
本地	远程	属性
n :	\\test-aan14c92-oyr46.cs-shangha≦_na	us~ UID=0, GID=0 rsize=1048576, wsize=1048576 mount=hard, timeout=10.0 retry=1, locking=no fileaccess=755, lang=GB2312-80 casesensitive=no sec=sys

How do I resolve the invalid device error that is returned when I try to rename a file on the Windows NFS client?

If you mount the NFS file system on a subdirectory of the ECS instance, the invalid device error is returned when you rename the file. To prevent this error, mount the file system on the root directory of the ECS instance.

# If the system fails to terminate processes that connect to a file system, how do I clear the handles that are exposed by a client?

To release all handles, you can use the following tool to remove all connections from an SMB file system.

Windows client

Use the tcpview tool to remove all connections from an SMB file system. For more information, see tcpview.

File Edit View Proces	s Connection (	Options Hel	р						
C 🕐 🖬 🗞	4 TCP v4 6	TCP v6	UDP v4	5 UDP v6 🏴	P Search				
Process Name	Process ID	Protocol	State	Local Address	Local Port	Remote Addre	ess	Remote Port	
System	4	TCP	Established	11. 9.60	50250	11	持要占回	445	б/
System	4	TCP	Established	11 9.60	61413	11	Process	Properties	6/
System	4	TCP	Established	11 9.60	50101	11	Kill Proc	ess	6/
CamberWin.exe	2584	TCP	Established	12	50877	127			1/16
🖆 java.exe	3800	TCP	Established	12	50860	127	Close Co	onnection	1/16
CamberWin.exe	2584	TCP	Established	12	50858	127	Whois		1/16
CcSvcHst.exe	3024	TCP	Established	11 9.60	50269	30	Сору	Ctrl+C	6/

• Linux client

Use the killcx tool to remove all connections from an SMB file system. For more information, see killcx.

# How do I prevent exceptions if multiple processes or clients concurrently write data to a log file?

Issue

NAS allows multiple clients to write data to different files in the same namespace over NFS. However, NFS does not support atomic appends. Some exceptions may occur if multiple processes or clients concurrently write data to the same file. This is because each process independently maintains context information. The file can be a log file. The context information includes file descriptors and write locations. These exceptions include overwrite, crossover, and disordered content.

- Solution
  - Recommended. Use different processes or clients to write data to different files in the same file system. When you analyze or process data, you can consolidate these files. This solution can fix the issues that are caused by concurrent write operations without the need to use file locks. In addition, the solution does not affect system performance.
  - Use a combination of the flock and seek functions. This ensures the atomicity and consistency of write operations. However, this solution requires a long period of time and may significantly affect system performance. The following steps describe this method.
- Use the flock and seek functions together

NFS does not support atomic appends. If multiple clients append data to the same file such as a log, data entries may overwrite each other. In Linux, you can use the flock and seek functions together to simulate atomic appends on an NFS file system. This ensures data consistency when multiple processes concurrently append data to the same file.

To use the flock and seek functions together, perform the following steps:

- i. Execute the fd=open(filename, O\_WRONLY | O\_APPEND | O\_DIRECT) statement to open a file by using the append method. This statement is used to set the write method to O\_DIRECT and obtain the file descriptor. O\_DIRECT specifies a write-only method. In this case, no page cache is used.
- ii. Call the flock(fd, LOCK\_EX|LOCK\_NB) function to obtain a file lock. If the function fails to obtain a file lock, an error message is returned. The failure may occur if the file lock is being used. You can try again or troubleshoot the failure.
- iii. After the file lock is obtained, call the lseek(fd, 0, SEEK\_END) function to set the current file offset of the file descriptor to the end of the file.
- iv. Write data to the file. The insert position is located at the end of the file. The file lock is used to prevent data entries from overwriting each other.

v. After data is written to the file, call the flock(fd, LOCK\_UN) function to release the file lock.

The following code shows a sample program that is written in C.

```
#define GNU SOURCE
#include<stdlib.h>
#include<stdio.h>
#include<fcntl.h>
#include<string.h>
#include<unistd.h>
#include<sys/file.h>
#include<time.h>
const char *OUTPUT FILE = "/mnt/blog";
int WRITE_COUNT = 50000;
int do lock(int fd)
{
   int ret = -1;
   while (1)
    {
       ret = flock(fd, LOCK_EX | LOCK_NB);
       if (ret == 0)
        {
            break;
        }
       usleep((rand() % 10) * 1000);
    }
   return ret;
}
int do unlock(int fd)
{
   return flock(fd, LOCK_UN);
}
int main()
{
        int fd = open(OUTPUT FILE, O WRONLY | O APPEND | O DIRECT);
        if (fd < 0)
        {
                printf("Error Open\n");
                exit(-1);
        }
        for (int i = 0; i < WRITE COUNT; ++i)</pre>
        {
                char *buf = "one line\n";
                /* Lock file */
                int ret = do lock(fd);
                if (ret != 0)
                {
                        printf("Lock Error\n");
                        exit(-1);
                }
                /* Seek to the end */
                ret = lseek(fd, 0, SEEK_END);
                if (ret < 0)
                {
                      printf("Seek Error\n");
```

```
exit(-1);
                 }
                 /* Write to file */
                int n = write(fd, buf, strlen(buf));
                 if (n <= 0)
                 {
                         printf("Write Error\n");
                         exit(-1);
                 }
                 /* Unlock file */
                 ret = do unlock(fd);
                 if (ret != 0)
                 {
                         printf("UnLock Error\n");
                         exit(-1);
                 }
        }
        return 0;
}
```

For more information about how to call the flock() function, see filck().

(?) Note Only Linux kernel versions 2.6.12 and later support the flock() function. If you use Linux of an earlier kernel version, call the fcntl() function.

# Why do I need to mount a NAS file system by using the noresvport parameter?

If network switchovers or high availability (HA) switchovers of backend services occur in a NAS file system, the network connection of the file system may be interrupted. If the network connection is affected, you may need to wait for several minutes for the connection to automatically recover. If the connection fails to automatically recover, you must restart the ECS instance. If you specify the noresvport parameter, the connection can automatically recover within several seconds.

#### How do I fix issues related to the noresvport parameter?

**Note** This question applies only to Linux users. If you are a Windows user or you mount NAS file systems by using CSI or Flexvolume plug-ins on Container Service for Kubernetes (ACK) clusters, you can skip this question.

1. Check whether the noresvport parameter is specified when you mount a file system.

i. Run the following command on a Linux ECS instance to download the check\_noresvport.py script:

wget -N https://code.aliyun.com/nas\_team/nas-client-tools/raw/master/linux\_client/c heck noresvport.py -P /tmp/

ii. Run the following Python command to execute the script:

python2.7 /tmp/check\_noresvport.py -e

If the "There is no issue for 'noresvport' on this ECS" message appears, skip the following steps.

2. Fix issues related to the noresvport parameter.

**?** Note We recommend that you fix these issues during off-peak hours.

Use one of the following mount scenario-specific solutions to fix the issues.

• If you mount a NAS file system on an ECS instance, run the following command to execute the script again:

```
python2.7 /tmp/check_noresvport.py -e -r
```

• If you mount a NAS file system on a container, run the following command to execute the script again on the node where the container resides:

```
python2.7 /tmp/check_noresvport.py -e -c
```

- 3. Update the settings for an automatic mount.
  - If you configured an automatic mount, add the noresvport parameter to the settings for the automatic mount. For more information, see Automatically mount the NFS file system.
  - If no automatic mount is configured, skip this step.

After the preceding steps are completed, repeat Step 1 and verify the result. If the issue persists, submit a ticket to contact Alibaba Cloud for technical support.

#### What happens when a network switchover or an HA switchover occurs on backend services?

NAS provides stable and continuous file storage services. However, network switchovers or HA switchovers of backend services may still occur in rare cases. HA switchovers of backend services may occur due to service upgrades. These switchovers can interrupt the client network. Before service upgrades, Alibaba Cloud sends notifications that include information about the upgrade schedule to all related users. This ensures that you have sufficient time to set the noresvport parameter. We recommend that you set the noresvport parameter at your earliest opportunity even if no upgrade is scheduled at the backend. This prevents failures that may occur on file systems due to other unexpected issues. These issues include changes on Server Load Balancer (SLB) instances, hardware failures at the backend, and other conditions that may trigger switchovers.

# Why do I need to remount a file system? Can I use an alternative solution?

Before you can use the noresyport parameter to mount a file system, you must remount the file system. This way, you can end all TCP connections to which the noresyport parameter is not applied. When you use the noresyport parameter to mount the file system, new TCP connections are established. To end all previous TCP connections, you must stop all services that use the NAS file system. Then, you can run the **umount** command to unmount the file system.

If you do not want to remount the file system, we recommend that you create a mount target based on a new file system. In this case, you must mount the new file system on a different local directory. You can migrate all services to the new local directory. Then, you can disable the old mount directory and mount target.

# How are files prefixed with .nfs generated? How do I delete files prefixed with .nfs?

If you delete a file when the file is being used by an application, a temporary file prefixed with .nfs is generated. When the process that uses the file ends, the temporary file is automatically deleted.

#### When I access a file in the directory of a NAS file system, the bind conn to session failed on NFSv4 server error message is returned. How do I fix this issue?

• Cause

The error message is returned because you mounted the file system by using the NFSv4.1 protocol. NAS does not support this protocol.

• Solution

Use the NFSv3.0 or NFSv4.0 protocol to remount the file system based on your business requirements. For more information, see Usage notes.