

# Alibaba Cloud Apsara File Storage NAS

FAQ

Issue: 20190617

## Legal disclaimer

---

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

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

or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, 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 the Alibaba Cloud website, product programs, or content shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates).
6. Please contact Alibaba Cloud directly if you discover any errors in this document.



## Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 <b>Danger:</b> Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 <b>Warning:</b> Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 <b>Notice:</b> Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 <b>Note:</b> You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
<b>Bold</b>	It is used for buttons, menus, page names, and other UI elements.	Click OK.
<code>Courier font</code>	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[ ] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	It indicates that it is a required value, and only one item can be selected.	<code>swich {stand   slave}</code>



# Contents

---

Legal disclaimer.....	I
Generic conventions.....	I
1 Known issues on an NFS client.....	1
2 FAQs.....	3
2.1 Regions and supported storage types and protocols.....	3
2.2 Why do I need RAM permissions to create a mount point in a classic network.....	6
2.3 How is NAS charged.....	6
2.4 What is NAS and where can I use it.....	7
2.5 Is there a NAS terminology.....	7
2.6 How many file systems can an account create.....	8
2.7 What protocols does NAS support.....	8
2.8 Does capacity affect performance in NAS.....	8
2.9 What is a mount point and where can I use it.....	8
2.10 What is a permission group and where can I use it.....	9
2.11 How to prevent NFS 4.0 from being mistaken for a Trojan attack.....	9
2.12 How do I avoid exceptions when multiple processes and multiple clients write the same log file at the same time.....	10
2.13 What impacts the I/O performance of Windows service SMB protocol.....	12
2.14 How to improve performance when using IIS to access NAS.....	13
2.15 Does NAS support inotify.....	14
2.16 Why does my NFS run in Linux at MB speed.....	15
2.17 How do I use the noac option.....	15
3 SMB FAQ.....	17
3.1 SMB basic operation FAQ.....	17
3.2 SMB mount issue FAQ.....	18
3.3 Performance issues of Server Message Block.....	21
3.4 Causes and solutions of SMB mount failures.....	22



# 1 Known issues on an NFS client

As an NFS client is a part of a kernel. Therefore, errors may occur when you use an NFS client due to a number of existing defects in a kernel. To improve the user experience of an NFS client, we recommend that you use kernels with latest versions. The following kernels are not recommended.

## No response from an NFS file system due to an abnormal TCP connection

If either one of the following conditions is met, an NFS file system may be inaccessible to end-users. This is caused by an abnormal TCP connection in situations, such as unstable network connections and server failovers.

- The kernel version is earlier than 3.10.0-693.2.2.
- The kernel version is 3.10.0-693.2.2 or later, but the `noresvport` option is not enabled for NFS mounting.



### Note:

You can run the `uname -a` command to check the version of a kernel.

If a connection is abnormal, you need to re-mount a file system on an ECS instance. For more information about how to mount a file system, see [Mount an NFS file system in Linux](#). You can use the recommended parameters to mount a file system.

After the `noresvport` option is enabled when you mount a file system, you must run the `ss -ant | grep 2049` command to view the number of the port used for mounting a file system as shown in the following figure. If the port number is greater than 1024, it indicates that the file system is mounted. Otherwise, you must unmount all file systems. You can run the `mount | grep nas` command to ensure no file system is mounted. Then, you can re-mount a file system.

```
[root@ ~]# ss -ant | grep 2049
ESTAB      0      0      192.168.0.163:30024      192.168.0.168:2049
```



### Notice:

We recommend that you do not use kernel versions that range from 2.6.32-696 to 2.6.32-696.10.1. The issue of no system response is more prone to occur on one of

the kernel versions within this range. For more information, see [RHEL6.9: NFSv4 TCP transport stuck in FIN\\_WAIT\\_2 forever](#).

#### No response from an NFS file system due to kernel defects

When the kernel version that you are using is one of the early versions, an NFS file system may have no response. When you perform operations, such as open, read, and write files on an NFS client, a deadlock issue occurs due to failover at the NFS backend. This results in no response being returned from an NFS file system.

- Redhat 6/CentOS 6 2.6.32-696.3.1.el6
- Redhat 7/CentOS 7 3.10.0-229.11.1.el7 and earlier versions
- Ubuntu 15.10 Linux 4.2.0-18-generic

When no response is returned from an operation, we recommend that you restart an ECS instance. For more information, see [RHEL7: NFSv4 client loops with WRITE / NFS4ERR\\_STALE\\_STATEID - if NFS server restarts multiple times within the grace period](#).

#### No support for the chown command and API operations

When the kernel version is 2.6.32, the chown command and system APIs are not supported on an NFS client.

#### Unresponsive ls operation

When the kernel version is 2.6.32-696.1.1.el6 or earlier, you cannot stop an ls operation. This occurs when you perform operations, such as adding and removing files or folders.

We recommend that you can upgrade the kernel version to fix the issue.

## 2 FAQs

---

### 2.1 Regions and supported storage types and protocols

The functions that NAS supports vary depending on regions.

A region refers to a physical data center. After you create a file system, you can no longer change the region of the file system.

The following table shows the storage types, protocols, and zones supported in different regions.

- Mainland China:

Region	City	Supported storage types, protocols, and zones
China North 1	Qingdao	Capacity: NFS, SMB Zones: C
China North 2	Beijing	- Performance: NFS, SMB Zones: C, E - Extreme Zones: A, B, C, D, E, F
China North 3	Zhangjiakou	- Performance: NFS, SMB Zones: B - Extreme Zones: A, B
China North 5	Hohhot	- Capacity: NFS, SMB Zones: A - Performance: NFS, SMB Zones: A

Region	City	Supported storage types, protocols, and zones
China East 1	Hangzhou	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB Zones: B, G</li> <li>- Performance: NFS, SMB Zones: G, F</li> <li>- Extreme Zones: B, E, F, H, I</li> </ul>
China East 2	Shanghai	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB Zones: B</li> <li>- Performance: NFS Zones: B</li> <li>- Extreme Zones: A, B, C, D, E, F</li> </ul>
China South 1	Shenzhen	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB Zones: A, D</li> <li>- Performance: NFS, SMB Zones: A, B, D</li> <li>- Extreme Zones: A, B, C, D, E</li> </ul>

· Other regions:

Region	City	Supported storage types, protocols, and zones
Hong Kong	Hong Kong	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB Zones: B</li> <li>- Performance: NFS, SMB Zones: B</li> </ul>

Region	City	Supported storage types, protocols, and zones
Asia Pacific SE 1	Singapore	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB</li> <li>Zones: A</li> <li>- Performance: NFS, SMB</li> <li>Zones: A</li> </ul>
Asia Pacific SE 2	Sydney	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB</li> <li>Zones: A</li> <li>- Performance: NFS, SMB</li> <li>Zones: A</li> </ul>
Asia Pacific SE 3	Kuala Lumpur	Capacity: NFS, SMB Zones: A
Asia Pacific SE 5	Jakarta	Capacity: NFS, SMB Zones: A
Asia Pacific NE 1	Tokyo	Capacity: NFS, SMB Zones: A, B
Asia Pacific SOU 1	Mumbai	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB</li> <li>Zones: A</li> <li>- Performance: NFS, SMB</li> <li>Zones: A</li> </ul>
EU Central 1	Frankfurt	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB</li> <li>Zones: A</li> <li>- Performance: NFS, SMB</li> <li>Zones: A</li> </ul>
UK	London	<ul style="list-style-type: none"> <li>- Capacity: NFS, SMB</li> <li>Zones: A</li> <li>- Performance: NFS, SMB</li> <li>Zones: A</li> </ul>

Region	City	Supported storage types, protocols, and zones
US West 1	Silicon Valley	<ul style="list-style-type: none"><li>- Capacity: NFS, SMB</li><li>Zones: B</li><li>- Performance: NFS, SMB</li><li>Zones: B</li></ul>
US East 1	Virginia	<ul style="list-style-type: none"><li>- Capacity: NFS, SMB</li><li>Zones: A</li><li>- Performance: NFS, SMB</li><li>Zones: A</li></ul>

## 2.2 Why do I need RAM permissions to create a mount point in a classic network

Unlike Virtual Private Cloud (VPC) environments, classic network environments are not isolated at the network layer. To ensure data security of your NAS file system, NAS must be authorized through RAM to list your ECS instances. This makes sure that only your own ECS instances can mount or access the NAS file system. Note that the NAS file system and the ECS instance must be under the same Alibaba Cloud account.



### Note:

- NAS is only granted permission to call your `DescribeInstances` interface and has no permission to call any other instances. ECS instances acquired by NAS through the `DescribeInstances` interface are only used for permission verification, and are not recorded in any form.
- Do not delete or edit `AliyunNASDefaultRole` in RAM as it may cause an operation exception error or failure when mounting the file system.

## 2.3 How is NAS charged

NAS storage can be charged as a Pay-As-You-Go service, or subscription packages.

- If you select Pay-As-You-Go, fees are based on the actual usage.

- If you select a subscription package, you can assign the package to a file system to cover usage. Any resource usage that exceeds the purchased capacity is charged at the Pay-As-You-Go billable level. You can upgrade or renew a subscription package during its validity period, however, you cannot downgrade a subscription package.

For more charging information about NAS, see [Billing method](#).

## 2.4 What is NAS and where can I use it

Alibaba Cloud Network Attached Storage (NAS) is a highly reliable, highly available file storage service featuring a distributed file system with unlimited capacity and performance scaling, with namespace and multiple client access support.

NAS supports standard file access protocols, so existing applications do not need to be modified. Furthermore, NAS supports multiple computing nodes (such as ECS instances, E-HPC, and Container Service) simultaneously reading or writing data to the file system.

Scenarios that benefit from Alibaba Cloud ECS instances using NAS include:

- Deploying services using Server Load Balancer and multiple ECS servers (such as web servers) for scenarios where multiple ECS servers need to access the same bucket to share data.
- Sharing logs for when apps on multiple ECS servers need to write logs to the same bucket to facilitate concentrated log data processing and analysis.
- Sharing files where NAS enables you to store your enterprises public files in a centralized manner and share data to multiple business groups, all while maintaining data security.

## 2.5 Is there a NAS terminology

Alibaba Cloud NAS (Network Attached Storage) provides an infinitely scalable file system to store data for ECS servers, and primarily involves the following concepts:

- **File system:** The file system is a NAS instance. You can mount the file system on an ECS server, E-HPC, or Container Service, and then use it like a local file system.
- **Mount point:** A mount point is the entry through which a computing node accesses NAS. It defines what type of network computing node can access NAS, and what permissions are required to access NAS.

- **Permission group:** A permission group defines NAS access permissions, including authorized IP addresses, read/write permissions, and user permissions.

## 2.6 How many file systems can an account create

Each Alibaba Cloud account can create up to 10 file systems. The maximum storage capacity of each file system is 1 PB (performance type) or 10 PB (capacity type).

## 2.7 What protocols does NAS support

NAS supports the following protocols:

1. NFS V3.0 and NFS v4.0.
2. SMB 2.1 and later versions, with corresponding support for Windows 7, Windows Server 2008 R2 and all later versions of Windows, but does not support Windows Vista, Windows Server 2008 and earlier versions.

## 2.8 Does capacity affect performance in NAS

Yes. The maximum throughput performance of an individual file system is linearly related to the storage space of the file system. The higher the storage capacity, the higher the throughput performance.

For more information on the detailed relationship between storage and performance, see the article "Relationship between Throughput Capacity and Storage Space" on the [Pricing Page](#).

## 2.9 What is a mount point and where can I use it

A mount point is the interface for computing nodes (such as ECS instance, E-HPC, or Container Service) to access a NAS file system.

Mount points define the network type of the computing nodes and the permissions required to access NAS.

One mount point can be simultaneously mounted by multiple computing nodes, enabling shared access.



## 2.10 What is a permission group and where can I use it

A permission group defines NAS access permissions, including authorized IP addresses, read/write permissions, and user permissions.

## 2.11 How to prevent NFS 4.0 from being mistaken for a Trojan attack

## Symptom

After NFS 4.0 is mounted to a NAS, NFS 4.0 listens to a random 0.0.0.0 port. Netstat is unable to identify the process that listens to this port.

**The changing listened port and the unidentified listening program make NFS 4.0 may be mistaken for a Trojan attack.**

```
[root@iz8v-1z0z /]# netstat -tanpl | grep tcp | grep LISTEN
tcp      0      0 0.0.0.0:111          0.0.0.0:*          LISTEN   1/systemd
tcp      0      0 0.0.0.0:22          0.0.0.0:*          LISTEN   2159/sshd
[root@iz8v-1z0z /]# mount -t nfs -o vers=4.0 9150-1z0z-1z0z-1z0z.nas.aliyuncs.com:/ /mnt
[root@iz8v-1z0z /]# netstat -tanpl | grep tcp | grep LISTEN
tcp      0      0 0.0.0.0:38862       0.0.0.0:*          LISTEN   -
tcp      0      0 0.0.0.0:111        0.0.0.0:*          LISTEN   1/systemd
tcp      0      0 0.0.0.0:22        0.0.0.0:*          LISTEN   2159/sshd
tcp6     0      0 :::41841           :::*                LISTEN   -
```

## Cause

NFS 4.0 listens to this random port to support callback. NFS 4.0 listens to this random port to support callback. Because the default value of the `fs.nfs.nfs_callback_tcpport` kernel parameter is 0, the NFS 4.0 client randomly chooses a port to listen. This random port does not constitute a security risk.

To facilitate port management, see [Solution to fix the callback port](#).

### Solution

Before mounting the file system, set the parameter `fs . nfs . nfs_callba`  
`ck_tcpport` to a non-zero value.

```
sudo sysctl fs . nfs . nfs callba ck tcpport =< port >
```

In the following example, the `fs . nfs . nfs_callba ck_tcpport` parameter is manually set to port 45450, and then NFS 4.0 is mounted. Netstat shows that the listened port is 45450.

Note that the following commands are run as user root, so running the `sysctl` command in `sudo` is unnecessary.

```

[root@iz8v...:/]# netstat -tanl | grep tcp | grep LISTEN
tcp        0      0 0.0.0.0:111          0.0.0.0:*           LISTEN    1/systemd
tcp        0      0 0.0.0.0:22          0.0.0.0:*           LISTEN    2159/sshd
[root@iz8v...:/]# sysctl fs.nfs.nfs_callback_tcpport=45450
fs.nfs.nfs_callback_tcpport = 45450
[root@iz8v...:/]# mount -t nfs -o vers=4.0 915... .nas.aliyuncs.com:/ /mnt
[root@iz8v...:/]# netstat -tanl | grep tcp | grep LISTEN
tcp        0      0 0.0.0.0:45450       0.0.0.0:*           LISTEN    -
tcp        0      0 0.0.0.0:111        0.0.0.0:*           LISTEN    1/systemd
tcp        0      0 0.0.0.0:22         0.0.0.0:*           LISTEN    2159/sshd
tcp6       0      0 :::45450           :::*                 LISTEN    -

```

## 2.12 How do I avoid exceptions when multiple processes and multiple clients write the same log file at the same time

NAS enables multiple clients to share and write files in the same namespace using the NFS protocol. However, the NFS protocol does not support Atomic Append semantics.

### Problem

When multiple processes/clients write the same file (for example, the same log) at the same time, and each process independently maintains context information (such as file descriptor or write location), there may be coverage, crossover, and disordered content.

### Solution

Two solutions are provided:

- (Recommended) Allow different processes/clients to write to different files in the same file system, and then consolidate these files during your analysis. This solution is preferable as it can mitigate problems caused by concurrent writing without the need for any filelock or alterations to performance.
- Use the flock with seek method to ensure the atomicity and consistency of writing. The flock with seek method is a relatively time-consuming operation, that may significantly affect performance. The following steps detail how to implement the flock with seek method in Linux.

### Flock with seek method

As the NFS protocol does not support Atomic Append semantics, when multiple processes/clients write the same file (for example, the same log) at the same time, there may be coverage content. In Linux, you can simulate Atomic Append on the NFS file system by using flock + seek to protect and support concurrent append writing to the same file.

You can use flock + seek as follows:

1. Call `fd = open ( filename , O_WRONLY | O_APPEND | O_DIRECT )` to open the file by means of append writing, and specify `O_DIRECT` (bypass Page Cache to write directly) to acquire file descriptor `fd`.
2. Call `flock ( fd , LOCK_EX | LOCK_NB )` to get the filelock. In case of failure (for example, the filelock is already in use), the system returns an error. Retry or perform error handling.
3. Call `lseek ( fd , 0 , SEEK_END )` to point the current file offset (cfo) of the `fd` to the end of the file.
4. Perform normal write operations. The insert location is the end of the file. The filelock can prevent overwriting.
5. Call `flock ( fd , LOCK_UN )` to release the filelock after the write operation.

The following is a simple C language sample program.

```
# 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 )
    {
        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, see [Linux file locking mechanisms - Flock, Lockf, and Fcntl](#).



#### Note:

To use `flock ()` on the NAS file system, your Linux kernel version must be 2.6.12 or later. If your Linux kernel uses an earlier version, use `fcntl ()`.

## 2.13 What impacts the I/O performance of Windows service SMB protocol

### Symptom

By default, the `large mtu` option is disabled on a Windows SMB client, which affects the increase in I/O performance.

## Solution

You can modify the following registry key to enable the `large mtu` option:

```
HKLM \ System \ CurrentControlSet \ Services \ LanmanWork station \  
Parameters
```

Create a `DWORD` at this location with the key named `DisableLargeMtu` and value set to `0`. Restart the file system to apply the change.

## 2.14 How to improve performance when using IIS to access NAS

### Problem description

When IIS accesses a file by using a NAS share, the backend of IIS will frequently access NAS. Unlike accessing a local file system, you must interact with networks when accessing NAS. Even if it takes a short time for each interaction, the total amount of time increases with an increasing number of clients.

### Solutions

For more information, see [SMB2 Client Redirector Caches Explained](#).

You can increase the values of the following registry keys. For example, you can change the values to 600 or above.

The path of the registry key is `HKEY_LOCAL_MACHINE \ SYSTEM \ CurrentControlSet \ services \ LanmanWork station \ Parameters`.

The registry keys are listed as follows:

- `FileInfoCacheLifetime`
- `FileNotFoundCacheLifetime`
- `DirectoryCacheLifetime`



#### Note:

- When none of the preceding keys exists, troubleshoot the issue as follows:
  1. Ensure that SMB is used rather than NFS.
  2. Ensure that the current version of Windows supports these registry keys. When the current version of Windows supports these registry keys but they do not

exist, you can manually create these registry keys. For more information, see [Performance tuning for file servers](#).

- For web files that are frequently accessed by IIS, such as js and css scripts, we recommend that you move these files to a local PC.

## 2.15 Does NAS support inotify

While inotifywait is commonly used in combination with rsync to backup/synchronize data on a quasi-real-time basis, it may not work properly on NAS file systems due to the implementation of inotify.

### How inotify works

inotify is a sub-module of the Linux kernel, and inotifywait is the user-mode interface of inotify. inotify is realized at the VFS layer. When file operations reach the VFS layer, the inotify module sends the operation type (creation/deletion/attribute change, and so on) and operation object (file name) to the user-mode, and the user-mode inotifywait then outputs the operation information to the user.

### Problem

Because inotify is implemented at the VFS layer of the kernel, the local kernel cannot recognize operations made by a remote client on the NFS file system. Therefore, inotify cannot recognize modifications on files made by the remote client.

If you Mount the same NAS file system simultaneously on Client A and Client B, and enable inotifywait at Client A to monitor the mounted directory, the following occurs:

- inotifywait recognizes operations on files in the mounted directory on Client A.
- inotify cannot recognize any operation on files in the mounted directory on Client B.

### Solution

An alternative solution is to use [FAM](#).

FAM is a library used for monitoring files or directories, and it is fully implemented in user-mode. You then only need to run a daemon in the background to regularly scan the directory and check for file changes.

However, using FAM has the following issues:

- You must write a program to call the FAM interface implementation function.

- In scenarios with a large number of files, using FAM may have poor performance and consume a lot of resources. Furthermore, it cannot ensure real-time monitoring.

## 2.16 Why does my NFS run in Linux at MB speed

The Linux nfs client has restrictions on the number of NFS requests which can be simultaneously initiated. If the parameter is too small, it causes relatively poor I/O performance. In such cases, check the following parameter:

```
cat /proc/sys/sunrpc/tcp_slot_table_entries
```

For the default kernel, the maximum value of this parameter is 256. For better performance, you can increase the value of this parameter. Run the following command as root:

```
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
sysctl -w sunrpc.tcp_slot_table_entries = 128
```

After you apply the change, you must re-mount the file system or restart the client.

## 2.17 How do I use the noac option

### Problem

A user mounts the same network file system on two ECS servers (ESC-A and ESC-B). The user writes data in append mode on ECS-A, and monitors file content changes with the `tail -f` command on ECS-B.

After data is written on ECS-A, the file content changes on ECS-B may experience latency of up to 30 seconds.

However, if a file is directly opened (such as using `vi`) on ECS-B under the same conditions, the updated content is visible immediately.

### Analysis

This is related to the `mount` option and the `tail -f` implementation.

The user uses the following mount command: `mount -t nfs4 /mnt /`

For file systems mounted on ECS-B using the NFS protocol, the kernel maintains a copy of metadata cache for the file and directory attributes. The cached file and directory attributes (including permission, size, and time stamp) are used to reduce the `NFSPROC_GETATTR` `RPC` requests.

The `tail -f` command uses `sleep + fstat` to monitor changes to the file attributes (primarily the file size), read files, and then output the results. However, file content output by using the `tail -f` command is dependent on the `fstat` result. Due to the metadata cache, the `fstat` command may not be monitoring real-time file attributes. Therefore, even if the file has been updated on the NFS server, the `tail -f` command cannot detect in real time whether the file has been changed or not, resulting in the latency.

### Solution

Use the `noac` option of the `mount` command to disable the caching of file and directory attributes. The command is as follows:

```
mount -t nfs4 -o noac /mnt /
```



## 3 SMB FAQ

---

### 3.1 SMB basic operation FAQ

Why is the disconnected state displayed when I use the net use command to view the status of a mount point?

If no operation is performed on a file system within 15 minutes, the connection is disconnected. The connection is established whenever an operation starts.

What is the maximum capacity and performance of a CIFS or SMB file system?

Currently, when an SMB file system is deployed on a NAS Capacity cluster, the maximum capacity and bandwidth for a single file system are subject to NAS Capacity. Other features, such as supports for a unique namespace, VPCs, and classic networks are the same as those of an NFS file system.

For more information, see [Network Attached Storage](#).

Supported protocols and operating systems for an SMB file system

For more information, see [Restrictions on the SMB protocol](#).

For more information about unsupported features for an SMB file system, see [Unsupported features for an SMB file system](#).

Restrictions when accessing an SMB file system

Similar to accessing an NFS file system, you cannot access an SMB file system from an ECS instance that is located in another region or from the Internet. You must connect to a VPC by using a dedicated leased line to access the file system.

To access a file system from external networks outside the VPC where the file system is located, see the following sections:

- [Access NAS from an on-premises IDC using a VPN network](#)
- [Access NAS from an on-premises IDC using NAT](#)
- [Mount NAS file systems on ECS instances that are located in multiple VPCs](#)
- [Mount NAS file systems on ECS instances that are owned by multiple accounts](#)

## 3.2 SMB mount issue FAQ

You can not mount a NAS file system whose protocol type is SMB on an ECS instance that runs Windows Server Version 1709.

- Symptoms

An error occurred while running the following command.

```
C : \ Users \ Administra tor > net use z : \\ xxxxxxxx - xxxx .
ap - southeast - 1 . nas . aliyuncs . com \ myshare
System error 1272 has occurred .
You can ' t access this shared folder because your
organizati on ' s security policies block unauthenti
cated guest access . These policies help protect
your PC from unsafe or malicious devices on the
network .
```

- Solutions

The preceding issue occurs when the security policies of Windows Server Version 1709 do not allow access to a remote shared directory by guest users.

You can use the following steps to resolve the issue:

- Locate and modify the following key.

```
[ HKEY_LOCAL _MACHINE \ SYSTEM \ CurrentCon trolSet \ Services \
LanmanWork station \ Parameters ]
" AllowInsec ureGuestAu th " = dword : 0
```

Change to:

```
[ HKEY_LOCAL _MACHINE \ SYSTEM \ CurrentCon trolSet \ Services \
LanmanWork station \ Parameters ]
" AllowInsec ureGuestAu th " = dword : 1
```

- Switch to Powershell and enter the following command:

```
New - ItemProper ty - Path $ registryPa th - Name $ name -
Value $ value - PropertyTy pe DWORD - Force
```

For more information, see [Guest access in SMB2 disabled by default in Windows 10, Windows Server 2016 version 1709, and Windows Server 2019](#).

How to mount an SMB file system on an ECS instance that runs Linux

Currently, you cannot mount a NAS file system whose protocol type is SMB on an ECS instance that runs Linux.

To mount a file system, use the following Linux command.

```
mount -t cifs [ MOUNT POINT ] [ LOCAL DIRECTORY ] -o  
username = guest , vers = 3 . 0
```



**Note:**

Use version-specific Linux commands to mount an SMB file system.

When using the `-o username = guest` switch for a mount command by using a later version of Linux, you are required to enter a password. We recommend that you use the following command.

```
mount -t cifs [ MOUNT POINT ] [ LOCAL DIRECTORY ] -o  
guest , vers = 3 . 0
```

An error occurred while connecting to the SMB file system

### Symptoms

This issue occurs when both an NFS and SMB file system are used. When you fail to mount an NFS file system by using the `net use` command, an error occurs when you mount an SMF file system.

### Solutions

Ensure that a file system to be mounted is running as expected. You can stop mounting the file system and retry in five minutes. If the issue still persists, open a ticket.

A mounted SMB directory is only available to the administrator.

This issue occurs when user accounts in Windows are isolated from one another.

You need to create a shared link for multiple users. For example, you can use the following command to create a shared link named `myshare` in drive C.

```
mklink /D C :\ myshare \\ xxxxxxxx - xxxx . cn - beijing . nas .  
aliyuncs . com \ myshare \
```

An error occurred while loading files that are located in a shared SMB volume by using IIS on an ECS instance that runs Windows Server 2016

For more information, see [Install and configure Active Directory domains](#).

## Best practices for deploying IIS in a NAS file system

An HTTP error 500.19 with error code 0x8007003a occurred while mounting an SMB file system on an ECS instance that runs Windows 2016. For more information, see [Best practices for deploying IIS in a NAS file system](#).

## Can I mount an NFS file system and SMB file system on the same ECS instance?

No.

We recommend that you do not access an SMB file system using a Linux client due to several encoding issues. 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 characters encoded with Unicode is supported. However, in Linux, a maximum length of 255 characters encoded with UTF-8 is supported.

If this feature is required for specific uses, you can mount an SMB file system on an ECS instance that runs Linux whose kernel supports SMB2.

Run either one of the following mount commands: `mount -t cifs -o vers=2.0 \< a mount point > \ myshare / mnt` or `mount -t cifs -o vers=2.0 //< a mount point > / myshare / mnt`. If a password is required, press Enter.

Ensure that a CIFS file system is supported by a specific Linux kernel: In the `/ boot` directory, check the value of `CONFIG_CIFS`. A value of `y` or `m` indicates that a CIFS file system can be mounted.

The details are shown in the following figure.

```
$grep -i cifs /boot/config-2.6.18-274.elios5.1
CONFIG_CIFS=m
```



### Note:

- You need to install `cifs-utils` before running the preceding command. Take CentOS as an example. You can run the following command to install `cifs-utils`.

```
yum install samba - client samba - common cifs - utils
```

- If one or more features of CIFS are not supported on the current version of Linux, we recommend that you upgrade the Linux kernel to version 3.10.0-514 or later.

## 3.3 Performance issues of Server Message Block

Latency issues before you perform I/O operations

When you access a Server Message Block (SMB) server by using a mount point, you need to wait for several minutes before performing I/O operations.

What can I do to reduce the waiting period when accessing an SMB server?

### Solutions

The waiting period that occurs is mainly caused by an NFS client or Web client.

- Check whether an NFS client is installed. If the NFS client is no longer used, we recommend that you delete it.
- Locate the following registry key. The path to the registry key is: `HKEY_LOCAL_MACHINE -> System -> CurrentControlSet -> Control -> NetworkProvider -> Order -> ProviderOrder`.

Assume that the value of the ProviderOrder key is `LanmanWorkstation`, `RDPNP`, `Nfsnp`. You need to remove `Nfsnp` and restart the ECS instance.

- When a Web client exists, this increases the latency when you access an SMB server by using a file manager. We recommend that you remove the Web client.



#### Note:

When a client connects to an SMB server for the first time, the latency is higher than expected. Check whether you can communicate with the mount address of the SMB server by using the ping command, or check if the latency for the communication is as expected.

- If a time-out error occurred while using the ping command, we recommend that you check the network settings.
- If the latency is higher than expected, we recommend that you ping the IP address of the SMB server. When the latency to ping the IP address of a mount point is lower than the latency to ping the domain name of the mount point, the issue may be caused by the DNS settings. We recommend that you check the DNS settings.

### Procedure to solve performance issues

1. Modify the value of the `ProviderOrder` key. When the latency to access the SMB server is longer than usual, we recommend that you check this value.

## 2. You can use fio to conduct a performance test to check the issue.

```
fio . exe -- name =./ iotest1 -- direct = 1 -- rwmixread = 0 --  
rw = write -- bs = 4K -- numjobs = 1 -- thread -- iodepth = 128  
-- runtime = 300 -- group_repo rting -- size = 5G -- verify =  
md5 -- randrepeat = 0 -- norandomma p -- refill_buf fers --  
filename =\\< mount point dns >\ myshare \ testfio1  
  
fio . exe -- name =./ iotest1 -- direct = 1 -- rwmixread = 0 --  
rw = write -- bs = 4K -- numjobs = 1 -- thread -- iodepth = 128  
-- runtime = 300 -- group_repo rting -- size = 5G -- verify =  
md5 -- randrepeat = 0 -- norandomma p -- refill_buf fers --  
filename =\\< mount point dns >\ myshare \ testfio1
```

3. For applications that use an SMB file system as data storage, try to perform read/write operations by using large data blocks. The smaller the data blocks, the more network resources are consumed. If you cannot modify the size of a data block, you can use `BufferedOutputStream`.

## 3.4 Causes and solutions of SMB mount failures

This topic describes the possible causes and solutions when an SMB file system fails to mount on an ECS instance that runs Windows.

### System error 53

#### Description

The network path was not found.

#### Cause

The network connection fails.

#### Solution

You can check the network connection as follows:

1. Use the following `ping` command to check whether you can access the IP address of a mount point and ensure that the latency is within the expected range.

```
ping < the IP address of a mount point >
```

- If you can ping the IP address, go to [Step 2](#).
- If you fail to ping the IP address, perform the following steps:
  - Check whether the mount command is valid. For example, no redundant or missing forward slash /, backslash \, space, and `myshare` exists.


The following provides the valid format of a mount command.

```
net use < the target drive > \\< the IP address of a mount point > \ myshare
```

Example:

```
net use z : \\ xxxx . cn - hangzhou . nas . aliyuncs . com \ myshare
```

- Ensure that the protocol type of a file system is SMB.

File System ID/Name	Storage Type	Protocol Type	Storage Capacity	Zone	Bound Storage Package	Number of Mount Points	Action
	Capacity-type	SMB	0 B	China North 1 Zone C	No	1	<a href="#">Add Mount Point</a>   <a href="#">Manage</a>   <a href="#">Delete</a>

- Ensure that the IP address of the mount point is valid.
  - Ensure that the ECS instance and the mount point are located in the same VPC.
  - Ensure the network configuration is valid. We recommend that you check the network configuration if the ECS instance and the mount point are located in different VPCs. We also recommend that you check the network configuration if the ECS instance is connect to the VPC by using a VPN.
2. Use the following `telnet` command to check whether the SMB file system is available.

```
telnet < the IP address of a mount point > 445
```

## System error 58

### Description

The specified server cannot perform the requested operation.

### Cause

The SMB protocol version used on the ECS instance is not compatible with the version used by the file system.

### Solution

Ensure that the ECS instance runs Windows Server 2008 R2 or later, excluding Windows Server 2008.

## System error 64

### Description

The specified network name is no longer available.


### Cause

- The IP address of the target ECS instance is not specified in any NAS permission group that is bound to the file system.
- The NAS service is overdue.
- The ECS instance and the NAS file system are located in classic networks, but they are created by different Alibaba Cloud accounts.
- The protocol type of a file system is not SMB.

### Solution

This issue occurs when you do not have permission to access NAS file system resources. You can troubleshoot the issue as follows:

1. Ensure that the private IP address or the VPC IP address of the ECS instance is specified in a permission group that is bound to the file system.
2. Confirm that the balance of the Alibaba Cloud account is positive.
3. Ensure that the ECS instance and the NAS file system are created by the same Alibaba Cloud account if they are located in classic networks.
4. Ensure that the protocol type of a file system is SMB.

File System ID/Name	Storage Type	Protocol Type	Storage Capacity	Zone	Bound Storage Package	Number of Mount Points	Action
	Capacity-type	SMB	0 B	China North 1 Zone C	No	1	<a href="#">Add Mount Point</a>   <a href="#">Manage</a>   <a href="#">Delete</a>

## System error 67

### Description

The network name cannot be found.

### Cause



Several important network services are not started.

#### Solution

Start the following services. For more information, see [Mount an SMB file system](#).

1. Workstation
2. TCP/IP NetBIOS Helper

#### System error 85

##### Description

The local device name is already in use.

##### Cause

The target drive letter is already in use.

##### Solution

You must change the target drive letter and re-mount the file system.

#### System error 1272

##### Description

Error message: You can't access this shared folder because your organization's security policies block unauthenticated guest access. These policies can help protect your computer from threats from insecure or malicious devices on the network.

##### Cause

Due to security policies, Windows denies access to an SMB file system from Guest users.

##### Solution

If the ECS instance runs a version later than Windows Server 2016 (excluding Windows Server 2016), modify the following registry key 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](#).