

# Dedicated Distributed Storage Service

## FAQs

**Issue** 01  
**Date** 2019-02-26



**Copyright © Huawei Technologies Co., Ltd. 2022. All rights reserved.**

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

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

## **Notice**

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

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

---

## Contents

---

<b>1 What Are the Risks of Not Expanding the Storage Pool Capacity?.....</b>	<b>1</b>
<b>2 How Many Statuses Does a Storage Pool Have?.....</b>	<b>2</b>
<b>3 How Many Statuses Does a Disk Have?.....</b>	<b>3</b>
<b>4 Troubleshooting and Impacts on the DSS Usage.....</b>	<b>5</b>
<b>5 Can I Attach a Disk to Multiple Servers?.....</b>	<b>6</b>
<b>6 Will Data in the DSS Disk Be Lost When the DSS Disk Is Detached?.....</b>	<b>7</b>
<b>7 What Should I Do If an Error Occurs on My DSS Disk?.....</b>	<b>8</b>
<b>8 How Can I Test DSS Disk Performance?.....</b>	<b>9</b>
<b>9 What Are the Restrictions on Attaching a Disk to an ECS?.....</b>	<b>17</b>
<b>10 What Are the Precautions for Detaching a Disk from an ECS?.....</b>	<b>18</b>
<b>11 Why My Disk Cannot Be Attached to a Server?.....</b>	<b>21</b>
<b>12 Can the Disk Space Usage Be Reduced If I Delete Files on a Server?.....</b>	<b>22</b>

# 1 What Are the Risks of Not Expanding the Storage Pool Capacity?

---

In the process of using a DSS storage pool, if the ratio of **Used Capacity (GB)** to **Total Available Capacity (GB)** exceeds 75%, the system will remind you to expand the storage pool capacity. If this ratio reaches 100%, write protection will be automatically enabled for the pool, data can no longer be written to disks, and your services will be affected. Therefore, you are advised to expand storage pool capacity when this ratio reaches 75%.

# 2 How Many Statuses Does a Storage Pool Have?

---

A storage pool has several statuses. [Table 2-1](#) lists the meaning of each status and the operations for each status.

**Table 2-1** Storage pool status

Status	Description	Allowed Operation
Deploying	The storage pool is being deployed.	-
Available	The storage pool is successfully created, and you can create disks in the storage pool.	Creating disks
Expanding	The storage pool capacity is being expanded and cannot be used.	-

# 3 How Many Statuses Does a Disk Have?

A disk has several statuses. [Table 3-1](#) lists the meaning of each status and the operations for each status.

**Table 3-1** Disk status description

Status	Description	Allowed Operation
In-use	The disk is attached to a server and in use.	<ul style="list-style-type: none"><li>• Detaching</li><li>• Creating backups</li></ul> <b>NOTE</b> If a shared disk is in the <b>In-use</b> state, the disk can be attached.
Available	The disk is successfully created and has not been attached to any server.	<ul style="list-style-type: none"><li>• Attaching</li><li>• Expanding</li><li>• Deleting</li></ul>
Creating	The disk is being created.	None
Attaching	The disk is being attached to a server.	None
Detaching	The disk is being detached from a server.	None
Deleting	The disk is being deleted.	None
Expanding	The capacity of the disk is being expanded.	None
Uploading	Data on the disk is being uploaded to an image. This status occurs when you create an image from a server.	None
Downloading	Data is being downloaded from an image to the disk. This status occurs when you create a server.	None
Error	An error occurs when you try to create a disk.	Deleting

Status	Description	Allowed Operation
Deletion failed	An error occurs when you try to delete a disk.	No operations can be performed.
Expansion failed	An error occurs when you try to expand the capacity of a disk.	Deleting

 **NOTE**

If a DSS disk status is **Error**, **Deletion failed**, or **Expansion failed**, you can rectify the error by following the steps provided in [What Should I Do If an Error Occurs on My DSS Disk?](#)

# 4 Troubleshooting and Impacts on the DSS Usage

---

- **Case one**  
If server or disk failures cause a disk being removed from the storage pool, the total available capacity becomes smaller. After the fault is rectified, the total available capacity can be restored to the original value.
- **Case two**  
If a server or disk is faulty, and no disk is removed from the storage pool, the storage pool is degraded, which does not affect the use of the storage pool. After the fault is rectified, the storage pool becomes normal.

 **NOTE**

Storage pool degradation refers to that the number of data copies for some data in a storage pool is reduced from three copies to two copies and cannot be automatically restored without manual troubleshooting.

- **Case three**  
If a server or disk is faulty, your services may be interrupted. For example, if the used capacity is 98% and a disk is removed due to server or disk failures, the total available capacity decreases. As a result, the proportion of the used capacity to the total available capacity reaches nearly 100%. The write protection is enabled for the storage pool, and your services are interrupted.



# 5 Can I Attach a Disk to Multiple Servers?

---

A non-shared disk can be attached to one server only.

A shared disk can be attached to a maximum of 16 servers.

 **NOTE**

Shared disks are a type of DSS disks and can be attached to multiple servers.

# 6 Will Data in the DSS Disk Be Lost When the DSS Disk Is Detached?

---

Not necessarily.

To prevent data loss when you detach a DSS disk, perform the following operations:

1. Stop the ECS to which the to-be-detached disk has been attached.
2. After server has been stopped, detach the disk.

# 7 What Should I Do If an Error Occurs on My DSS Disk?

If an error occurs, the disk may show one of the states listed in [Table 7-1](#). Take the following measures to handle the exceptions.

**Table 7-1** Solutions for disk errors

Error Status	Handling Suggestion
Error	Delete the disk in the <b>Error</b> state and create another one.
Deletion failed	Contact customer service.
Expansion failed	Customer service personnel will contact you and help you handle this error. Do not perform any operations on the disk before they contact you. If you require that the error be handled as soon as possible, contact customer service.

# 8 How Can I Test DSS Disk Performance?

## Precautions

In the disk performance test, if the start sector number is not 4-KB aligned, the disk performance will be greatly affected. Ensure that the start sector number is 4-KB aligned before you start the test.

### NOTE

To test the performance of a shared disk, the following requirements must be met:

- The shared disk must be attached to multiple servers (ECSs or BMSs).
- If the shared disk is attached to multiple ECSs, these ECSs must belong to the same anti-affinity ECS group.

If these ECSs fail to meet the anti-affinity requirement, the optimal performance of the shared disk cannot be achieved.

The testing process for Windows and Linux is different.

- [Windows](#)
- [Linux](#)

If the test results do not meet your expectation, see [Why Does My Disk Performance Test Using Fio Have Incorrect Results?](#)

## Windows

The way you test disk performance depends on the server OS. This section uses Windows 7 Professional 64-bit as an example. For other Windows OSs, see the corresponding OS documentations.

Install the performance measurement tool Iometer before the test. You can obtain the tool at <http://www.iometer.org/>.

**Step 1** Log in to the server.

**Step 2** Press **win+R** to open the **Run** window. Enter **msinfo32** and click **OK**.

The system information window is displayed.

**Step 3** Choose **components > storage > disks**. In the right pane, view the partition offset value.

- If 4096 can be divided by the parameter value, the partition is 4-KB aligned. Go to [Step 4](#).
- If 4096 cannot be divided by the parameter value, the partition is not 4-KB aligned. Ensure the 4-KB alignment for the partition before continuing the test.

### NOTICE

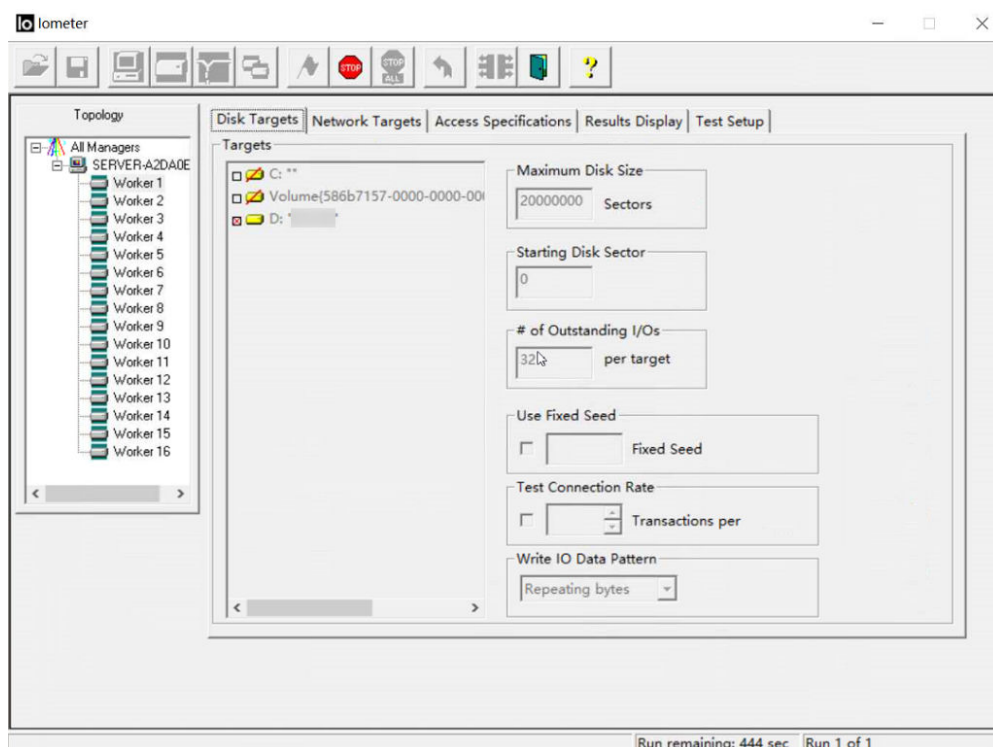
If you delete the partition and select another start sector number for 4-KB alignment, you will lose all the data on that partition. Exercise caution when performing this operation.

**Step 4** Use iometer to test the disk performance. For details, see the iometer product documentation.

When the disk IOPS and throughput are tested, the parameter settings for iometer and fio are the same. For details, see [Table 8-1](#).


The following example uses iometer to test the disk performance.

1. Set the workflow.




2. Set the test run time.


In this example, the test run time is set to 10 minutes, with 60 seconds ramp up time. Disk performance is tested after the writes are stable.



3. Set the data block size and read/write policy. In this example, the disk size is set to 64 KB, the policy is 100% sequential write.



4. View the test results.



----End

## Linux

If you use an old version Linux OS, for example CentOS 6.5, and run **fdisk** to create partitions, the default start sector number will not be 4-KB aligned, which will greatly affect the test performance. For this reason, if such an OS is used, you are advised to select a new start sector number, one that is 4-KB aligned, when creating partitions.

The way you test disk performance depends on the server OS. This section uses CentOS 7.2 64-bit as an example. For other Linux OSs, see the corresponding OS documentations.

Install the performance measurement tool, fio, before the test.

**Step 1** Log in to the server and switch to user **root**.

**Step 2** Before you start the test, run the following command to check whether the start sector number is 4-KB aligned:

**fdisk -lu**

Information similar to the following is displayed:

```
[root@ecs-centos sdc]# fdisk -lu
Disk /dev/xvda: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x7db77aa5

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1 *        2048     20968919    10483436   83  Linux
```

```
Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk /dev/xvdc: 53.7 GB, 53687091200 bytes, 104857600 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x3cf3265c
```

```
Device Boot      Start         End      Blocks   Id  System
/dev/xvdc1        2048     41943039     20970496   83  Linux
```

- If 8 can be divided by the start sector number, the number is 4-KB aligned. Go to [Step 3](#).
- If 8 cannot be divided by the start sector number, the number is not 4-KB aligned. Delete the partition and select a 4-KB aligned start sector number for the new partition before continuing the test.

---

**NOTICE**

If you delete the partition and select another start sector number for 4-KB alignment, you will lose all the data on that partition. Exercise caution when performing this operation.

---

**Step 3** Run the following commands and use fio to test the disk performance:

- To test random write IOPS, run the following command: **fio -direct=1 -iodepth=128 -rw=randwrite -ioengine=libaio -bs=4k -size=10G -numjobs=1 -runtime=600 -group\_reporting -filename=/opt/fiotest/fiotest.txt -name=Rand\_Write\_IOPS\_Test**
- To test random read IOPS, run the following command: **fio -direct=1 -iodepth=128 -rw=randread -ioengine=libaio -bs=4k -size=10G -numjobs=1 -runtime=600 -group\_reporting -filename=/opt/fiotest/fiotest.txt -name=Rand\_Read\_IOPS\_Test**
- To test write throughput, run the following command: **fio -direct=1 -iodepth=32 -rw=write -ioengine=libaio -bs=1024k -size=10G -numjobs=1 -runtime=600 -group\_reporting -filename=/opt/fiotest/fiotest.txt -name=Write\_BandWidth\_Test**
- To test read throughput, run the following command: **fio -direct=1 -iodepth=32 -rw=read -ioengine=libaio -bs=1024k -size=10G -numjobs=1 -runtime=600 -group\_reporting -filename=/opt/fiotest/fiotest.txt -name=Read\_BandWidth\_Test**
- To test random read latency, run the following command: **fio -direct=1 -iodepth=1 -rw=randread -ioengine=libaio -bs=4k -size=10G -numjobs=1 -runtime=60 -group\_reporting -filename=/opt/fiotest/fiotest.txt -name=Rand\_Read\_LATE\_Test**



**NOTICE**

- When using fio to perform a raw disk performance test, ensure that no partitions and file systems have been created on the disk and there is no data stored on the disk. Or, the raw disk test will damage the file system, and data on the disk will become read-only. In this case, your only option will be to delete the disk and buy a new one to continue the test.
- Do not perform the test on a disk with service data on it. If such test is a must, you are advised to perform the test as follows:
  - Back up the disk data before the test as you may damage the data on the disk.
  - Specify a file, for example **-filename=/opt/fiotest/fiotest.txt**, to test the performance of the file system.

**Table 8-1** lists the fio test parameters.

**Table 8-1** Parameter description

Parameter	Description
direct	Defines whether direct I/O is used. <ul style="list-style-type: none"> <li>- Set to <b>0</b>: buffered I/O is used.</li> <li>- Set to <b>1</b>: direct I/O is used.</li> </ul>
iodepth	Defines the I/O queue depth. The default value is <b>1</b> . This queue depth refers to the queue depth of each thread. That said, when multiple threads are tested, this parameter defines the queue depth of each thread. Total concurrent I/Os = iodepth x numjobs
rw	Defines the test read/write policy. <ul style="list-style-type: none"> <li>- <b>randread</b>: random read</li> <li>- <b>randwrite</b>: random write</li> <li>- <b>read</b>: sequential read</li> <li>- <b>write</b>: sequential write</li> <li>- <b>randrw</b>: mixed random read/write</li> </ul>

Parameter	Description
ioengine	<p>Defines how fio delivers the I/O request (synchronously or asynchronously).</p> <ul style="list-style-type: none"> <li>- Synchronous I/O: Only one I/O request is delivered at a time, and the response is returned after the kernel has processed the request. That said, the single-thread I/O queue depth is always less than 1, and multi-thread concurrent processing can be used to handle such issues. Normally, 16 to 32 concurrent working threads fully occupy the I/O queue depth.</li> <li>- Asynchronous I/O: Multiple I/O requests are delivered using libaio at a time. Wait for the process to complete and reduce the interaction times to improve efficiency.</li> </ul>
bs	<p>Defines the I/O block size. The unit can be KB, Kb, MB, and Mb, and the default value is 4 KB.</p>
size	<p>Defines the amount of data processed by the test I/Os. If parameters, such as <b>runtime</b>, are not specified, the test ends when fio has processed all the specified data amount.</p> <p>The value can be a number with a unit or percentage. A number with a unit indicates the read/write data amount, for example <b>size=10G</b>, indicating a 10-GB read/write data amount. A percentage indicates the ratio of read/write data amount to the capacity of total files, for example <b>size=20%</b>, indicating the read/write data amount takes 20% of the total file space.</p>
numjobs	<p>Defines the number of concurrent threads.</p>
runtime	<p>Defines the test time.</p> <p>If this parameter is not specified, the test ends until the specified data amount is processed by the block size defined using parameter <b>size</b>.</p>
group_reporting	<p>Defines the test result display mode. The parameter value displays the statistics on a single thread instead of that on all jobs.</p>
filename	<p>Defines the name of the test file or device.</p> <ul style="list-style-type: none"> <li>- If a file is specified, the performance of the file system is tested. Example: <b>-filename=/opt/fiotest/fiotest.txt</b></li> <li>- If a device name is specified, the performance of the raw disk is tested. Example: <b>-filename=/dev/vdb1</b></li> </ul> <p><b>NOTICE</b> If the test is performed on a disk already has partitions and file systems created as well as data on it, user parameter <b>filename</b> to specify a file so that the original file system is not damaged and the data is not overwritten.</p>

Parameter	Description
name	Defines the test task name.

----End

# 9 What Are the Restrictions on Attaching a Disk to an ECS?

---

- The disk and the target ECS must be located in the same AZ.
- The target ECS must be in **Running** or **Stopped** state.
- A frozen disk cannot be attached to an ECS.
- For ECSs charged in yearly/monthly mode:
  - If you detach the system disk purchased when you buy an ECS and want to continue using it as a system disk, you can only attach it to an ECS that has the same specifications as those of the original ECS. If you want to use it as a data disk, you can attach it to any ECS.
  - If you detach the non-shared data disk purchased when you create an ECS and want to attach it again, you can only attach it to the original ECS as a data disk.

# 10 What Are the Precautions for Detaching a Disk from an ECS?

---

## Scenarios

A disk attached to an ECS can function as a system disk or data disk.

- Disks mounted on **/dev/sda** or **/dev/vda** function as system disks. Currently, system disks must be detached offline. In this case, you must stop the ECS before detaching a system disk.
- Disks mounted on other mount points function as data disks. In addition to offline detachment, a data disk can also be detached online if the server OS supports online detachment. In this case, you do not need to stop the running ECS.

This FAQ describes how to detach a disk from a running ECS.

## Constraints

- The disk to be detached must be mounted on a mount point other than **/dev/sda** or **/dev/vda**.  
Disks mounted on **/dev/sda** or **/dev/vda** are system disks and cannot be detached from running ECSs.
- Before detaching a disk from a running Windows ECS, make sure that VMTtools have been installed on the ECS and that the tools are running properly.
- Before detaching a disk from a running Windows ECS, ensure that no program is reading data from or writing data to the disk. Otherwise, data will be lost.
- SCSI disks cannot be detached from running Windows ECSs.
- Before detaching a disk from a running Linux ECS, you must log in to the ECS and run the **umount** command to cancel the association between the disk and the file system. In addition, ensure that no program is reading data from or writing data to the disk. Otherwise, the disk will fail to be detached.

## OSs Supporting Disk Detachment from a Running ECS

OSs supporting disk detachment from a running ECS include two parts:

- For the first part, see [Formats and OSs Supported for External Image Files](#).
- For the second part, see [Table 10-1](#).

**Table 10-1** OSs supporting disk detachment from a running ECS

OS	Version
CentOS	7.3 64bit
	7.2 64bit
	6.8 64bit
	6.7 64bit
Debian	8.6.0 64bit
	8.5.0 64bit
Fedora	25 64bit
	24 64bit
SUSE	SUSE Linux Enterprise Server 12 SP2 64bit
	SUSE Linux Enterprise Server 12 SP1 64bit
	SUSE Linux Enterprise Server 11 SP4 64bit
	SUSE Linux Enterprise Server 12 64bit
OpenSUSE	42.2 64bit
	42.1 64bit
Oracle Linux Server release	7.3 64bit
	7.2 64bit
	6.8 64bit
	6.7 64bit
Ubuntu Server	16.04 64bit
	14.04 64bit
	14.04.4 64bit
Windows	Windows Server 2008 R2 Enterprise 64bit
	Windows Server 2012 R2 Standard 64bit
	Windows Server 2016 R2 Standard 64bit
Redhat Linux Enterprise	7.3 64bit
	6.8 64bit

 **NOTE**

Online detachment is not supported by the ECSs running OSs not listed in the preceding table. For such ECSs, stop the ECSs before detaching disks from them to prevent any possible problems from occurring.

# 11 Why My Disk Cannot Be Attached to a Server?

---

If your disk cannot be attached, verify that your disk meets the following conditions:

- The disk is in the **Available** state if it is a non-shared disk.
- The disk is in the **Available** or **In-use** state if it is a shared disk, and the maximum number of servers that the disk can be attached to is not reached.
- The disk is not added to any replication pair in SDRS. If it has been added to a replication pair, delete the replication pair and try again.
- The disk is not frozen.
- The disk is not locked by any service.



# 12 Can the Disk Space Usage Be Reduced If I Delete Files on a Server?

---

No.

When you delete a file, the system adds a deletion label to the header of the file. This does not reduce the disk space usage.