

Alibaba Cloud Elastic Compute Service

Block Storage

Issue: 20190916

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.	 Danger: 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.	 Warning: 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.	 Notice: 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.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK.
Courier font	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 What is block storage?.....	1
2 Block storage performance.....	3
3 Block storage.....	12
3.1 Cloud disk overview.....	12
3.2 ESSD cloud disk.....	15
3.3 Triplicate technology.....	18
3.4 ECS disk encryption.....	20
3.5 Create a cloud disk.....	24
3.5.1 Create a Pay-As-You-Go cloud disk.....	24
3.5.2 Create a Subscription cloud disk.....	26
3.5.3 Create a cloud disk by using a snapshot.....	29
3.6 Attach a cloud disk.....	31
3.7 Format a data disk.....	35
3.7.1 Format a data disk for a Windows instance.....	35
3.7.2 Format a data disk of a Linux instance.....	37
3.7.3 Partition and format data disks larger than 2 TiB.....	42
3.7.4 Create a file system on a raw data disk.....	49
3.8 Resize cloud disks.....	52
3.8.1 Overview.....	52
3.8.2 Resize cloud disks online.....	54
3.8.3 Resize cloud disks offline.....	57
3.8.4 Extend a Windows file system.....	59
3.8.5 Resize partitions and file systems of Linux system disks.....	64
3.8.6 Resize partitions and file systems of Linux data disks.....	68
3.8.7 Update the Red Hat VirtIO SCSI driver for Windows ECS instances.....	82
3.9 Change the operating system.....	85
3.9.1 Replace the system disk (non-public image).....	85
3.9.2 Replace the system disk by using a public image.....	91
3.10 Roll back a cloud disk.....	98
3.11 Reinitialize a cloud disk.....	100
3.11.1 Reinitialize a cloud disk.....	100
3.11.2 Reinitialize a data disk.....	108
3.12 Convert billing methods of cloud disks.....	110
3.13 Modify the performance level of an ESSD cloud disk.....	110
3.14 Monitor a cloud disk.....	112
3.15 Detach a cloud disk.....	114
3.16 Release a data disk.....	117

4 Shared Block Storage.....	119
5 Local disks.....	122
6 Block Storage FAQ.....	128

1 What is block storage?

Overview

Block storage is a high-performance, low latency block storage service for Alibaba Cloud ECS. Similar to a hard disk, you can format block storage and create a file system on it to easily meet the data storage needs of your business.

Alibaba Cloud provides a variety of block-level storage products based on a distributed storage architecture and local disks located on the physical servers where ECS instances are hosted. Specifically, the storage products are as follows:

- [Cloud Disk](#), which is a block-level data storage product provided by Alibaba Cloud for ECS, uses a [multiple distributed system](#), and features low latency, high performance, persistence, high reliability, and more. Cloud disks can be created, resized, and released at any time.
- Shared block storage is a block-level data storage device that supports simultaneous read and write access to multiple ECS instances. Similar to the cloud disk, shared block storage uses a [multiple distributed system](#). It supports simultaneous access to multiple instances, and features low latency, high performance, and high reliability. Shared Block Storage applies to shared access scenarios for block storage devices under a shared everything architecture.
- Local disks are the disks attached to the physical servers (host machines) on which ECS instances are hosted. They are designed for business scenarios requiring high storage I/O performance and massive storage cost performance. Local disks provide local storage and access for instances, and features low latency, high random IOPS, high throughput, and cost-effective performance.

For more information about the performance of block-level storage products, see [#unique_6](#).

Block storage, OSS and NAS

Currently, Alibaba Cloud provides three types of data storage products: block storage, [Object Storage Service \(OSS\)](#), and [Network Attached Storage \(NAS\)](#).

the following three types of data storage products:

- Block storage: A high-performance and low-latency block-level storage device for ECS. It supports random reads and writes. You can format block storage and create

a file system on it as you would with a hard disk. , thereby enabling block storage to meet the data needs of numerous business scenarios.

- **OSS:** A huge storage space designed for storing massive amounts of unstructured data on the Internet, including images, audio, and video. You can access the data stored in OSS anytime, anywhere, by using APIs. Generally, OSS is applicable to business scenarios as website construction, separation of dynamic and static resources, and CDN acceleration.
- **NAS:** A storage space designed to store massive amounts of unstructured data that can be accessed by using standard file access protocols , such as the Network File System (NFS) protocol for Linux, and the Common Internet File System (CIFS) protocol for Windows. You can set permissions to allow different clients to access the same file at the same time. NAS is suitable for business scenarios such as file sharing across departments, non-linear file editing, high-performance computing, and containerization (such as with Docker).

2 Block storage performance

This topic describes the performance, metrics, and test methods of block storage devices such as disks, shared block storage, and local disks.

Performance metrics

The metrics for measuring storage performance include IOPS, throughput, and latency.

- IOPS

Input/output operations per second (IOPS) measures the number of write/read operations that can be performed each second. IOPS directly affects the performance of transaction-intensive applications, such as databases.

Common IOPS performance characteristics that are measured include sequential and random operations. The following table describes these performance characteristics.

Metric	Description	Data access method
Total IOPS	The total number of I/O operations per second	Randomly and sequentially access locations in storage devices
Random read IOPS	The average number of random read I/O operations per second	Randomly access locations in storage devices
Random write IOPS	The average number of random write I/O operations per second	
Sequential read IOPS	The average number of sequential read I/O operations per second	Sequentially access locations in storage devices
Sequential write IOPS	The average number of sequential write I/O operations per second	

- Throughput

Throughput measures the size of data transferred per second. Unit: Mbit/s.

Throughput directly affects the performance of applications that require a large number of read/write operations, such as Hadoop offline computing applications.

- Latency

Latency is the period of time required to process an I/O request. Unit: s, ms, or us. High latency may cause performance deterioration or errors in applications with high requirements for latency, such as databases.

- Therefore, we recommend that you use enhanced SSDs, standard SSDs, shared SSD block storage, or local SSDs for latency-sensitive applications.
- For applications that require high throughput but do not have high requirements for latency, such as Hadoop offline computing applications, we recommend that you use ECS instances equipped with local SATA HDDs in the d1 or d1ne instance family.

You cannot use capacity as a metric to measure the performance of block storage, but the performance of block storage devices varies with their capacities. The larger the capacity of a block storage device is, the more powerful data processing capability the device has. Block storage devices of the same type have the same I/O performance per unit capacity. However, the performance of each disk increases linearly with its capacity until the upper limit for a single disk is reached.



Note:

Block storage capacity is measured in binary units, such as kibibyte (KiB), mebibyte (MiB), gibibyte (GiB), or tebibyte (TiB). 1 KiB equals 1,024 bytes. 1 MiB equals 1,024 KiB. 1 GiB equals 1,024 MiB. 1 TiB equals 1,024 GiB.

Disks

The following table lists the performance and typical scenarios of different types of disks.

Performance category	Enhanced SSD			Standard SSD	Ultra disk	Disk***
Performance Level (PL)	PL3	PL2	PL1	N/A	N/A	N/A
Capacity of a single disk	1,261 GiB to 32,768 GiB	461 GiB to 32,768 GiB	20 GiB to 32,768 GiB	32,768 GiB	32,768 GiB	2,000 GiB
Maximum IOPS	1,000,000	100,000	50,000	25,000*	5,000	Several hundreds

Performance category	Enhanced SSD			Standard SSD	Ultra disk	Disk***
Maximum throughput	4,000 Mbit/s	750 Mbit/s	350 Mbit/s	300 Mbit/s*	140 Mbit/s	30 Mbit/s to 40 Mbit/s
Formulas for calculating the performance of a single disk**	IOPS = $\min\{1,800 + 50 \times \text{capacity, 1,000,000}\}$	IOPS = $\min\{1,800 + 50 \times \text{capacity, 100,000}\}$	IOPS = $\min\{1,800 + 50 \times \text{capacity, 50,000}\}$	IOPS = $\min\{1,800 + 30 \times \text{capacity, 25,000}\}$	IOPS = $\min\{1,800 + 8 \times \text{capacity, 5,000}\}$	N/A
	Throughput = $\min\{120 + 0.5 \times \text{capacity, 4,000}\}$ Mbit/s	Throughput = $\min\{120 + 0.5 \times \text{capacity, 750}\}$ Mbit/s	Throughput = $\min\{120 + 0.5 \times \text{capacity, 350}\}$ Mbit/s	Throughput = $\min\{120 + 0.5 \times \text{capacity, 300}\}$ Mbit/s	Throughput = $\min\{100 + 0.15 \times \text{capacity, 140}\}$ Mbit/s	N/A
Data reliability	99.999999%	99.999999%	99.999999%	99.999999%	99.999999%	99.999999%
API name	cloud_essd			cloud_ssd	cloud_efficiency	cloud
Scenarios	<ul style="list-style-type: none"> OLTP databases: relational databases such as MySQL, PostgreSQL, Oracle, and SQL Server databases NoSQL databases: non-relational databases such as MongoDB, HBase, and Cassandra databases Elasticsearch distributed logs : Elasticsearch, Logstash, and Kibana (ELK) log analysis 			Large and medium-sized development and testing applications that require high data reliability	<ul style="list-style-type: none"> Development and testing applications System disks 	<ul style="list-style-type: none"> Applications that are not frequently accessed or have low I/O loads Applications that require low costs and random read and write I/O operations

* The performance of standard SSDs varies according to the sizes of data blocks. Smaller data blocks result in lower throughput and higher IOPS, as described in the following table. A standard SSD can achieve the committed IOPS performance only when it is attached to an I/O optimized instance.

Data block size	Maximum IOPS	Throughput
4 KiB	25,000	100 Mbit/s
16 KiB	17,200	260 Mbit/s
32 KiB	9,600	300 Mbit/s
64 KiB	4,800	300 Mbit/s

** A standard SSD is used as an example to describe how to calculate the performance of a single disk:

- Maximum IOPS: The baseline is 1,800 IOPS. It increases by 30 IOPS per GiB of storage. The maximum IOPS is 25,000.
- Maximum throughput: The baseline is 120 Mbit/s. It increases by 0.5 Mbit/s per GiB of storage. The maximum throughput is 300 Mbit/s.

*** Basic disks are the previous generation of disks and are no longer available for purchase. You can choose from a variety of disks, such as ultra disks, standard SSDs, or enhanced SSDs based on your actual needs.

The random write latency varies according to disk categories. For the test commands, see [Performance tests](#):

- Enhanced ESSDs: 0.2 ms
- Standard SSDs: 0.5–2 ms
- Ultra disks: 1–3 ms
- Basic disks: 5–10 ms

Shared block storage

The following table lists the performance and typical scenarios of different types of shared block storage.

Performance category	Shared SSD block storage	Shared ultra block storage
Maximum capacity of a single disk	32,768 GiB	32,768 GiB

Performance category	Shared SSD block storage	Shared ultra block storage
Maximum capacity of a single instance	128 TiB	128 TiB
Maximum random read/write IOPS*	30,000	5,000
Maximum sequential read/write throughput*	512 Mbit/s	160 Mbit/s
Formulas for calculating performance of a single disk**	$\text{IOPS} = \min\{1,600 + 40 \times \text{capacity}, 30,000\}$	$\text{IOPS} = \min\{1,000 + 6 \times \text{capacity}, 5,000\}$
	$\text{Throughput} = \min\{100 + 0.5 \times \text{capacity}, 512\} \text{ Mbit/s}$	$\text{Throughput} = \min\{50 + 0.15 \times \text{capacity}, 160\} \text{ Mbit/s}$
Scenarios	<ul style="list-style-type: none"> • Oracle RAC • Failover cluster • High-availability servers 	High-availability architecture of servers

* The maximum IOPS and throughput listed in the preceding table are the maximum performance of a bare shared block storage device that is attached to two or more instances at the same time during stress tests.

** A shared SSD block storage device is used as an example to describe how to calculate the performance of a single disk:

- Maximum IOPS: The baseline is 1,600 IOPS. It increases by 40 IOPS per GiB of storage. The maximum IOPS is 30,000.
- Maximum throughput: The baseline is 100 Mbit/s. It increases by 0.5 Mbit/s per GiB of storage. The maximum throughput is 512 Mbit/s.

The latency varies according to the shared block storage categories as follows:

- Shared SSD block storage: 0.5–2 ms
- Shared ultra block storage: 1–3 ms

Local disks

For the performance of local disks, see [Local disks](#).

Performance tests

We recommend that you use the FIO tool to test block storage performance for both Linux and Windows instances.

**Note:**

You can also use other tools to test block storage performance, but the obtained benchmark may be different. For example, tools such as `dd`, `sysbench`, and `iometer` may be affected by test parameter settings and file systems, so the tools may not reflect the actual disk performance. The performance results in this topic are all from a Linux instance tested by using FIO, and are used as the performance reference for block storage products.

This topic takes a Linux instance as an example to describe how to use the FIO tool to test block storage performance. Before you test a block storage device, ensure that it is 4 KiB aligned.

**Warning:**

You can test bare block storage disks to obtain more accurate performance data, but such tests may damage the structures of the file systems. Make sure that you have backed up your data before the test. We recommend that you use a new ECS instance which contains no data to test block storage performance to avoid data loss.

- Run the following command to test random write IOPS:

```
fio - direct = 1 - iodepth = 128 - rw = randwrite - ioengine
= libaio - bs = 4k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Rand_Write
_Testing
```

- Run the following command to test random read IOPS:

```
fio - direct = 1 - iodepth = 128 - rw = randread - ioengine
= libaio - bs = 4k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Rand_Read_
_Testing
```

- Run the following command to test sequential write throughput:

```
fio - direct = 1 - iodepth = 64 - rw = write - ioengine =
libaio - bs = 1024k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Write_PPS_
_Testing
```

- Run the following command to test sequential read throughput:

```
fio - direct = 1 - iodepth = 64 - rw = read - ioengine = libaio
- bs = 1024k - size = 1G - numjobs = 1 - runtime = 1000 -
```

```
group_repo rting - filename = iotest - name = Read_PPS_T esting
```

- Run the following command to test random write latency:

```
fio - direct = 1 - iodepth = 1 - rw = randwrite - ioengine =  
libaio - bs = 4k - size = 1G - numjobs = 1 - group_repo rting  
- filename = iotest - name = Rand_Write _Latency_T esting
```

- Run the following command to test random read latency:

```
fio - direct = 1 - iodepth = 1 - rw = randread - ioengine =  
libaio - bs = 4k - size = 1G - numjobs = 1 - group_repo rting  
- filename = iotest - name = Rand_Read_ Latency_Te stingrandw  
rite - ioengine = libaio - bs = 4k - size = 1G - numjobs = 1  
- group_repo rting - filename = iotest - name = Rand_Write  
_Latency_T esting
```

The following table uses the command for testing random write IOPS as an example to describe the parameter settings in FIO commands.

Parameter setting	Description
-direct=1	Indicates that the I/O buffer is ignored during the test and data is written directly.
-iodepth=128	Indicates that when you use AIO, a maximum of 128 I/O requests can be made at the same time.
-rw=randwrite	Indicates that the read/write policy is random writes. Other valid values for rw: <ul style="list-style-type: none"> • randread (random reads) • read (sequential reads) • write (sequential writes) • randrw (random reads and writes)

Parameter setting	Description
-ioengine=libaio	<p>Indicates that libaio (the Linux-native asynchronous I/O facility) is used as the testing method. There are two ways for an application to use I/O:</p> <ul style="list-style-type: none"> · Synchronous <p>In synchronous I/O mode, a thread sends only one I/O request at a time and waits until the I/O request is completed. In this case, the iodepth value is always smaller than 1 for a single thread. You can increase the iodepth value by using multiple concurrent threads. The iodepth value will reach its upper limit when there are 16 to 32 threads running concurrently.</p> <ul style="list-style-type: none"> · Asynchronous <p>In asynchronous I/O mode, a thread uses libaio to send multiple I/O requests to the kernel at a time and waits until all these I/O requests are completed. Asynchronous I/O reduces the number of interactions and increases interaction efficiency.</p>
-bs=4k	<p>Indicates that the size of each block for one I/O is 4 KB (4k). If this parameter is not specified, default value 4k is used.</p> <p>When IOPS is tested, we recommend that you set <code>- bs</code> to a small value, such as 4k in this example command.</p> <p>When throughput is tested, we recommend that you set <code>- bs</code> to a large value, such as 1024k in other example commands.</p>
-size=1G	Indicates that the size of the testing file is 1 GiB.
-numjobs=1	Indicates that the number of testing threads is 1.
-runtime=1000	Indicates that the test duration is 1,000 seconds. If this parameter is not specified, the file of the size specified by <code>- size</code> is written in blocks of the size specified by <code>- bs</code> .
-group_reporting	Indicates that in the testing results, the statistics are displayed for a group of threads as a whole, not for each individual thread.

Parameter setting	Description
-filename=iotest	Indicates that the name of the test file is iotest. You can test bare block storage disks to obtain more accurate performance data, but such tests may damage the structures of the file systems. Make sure that you have backed up your data before the test.
-name=Rand_Write_Testing	Indicates that the name of the testing task is Rand_Write_Testing. You can specify any name you want.

3 Block storage

3.1 Cloud disk overview

Cloud disks are block-level data storage products provided by Alibaba Cloud for ECS. They feature low latency and high durability and reliability.

Cloud disks provided by Alibaba Cloud support data redundancy through the replication of data three times by default, and the storage of the data copies across different servers, to provide 99.9999999% data reliability for ECS instances. This means that if service disruptions occur (for example, due to hardware failure), data within the target zone is copied to an unaffected cloud disk to help ensure data availability. You can partition and format cloud disks attached to an ECS instance, and create a file system and store data on the cloud disk. For more information, see [Triplicate technology](#).



Note:

You can only attach a cloud disk to one ECS instance that is in the same zone of the same region as the cloud disk.

Cloud disk types


The following tables classify cloud disk types by performance and function.

- By performance

Type	Description
ESSD cloud disk	An Enhanced SSD (ESSD) cloud disk is an ultra-high performance cloud disk provided by Alibaba Cloud. ESSD cloud disks are based on the next-generation distributed block storage architecture, a 25GE network, and remote direct memory access (RDMA) technology. They deliver a maximum random IOPS of 1 million per disk and low latency. For more information, see #unique_11 .
SSD cloud disk	High-performance disks with stable and high random I/O performance, and high data reliability.
Ultra cloud disk	Cloud disks with high cost performance, medium random I/O performance, and high data reliability.

Type	Description
Basic cloud disk	Cloud disks with high data reliability and standard random I/O performance.

· By function

Type	Description
System disk	<p>A system disk has the same life cycle as the ECS instance to which it is mounted, and is created and released along with the ECS instance. Shared access to the system disk is not allowed. The available size of a system disk varies according to the operating system as follows:</p> <ul style="list-style-type: none"> - Linux (excluding CoreOS) and FreeBSD: 20–500 GiB - CoreOS: 30–500 GiB - Windows: 40–500 GiB
Data disk	<p>Data disks can be created separately or at the same time as ECS instances.</p> <ul style="list-style-type: none"> - Data disks created along with ECS instances have the same life cycle as the corresponding instance, and are released along with the instance. - Data disks created separately can be released separately or at the same time as the corresponding ECS instance. <p>Shared access to a data disk is not allowed. The performance of data disks depends on the cloud disk type. For more information, see #unique_6.</p> <div>  <p>Note: When used as data disks, up to 16 cloud disks can be attached to one ECS instance.</p> </div>

Performance

For the performance of each type of cloud disk, see [#unique_6](#).

Billing method

- Cloud disks created along with Subscription instances or created separately for Subscription instances are billed according to the selected billing cycle and are charged before you use the service. For more information, see [#unique_13](#).

- Cloud disks created separately or created along with Pay-As-You-Go instances are billed based on the amount of resources you actually use. For more information, see [Pay-As-You-Go](#).

You can change the billing method of a cloud disk, as described in the following table.

Conversion of billing methods	Related Feature	Effective time	Suitable for
From Subscription to Pay-As-You-Go	#unique_15	Effective from the next billing cycle	Subscription cloud disks mounted to Subscription instances. The billing method of the system disk cannot be changed.
	Switch from Subscription to Pay-As-You-Go billing	Effective immediately	The system disk of Subscription instances and data disks created along with Subscription instances.
From Pay-As-You-Go to Subscription	Upgrade configurations	Effective immediately	Pay-As-You-Go data disks mounted to Subscription instances. The billing method of the system disk cannot be changed.
	#unique_17		System disks and data disks mounted to Pay-As-You-Go instances.

Related operations

You can perform the following operations on a cloud disk:



Note:

Some of the following operations are dependent on whether the cloud disk functions as a system disk or a data disk.

- [Attach the cloud disk to an ECS instance](#) in the ECS console and then connect to the ECS instance to [partition and format the cloud disk](#) if you [created the cloud disk separately as a data disk](#).
- [Encrypt the cloud disk](#).
- [Resize the system disk](#).
- [Resize the data disk](#).
- [Replace the system disk](#) to change the operating system of the system disk.
- [Manually create snapshots of the cloud disk](#) to back up data on the cloud disk, or [apply an automatic snapshot policy to it](#) to automatically create snapshots.
- [Create a custom image by using the snapshots of the system disk](#).
- [Roll back the cloud disk by using a snapshot](#).
- [Reinitialize the cloud disk](#).
- [Detach](#) and then [release the cloud disk](#) if you no longer need it.
- [Change the billing method of a Subscription cloud disk](#), and then [detach](#) and [release the cloud disk](#) if you no longer need it.

3.2 ESSD cloud disk

This topic provides general information about Alibaba Cloud enhanced SSD (ESSD) cloud disks, such as performance levels, scenarios, and performance specifications. ESSD cloud disks are based on a 25 GE network and remote direct memory access (RDMA) technology. They deliver a maximum random IOPS of 1 million per disk and low one-way latency.

Scenarios

ESSD cloud disks apply to the following scenarios where short latency and high throughput are required:

- Online transaction processing (OLTP) databases (relational databases), for example , MySQL, PostgreSQL, Oracle, and SQL Server.
- NoSQL databases (non-relational databases), for example, MongoDB, HBase, and Cassandra.
- Real-time analysis applications, for example, Elasticsearch and ELK (Elasticsearch , Logstash, and Kibana) logging tools.

Specifications

The API value of ESSD cloud disks is `cloud_essd`. ESSD cloud disks are available in three performance levels (PLs). For more information, see [#unique_11/unique_11_Connect_42_ESSDperfor](#)

The following table describes the specifications of ESSD cloud disks of different performance levels.

ESSD cloud disk attribute	PL3	PL2	PL1
Capacity (GiB)	1,261-32,768	461-32,768	20-32,768
Data reliability	99.9999999%	99.9999999%	99.9999999%
Maximum IOPS of a single cloud disk	1,000,000	100,000	50,000
Maximum throughput of a single cloud disk (Mbit/s)	4,000	750	350
Performance formula for a single cloud disk	$\text{IOPS} = \min\{1800 + 50 * \text{capacity}, 1000000\}$	$\text{IOPS} = \min\{1800 + 50 * \text{capacity}, 100000\}$	$\text{IOPS} = \min\{1800 + 50 * \text{capacity}, 50000\}$
	$\text{Throughput} = \min\{120 + 0.5 * \text{capacity}, 4000\}$ Mbit/s	$\text{Throughput} = \min\{120 + 0.5 * \text{capacity}, 750\}$ Mbit/s	$\text{Throughput} = \min\{120 + 0.5 * \text{capacity}, 350\}$ Mbit/s

Billing

ESSD cloud disks support Pay-As-You-Go and Subscription billing methods. For more information, see [Create a Pay-As-You-Go cloud disk](#) and [#unique_25](#).

For more information, see [ECS products](#).

PL selection for ESSD cloud disks

ESSD cloud disks of the three PLs have the same single-disk capacity that ranges from 20 GiB to 32,768 GiB. The key difference lies in the maximum performance of a single disk. You can use ESSD cloud disks as system disks or data disks according to your specific needs. The scenarios where ESSD cloud disks are used as data disks are described as follows:

- ESSD PL1 cloud disks are suitable for scenarios where moderate concurrent IOPS and stable read/write latency are required.

Such scenarios include small and medium MySQL and SQLServer databases, ELK clusters, enterprise-level commercial software (such as SAP and Oracle), and container applications. In these scenarios, you can use ESSD PL1 cloud disks to replace SSD or Ultra cloud disks.

- ESSD PL2 cloud disks are suitable for scenarios where high concurrent IOPS and stable read/write latency are required.

Such scenarios include medium relational databases and NoSQL databases, ELK clusters, and enterprise-level commercial software (such as SAP and Oracle). In these scenarios, you can use ESSD PL2 cloud disks to replace SSD cloud disks or local SSD instance types (i1, i2, and i2g).

- ESSD PL3 cloud disks are suitable for scenarios where ultra-high concurrent IOPS and ultra-stable read/write latency are required.

Such scenarios include medium and large relational databases and NoSQL databases, and large enterprise-level commercial software (such as SAP and Oracle). In these scenarios, you can use ESSD PL3 cloud disks to replace local SSD instance types (specifically the instance types i1, i2, and i2g, with at least 16 vCPUs).

Disk capacity and performance

For ESSD cloud disks, larger capacity positively corresponds with higher data-processing performance. The capacity, IOPS, and throughput provided by ESSD cloud disks vary depending on their performance levels. The specific capacity and performance are detailed in the following table.

Performance level	Capacity (GiB)	Maximum IOPS	Maximum throughput (Mbit/s)
PL1	20-32,768	50,000	350
PL2	461-32,768	100,000	750
PL3	1,261-32,768	1,000,000	4,000

- Example 1: If you select 20 GiB memory when you create an ESSD cloud disk on the ECS console, you can only select a PL1 ESSD cloud disk whose maximum IOPS is 50,000.

- Example 2: If you select 32,000 GiB memory, you can select an ESSD cloud disk of any performance level. The maximum IOPS is 50,000, 100,000, and 1000,000 in ascending order of the performance level.

If you require higher performance, [resize the ESSD cloud disk](#) or [change the performance level of the ESSD cloud disk](#).


Related operations

- [Create a Pay-As-You-Go cloud disk](#).
- [Create a Subscription cloud disk](#).
- [Create a cloud disk by using a snapshot](#).
- [#unique_18](#).
- [#unique_29](#).
- [Resize a cloud disk](#).
- [#unique_27](#).

3.3 Triplicate technology

The Alibaba Cloud Distributed File System provides stable and efficient data access and reliability for ECS. Triplicate technology, that is, the process of making and distributing three copies of data, is the principle concept implemented in the Alibaba Cloud Distributed File System.

When you perform read and write operations on cloud disks, the operations are translated into the corresponding processes on the files stored in Alibaba Cloud data storage system. The Distributed File System of Alibaba Cloud uses a flat design in which a linear address space is divided into slices, also called chunks. Each chunk has three copies stored on different server nodes on different racks. This guarantees data reliability.



How triplicate technology works


Triplicate technology involves three key components: Master, Chunk Server, and Client. To demonstrate how triplicate technology works, in this example, the write operation of an ECS user undergoes several conversions before being executed by the Client. The process is as follows:

1. The Client determines the location of a chunk corresponding to a write operation.
2. The Client sends a request to the Master to query the storage locations (that is, the Chunk Servers) of the three copies of the chunk.
3. The Client sends write requests to the corresponding three Chunk Servers according to the results returned from the Master.
4. The Client returns a message that indicates whether the operation was successful.

This strategy guarantees that all the copies of a chunk are distributed on different Chunk Servers on different racks, effectively reducing the potential of total data loss caused by failure of a Chunk Server or a rack.

Data protection

If a system failure occurs because of a corrupted node or hard drive failure, some chunks may lose one or more of the three valid chunk copies associated with them. If this occurs and triplicate technology is enabled, the Master replicates data between Chunk Servers to replace the missing chunk copies across different nodes.



To summarize, all your operations (additions, modifications, or deletions) on cloud disk data are synchronized to the three chunk copies at the bottom layer. This mode ensures the reliability and consistency of your data.

Furthermore, we recommend you implement appropriate backup strategies, [snapshots](#), and other precautionary actions to restore and protect your data and guarantee its availability against other types of failures, such as viruses, human error, or malicious activity on your account. No single technology can solve all the problems, so you must choose appropriate data protection measures to establish a solid defense line for your valuable business data.

3.4 ECS disk encryption

ECS disks in this article refer to cloud disks and Shared Block Storage devices. They are referred to as ECS disks in the following contents, unless otherwise specified.

What is ECS disk encryption?

The ECS disk encryption feature allows you to encrypt new ECS disks so that you can meet encryption needs for scenarios such as certification requirements and business security. The ECS disk encryption feature means you do not have to create, maintain, or protect your own key management infrastructure, nor change any of your existing applications or maintenance processes. In addition, no extra encryption or decryption operations are required, making ECS disk encryption operations invisible to your applications or other operations.

Encryption and decryption processes hardly degrade ECS disk performance. For information on the performance testing method, see [storage parameters and performance test](#).

After an encrypted ECS disk is created and attached to an ECS instance, you can encrypt data that is:

- Stored directly on the ECS disk.
- Transmitted between the ECS disk and the instance. However, data in the instance operating system is not encrypted.
- Created from the encrypted ECS disk, such as snapshots. These snapshots are called encrypted snapshots.

Encryption and decryption are performed on the host that runs the ECS instance, so the data transmitted from the ECS instance to the cloud disk is encrypted.

ECS disk encryption supports all available cloud disks (Basic Cloud Disks, Ultra Cloud Disks, SSD Cloud Disks, and ESSDs) and shared block storage (Ultra Shared Block Storage and SSD Shared Block Storage).

ECS disk encryption supports all available instance types and is supported in all regions.

ECS disk encryption dependencies

ECS disk encryption is dependent on the Key Management Service (KMS), which must be in the same region. However, you do not need to perform any additional operations in the KMS console to activate ECS disk encryption.

The first time you use the ECS disk encryption function (such as when you are creating ECS instances or ECS disks), you must first authorize and activate KMS. Otherwise, you cannot create encrypted ECS disks or instances with encrypted disks.

If you use an API or the CLI to use the ECS disk encryption function, such as [CreateInstance](#) or [CreateDisk](#), you must first activate KMS on the Alibaba Cloud console.

The first time you encrypt a disk in a target region, Alibaba Cloud automatically creates a Customer Master Key (CMK) in the KMS region, exclusively for ECS. The CMK cannot be deleted. You can query the CMK in the KMS console.

Key management for ECS disk encryption

ECS disk encryption handles key management for you. Each new ECS disk is encrypted by using a unique 256-bit key (derived from the CMK). This key is also associated with all snapshots created from this ECS disk and any ECS disks subsequently created from these snapshots. These keys are protected by the key management infrastructure of Alibaba Cloud provided by KMS. This approach implements strong logical and physical security controls to prevent unauthoriz

ed access. Your data and the associated keys are encrypted based on the industry standard AES-256 algorithm.

You cannot change the CMK associated with encrypted ECS disks and snapshots.

The key management infrastructure of Alibaba Cloud conforms to the recommendations in (NIST) 800-57 and uses cryptographic algorithms that comply with the (FIPS) 140-2 standard.

Each Alibaba Cloud account has a unique CMK in each region. This key is separate from the data and is stored in a system protected by strict physical and logical security controls. Each encrypted disk and its snapshots use an encryption key that is unique to the specific disk. The encryption key is created from and encrypted by the CMK for the current user in the current region. The disk encryption key is only used in the memory of the host that runs your ECS instance. The key is never stored in plaintext in any permanent storage media (such as an ECS disk).

Fees

The ECS disk encryption features incur no additional fees.

The CMK that ECS creates for you in each region is a service key. It does not consume your master key quota in a given region, meaning no additional fees are incurred.



Note:

No additional fees are charged for any read/write operations on a disk, such as mounting/umounting, partitioning, and formatting. However, if you perform operations on a disk in the ECS console or by using APIs, KMS APIs are called and such calls consume the KMS API quota in the current region.

These operations include:

- Creating encrypted disks by calling [CreateInstance](#) or [CreateDisk](#).
- Attaching an encrypted disk to an instance by calling [AttachDisk](#).
- Detaching an encrypted disk from an instance by calling [DetachDisk](#).
- Creating a snapshot by calling [CreateSnapshot](#).
- Restoring a disk by calling [ResetDisk](#).
- Re-initializing a disk by calling [ReInitDisk](#).

Create an encrypted ECS disk

Currently, only cloud disks can be encrypted. You can create an encrypted cloud disk in the following ways:

- Create a cloud disk as a data disk when creating an ECS instance or :
 - Check Encrypted to create a encrypted blank cloud disk.
 - Select an encrypted screenshot to create a cloud disk.
- When using APIs or the CLI:
 - Set the parameter `DataDisk . n . Encrypted` ([CreateInstance](#)) or `Encrypted` ([CreateDisk](#)) to `true` .
 - Specify the `SnapshotId` parameter of the encrypted snapshot in `CreateInstance` or `CreateDisk`.

Convert unencrypted data to encrypted data

You cannot directly convert an unencrypted disk to an encrypted disk, or perform the converse operation.

You cannot convert a snapshot created from an unencrypted disk to an encrypted snapshot, or perform the converse operation.

Therefore, if you must switch the existing data from status unencrypted to encrypted , we recommend that you use the `rsync` command in a Linux instance or the `robocopy` command in a Windows instance to copy data from an unencrypted disk to a (new) encrypted disk.

Therefore, if you must switch the existing data from status encrypted to unencrypted, we recommend that you use the `rsync` command in a Linux instance or the `robocopy` command in a Windows instance to copy data from an encrypted disk to a (new) unencrypted disk .

Limits

ECS disk encryption has the following limits:

- You can only encrypt ECS disks, not local disks or ephemeral disks.
- You can only encrypt data disks, not system disks.
- You cannot directly convert existing unencrypted disks into encrypted disks.
- You cannot convert encrypted disks into unencrypted disks.
- You cannot convert unencrypted snapshots to encrypted snapshots.

- You cannot convert encrypted snapshots to unencrypted snapshots.
- You cannot share images created from encrypted snapshots.
- You cannot copy images created from encrypted snapshots across regions.
- You cannot export images created from encrypted snapshots.
- You cannot define CMKs for each region. They are generated by the system.
- The ECS system creates CMKs for each region. You cannot delete these keys, and you do not incur fees from them.
- After a cloud disk is encrypted, you cannot change the CMK used for encryption and decryption.

3.5 Create a cloud disk

3.5.1 Create a Pay-As-You-Go cloud disk

This topic describes how to create a Pay-As-You-Go cloud disk that contains no data by using the ECS console. Adding a Pay-As-You-Go cloud disk can increase the storage space of an instance.

Limits

Before you create a Pay-As-You-Go cloud disk, note the following:

- You cannot create a separate system disk in the ECS console.
- Across all regions, the number of Pay-As-You-Go data disks that you can create cannot be more than five times the number of Pay-As-You-Go instances under your account. For more information, see [#unique_41](#).
- You cannot merge multiple cloud disks by formatting them because they are independent of each other. Therefore, we recommend that you estimate the number and size of cloud disks required before you create them.
- We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes for multiple cloud disks. This is because a snapshot can only back up data of a single cloud disk. If you use LVM, data discrepancies will occur when you roll back these cloud disks.

Billing

The cloud disk created in this topic is billed as Pay-As-You-Go. For more information, see [#unique_14](#).

If you need to convert the billing method of a Pay-As-You-Go cloud disk to Subscription, see [#unique_29](#).

**Note:**

If you need to create a Subscription cloud disk, use either of the following methods:

- [Create a Subscription cloud disk and mount it to](#) a Subscription instance.
- When you [create an ECS instance](#), set its billing method to Subscription. In this case, all the cloud disks attached to the instance use the Subscription billing method.

You cannot [unmount](#) or [release](#) a Subscription cloud disk on demand. The Subscription cloud disk is released along with the corresponding instance.

Procedure

1. In the upper-right corner of the Disks page, click Create Disk.
2. Select the region and zone that are the same as the target instance.
3. Specify the disk type, capacity, disk encryption mode, and quantity.

Disk

SSD 20 GIB 2400 IOPS [Create from Snapshot](#) ☐ Disk Encryption

[Click here>](#) for guidelines on how to select an appropriate disk for your scenario.
[Learn how to create a subscription disk>](#)

Quantity: 1 Unit
You have created 0 disks. You can create 25 more disks.

Disk Name:
The name must be 2 to 128 characters in length and can contain letters, Chinese characters, digits, and the following special characters: -_
The name must start with a letter or Chinese character.

Description:

**Note:**

- To check whether a cloud disk needs to be encrypted, see [#unique_20](#).
- You can also [create a cloud disk by using a snapshot](#).
- For information about how to select between ESSD cloud disks of different specifications, see [ESSD cloud disk](#).

4. Select the *ECS Service Level Agreement* check box.
5. Confirm the cloud disk configuration and fee.
6. Click Preview Order.
7. In the displayed dialog box, click Create.

8. After the cloud disk is created, return to the Disks page and refresh the disk list.

The Status of the cloud disk is displayed as Unattached.

What to do next

[Attach a cloud disk.](#)

Related APIs

- [CreateDisk](#): Creates a Pay-As-You-Go or Subscription cloud disk.
- [RunInstances](#) or [CreateInstance](#): Creates a data disk along with an ECS instance.

References

- [Create a cloud disk by using a snapshot](#)
- [Create an instance by using the wizard](#)

3.5.2 Create a Subscription cloud disk

This topic describes how to create a Subscription cloud disk that contains no data by using the ECS console. Adding a Subscription cloud disk can increase the storage space of an instance.

Limits

Before you create a Subscription Cloud disk, note the following:

- You cannot merge multiple cloud disks by formatting them because they are independent of each other. Therefore, we recommend that you estimate the number and size of cloud disks required before you create them.
- We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes for multiple cloud disks. This is because a snapshot can only back up data of a single cloud disk. If you use LVM, data discrepancies will occur when you roll back these cloud disks.
- You cannot [detach a Subscription cloud disk](#) on demand. The Subscription cloud disk is released along with the instance. If you need to release a Subscription cloud disk, convert it to a Pay-As-You-Go cloud disk, and then detach and release it.

Billing details

The cloud disk created in this topic is billed as Subscription. For more information, see [Subscription](#).

If you need to convert a Subscription cloud disk to a Pay-As-You-Go cloud disk, see [#unique_29](#).

**Note:**

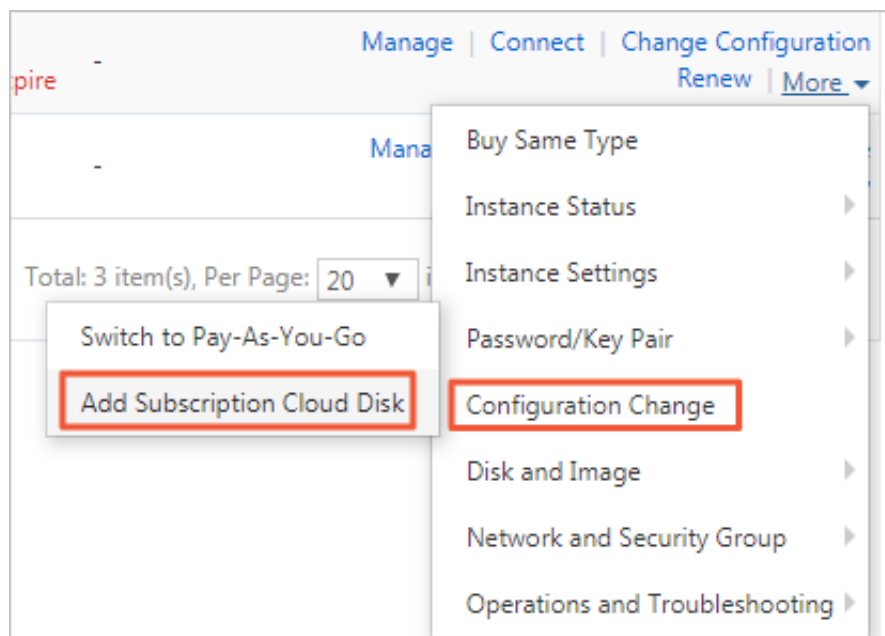
If you need to create a Pay-As-You-Go cloud disk, use either of the following methods:

- [Create a Pay-As-You-Go cloud disk](#) on the Disks page of the ECS console.
- When you [create an ECS instance](#), set its billing method to Pay-As-You-Go. In this case, all the cloud disks attached to the instance use the Pay-As-You-Go billing method.

Procedure

To create a Subscription cloud disk, follow these steps:

1. On the Instances page, find the target Subscription instance. In the Actions column, choose More > Configuration Change > Add Subscription Cloud Disk.



2. In the Disk area of the displayed page, complete the following configurations:

- Cloud disk type: Select a [cloud disk type](#) in the drop-down list. For information about how to select between ESSD cloud disks of various specifications, see [ESSD cloud disk](#).
- Cloud disk capacity: Enter a cloud disk capacity in the text box. The disk capacity ranges from 20 GiB to 32,768 GiB.
- Cloud disk encryption: If [#unique_50](#) is needed, select the Disk Encryption check box.
- Quantity: Enter the number of cloud disks you want to purchase.



Note:

You can create up to 16 data disks for a single instance, including cloud disks and Shared Block Storage devices.

- Disk Name: optional. The name for the cloud disk.
- Description: optional. The description of the cloud disk.
- If you need to [create a cloud disk from a snapshot](#), click Create from snapshot.

3. Select the *ECS Service Level Agreement* check box.
4. Click Preview Order.
5. In the displayed dialog box, click Create.
6. Select a payment method and click Confirm to Pay to complete the creation.
7. Click Console to return to the Instances page and click the instance for which the Subscription cloud disk has just been added.
8. Click Disks to view the newly added Subscription cloud disk.

The Subscription cloud disk is mounted to the instance and is in the Running state.

What to do next

Format and partition the cloud disk. For more information, see [#unique_19](#) or [#unique_51](#).

Related APIs

- [CreateDisk](#): Creates a Pay-As-You-Go or Subscription cloud disk.
- [RunInstances](#) or [CreateInstance](#): Creates a data disk along with an ECS instance.

References

- [Create a Pay-As-You-Go cloud disk](#)
- [Create a cloud disk by using a snapshot](#)

3.5.3 Create a cloud disk by using a snapshot

This topic describes how to create a cloud disk by using a snapshot of a system disk or data disk in the ECS console. After you create a cloud disk, you can attach it to any ECS instance in the same zone of the same region. Note that a cloud disk created by a snapshot can only be used as a data disk.

Scenarios

We recommend that you create a cloud disk by using a snapshot in the following scenarios:

- If you need to obtain data from a snapshot, but you do not want to [roll back a cloud disk](#).
- If your instance encounters a system disk failure, you can use an existing system disk snapshot to create a cloud disk. Then, you can attach the cloud disk as a data disk to a healthy ECS instance so that you can continue to read the data of the system disk.

Limits

- By default, a cloud disk created by using a snapshot uses the [#unique_14](#) billing method, and can only be used as a data disk. However, you can change the billing method of the cloud disk. For more information, see [What to do next](#).
- When you access a cloud disk created by a snapshot the first time, the performance is reduced because it takes some time for ECS to read data from OSS and write data to the cloud disk. Therefore, we recommend that you do not use the cloud disk until it has read from all data blocks. For more information, see [#unique_54](#)

- Across all regions, the number of Pay-As-You-Go data disks that you can create cannot be more than five times the number of Pay-As-You-Go instances under your account. For more information, see [#unique_41](#).
- You cannot merge multiple cloud disks by formatting them because they are independent of each other. Therefore, we recommend that you estimate the number and size of cloud disks required before you create them.
- We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes for multiple cloud disks. This is because a snapshot can only back up data of a single cloud disk. If you use LVM, data discrepancies will occur when you roll back these cloud disks.

Prerequisites

A system disk snapshot or a data disk snapshot is created and its ID is obtained. For more information, see [Create a snapshot](#).

Procedure

You can create a cloud disk through the ECS console or by calling [CreateDisk](#). To create a cloud disk through the ECS console, follow these steps:

1. In the upper-right corner of the Disks page, click Create Disk.
2. Select the target region and zone.



Note:

If you want to attach the cloud disk to an ECS instance, they must be in the same zone of the same region.

3. Complete the following settings:
 - a. Select a disk type: Ultra Disk, ESSD Cloud Disk, or SSD Cloud Disk.
 - b. Click Create from snapshot, and then select a snapshot.
 - c. Set the capacity of the cloud disk. Note that it must be greater than 20 GiB and less than 32768 GiB.



Note:

- If you do not set the capacity of the cloud disk, the system automatically sets a capacity for the new cloud disk according to the cloud disk corresponding to the snapshot.

- If you set a capacity greater than the snapshot size, you must partition the cloud disk before you can use the full capacity.
- If the snapshot size is less than 2048 GiB and the cloud disk to be created is greater than 2048 GiB, you must ensure that the cloud disk corresponding to the snapshot uses GPT partitions (Globally Unique Identifier Partition Table). Otherwise, we recommend that you set the cloud disk size to less than 2048 GiB to avoid data loss that may occur in partitioning. For more information, see [#unique_56](#).

- d. Set the Quantity of cloud disks that you want to create.
 - e. Read and confirm that you agree with the ECS Service Level Agreement by selecting the check box.
4. Confirm your settings and the estimated cost displayed.
 5. Click Preview, and then click Create in the displayed dialog box.

After you complete the payment, go back to the Disks page and refresh the disk list. The new disk is displayed in the Unmounted state.

What to do next

After you create a cloud disk, you can:

- [Attach the cloud disk](#).
- Change the billing method of the cloud disk from Pay-As-You-Go to Subscription.
 - If the cloud disk is attached to a Subscription instance, see [#unique_57](#).
 - If the cloud disk is attached to a Pay-As-You-Go instance, see [#unique_17](#).

You can also create a cloud disk by using a system disk snapshot or a data disk snapshot when you [create an ECS instance](#). In this case, the cloud disk uses the same billing method as the ECS instance.

3.6 Attach a cloud disk

This topic describes how to attach a cloud disk. You can create a cloud disk and attach it to an ECS instance as a data disk.

Limits

Before you attach a cloud disk to an ECS instance, consider the following:

- You can attach a cloud disk as a data disk only. You cannot attach a cloud disk as a system disk.
- To attach a cloud disk to an ECS instance, the instance must meet the following requirements:
 - The instance must be in the Running or Stopped state, but not in the Locked state.
 - The instance must not have any overdue payments.
- The disk to be attached must be in the Unmounted state.
- The cloud disk and the ECS instance must be in the same region and the same zone.
- Up to 16 cloud disks can be attached to an ECS instance to work as data disks. One cloud disk cannot be attached to multiple instances simultaneously.
- If a cloud disk is created independently on the Disks page in the [ECS console](#), it can be attached to any ECS instance in the same region and the same zone, regardless of the billing method of the instance.

Prerequisites

You have created an ECS instance and a cloud disk in the same region and zone. For more information, see [#unique_42](#) and [Create a Pay-As-You-Go cloud disk](#).

Attach a cloud disk on the Instances page

If you want to attach multiple cloud disks to one ECS instance, we recommend that you do so on the Instances page. To attach cloud disks to a specified ECS instance, follow these steps:

1. Find the target ECS instance and click its ID to go to the Instance Details page.
2. In the left-side navigation pane, click Disks. Then on the displayed page, click Mount in the upper-right corner.
3. In the dialog box, complete the following configurations:
 - Target Disk: Select a cloud disk in the Unmounted state in the same zone.
 - Release Disk with Instance: If you select this option, the disk is released when you release the corresponding instance.
 - Delete Automatic Snapshots While Releasing Disk: If you select this option, all the automatic snapshots of the selected disk are deleted when you release the

disk. However, all the manual snapshots are retained. To keep a complete data backup history, we recommend that you do not select this option.

Click OK, and then Mount.

4. Refresh the Disks page.

When the cloud disk is in the In Use state, the attachment is successful.

5. According to the content of the cloud disk and the operating system of the ECS instance, perform appropriate operations to make the disk ready for use. The following table details the different follow-up operations.

Disk content	Operating system of the ECS instance	Follow-up operations
A new empty cloud disk	Linux	#unique_19 . If the cloud disk is larger than 2 TiB, see #unique_56 .
	Windows	#unique_51 . If the cloud disk is larger than 2 TiB, see #unique_56 .
A cloud disk created from a snapshot	Linux	Connect to the Linux instance and run the <code>mount < partition > < mount point ></code> command to mount the partitions to the target mount points make the disk ready for use.
	Windows	No follow-up operations are required. The cloud disk is ready for use.

Attach a cloud disk on the Disks page

If you want to attach multiple cloud disks to different ECS instances, we recommend that you do so on the Disks page. To attach a cloud disk to an ECS instance, follow these steps:

1. Find a cloud disk in the Unmounted state. Then in the Actions column, select More > Mount.

2. In the dialog box, complete the following configurations:

- **Target Instance:** Select an ECS instance in the same zone.
- **Release Disk with Instance:** If you select this option, the disk is released when you release the corresponding instance.
- **Delete Automatic Snapshots While Releasing Disk:** If you select this option, all the automatic snapshots of the selected disk are deleted when you release the disk. However, all the manual snapshots are retained. To keep a complete data backup history, we recommend that you do not select this option.

Click Mount.

3. Refresh the Disks page.

When the cloud disk is in the In Use state, the attachment is successful.

4. According to the content of the cloud disk and the operating system of the ECS instance, perform appropriate operations to make the disk ready for use. The following table details different follow-up operations.

Disk content	Operating system of the ECS instance	Follow-up operations
A new empty cloud disk	Linux	#unique_19 . If the cloud disk is larger than 2 TiB, see #unique_56 .
	Windows	#unique_51 . If the cloud disk is larger than 2 TiB, see #unique_56 .
A cloud disk created from a snapshot	Linux	Connect to the Linux instance and run the <code>mount</code> command to mount the partitions to make the disk ready for use.
	Windows	No follow-up operations are required. The cloud disk is ready for use.

What to do next

After a cloud disk is attached to an ECS instance, you can perform any of the following operations as needed:

- You can [reinitialize a cloud disk](#) to restore its initial status.
- You can increase the size of a cloud disk. For more information, see [Extend the file system of a Linux data disk](#) or [Extend a Windows file system](#).
- You can [create a snapshot](#) of a cloud disk to back up its data. Alternatively, you can [use automatic snapshot policies](#) to create automatic snapshots.
- You can use a snapshot to [roll back a cloud disk](#) to restore the cloud disk to a previous state.
- You can [detach a cloud disk](#) and [release a cloud disk](#) when you no longer require the cloud disk, thus reducing the costs.

Related APIs

[AttachDisk](#)

3.7 Format a data disk

3.7.1 Format a data disk for a Windows instance

This topic describes how to format a new data disk for a Windows instance, create a single partition in the data disk, and mount a file system on the data disk. You can use the procedures provided in this topic to create multiple partitions as needed.



Warning:

- Disk partitioning and formatting are high-risk operations. Exercise caution when performing these operations. The following procedure uses a newly purchased data disk as an example. If you partition or format an existing data disk, [create a snapshot of the data disk](#) to avoid data loss.
- We strongly recommend that you do not partition the system disk of an ECS instance. Failure to comply can result in unknown risks, such as system failure and data loss. You can only extend a partition of, or add a partition to, a system disk after you resize the system disk. For more information, see [#unique_64](#).

Procedure

The following procedure applies only to data disks less than 2 TiB. If your data disk is greater than 2 TiB, see [Partition and format a data disk greater than 2 TiB](#).

In this example, a new 20 GiB data disk is partitioned and formatted. An instance running Windows Server 2012 R2 64-bit is used.

1. [Connect to the instance](#).
2. On the Windows Server desktop, right-click the Start icon, and then select Disk management.
3. Locate the data disk that is not formatted (Disk 2 in this example) and check that the disk is in the Offline state.
4. Right-click in the blank space in the Disk 2 area, and then select Online.

After Disk 2 goes online, it enters the Not Initialized state.

5. Right-click in the blank space in the Disk 2 area, and then select Initialize Disk.
6. In the Initialize Disk dialog box, select Disk 2, and then select either of the following partitioning methods:

- MBR

This method only supports data disks less than 2 TiB and can divide a disk into a maximum of four primary partitions. If you want to divide a disk into more than four partitions, you need to use a primary partition as an extended partition and create logical partitions in it.

- GPT

Although GPT is a newer partitioning method, it is not recognized by earlier versions of Windows. The data disk size that GPT can support is determined by the operating system and the file system. In Windows, GPT supports up to 128 primary partitions.


In this example, select MBR, and then click OK.

7. In the Disk Management window, right-click the Unallocated area of Disk 2, and then select New Simple Volume.
8. In the New Simple Volume Wizard dialog box, complete the following operations:
 - a. Click Next.
 - b. Specify the size of the simple volume. If you only need one primary partition, use the default value. Then, click Next.
 - c. Select a drive letter (in this example, F), and then click Next.
 - d. Select the required formatting settings (including the file system, allocation unit size, and volume label), and confirm whether to enable Quick Formatting and

File and Folder Compression. In this example, use the default settings, and then click Next.

- e. After the simple volume is created, click Finish to close the New Simple Volume Wizard.

After Disk 2 is partitioned and formatted, its status is displayed in the Disk Management window, as shown in the following figure.



In This PC, you can view a new drive named New Volume (F:). The data disk is ready to use.

3.7.2 Format a data disk of a Linux instance

This topic describes how to format a data disk of a Linux instance. A newly created or purchased data disk cannot be used unless you format it, create one or more partitions in it, and mount a file system on it.



Warning:

- Disk partitioning and formatting are high-risk operations. Exercise caution when performing these operations. The following procedure uses a newly purchased data disk as an example. If you partition or format an existing data disk, make sure that you have [created a snapshot of the data disk](#) to avoid data loss.
- Do not partition the system disk of an ECS instance. Failure to comply can result in unknown risks, such as system failure and data loss. You can only extend a partition of, or add a partition to a system disk after you resize the system disk. For more information, see [#unique_67](#).



Note:

The following procedure applies only to data disks less than 2 TiB. If your data disk is greater than 2 TiB, see [Partition and format data disk greater than 2 TiB](#).

Prerequisites

- The ECS instance is **attached with a data disk** that was **created separately**. You do not need to perform this operation for data disks created along with ECS instances.
- The device name of the data disk is obtained.

You can obtain the device name of the data disk by choosing ECS Console > Block Storage > Disks > (Disk ID specific) More > Modify Attributes.



Note:

By default, device names are assigned by the system. The device name for I/O-optimized instances starts from `/dev/vdb` to `/dev/vdz`. If the device name is `/dev/xvd*` (where, * is a lowercase letter), then a non-I/O-optimized instance is being used.

Procedure

In this example, a new 20 GiB data disk with the device name of `/dev/vdb` is used to create a single-partition data disk and format the disk to an ext4 file system. An I/O-optimized instance with CentOS 7.6 is used.

1. **Connect to the instance** to which the data disk is attached.
2. Run the `fdisk -l` command to view the data disks of the instance.



Note:

If `/dev/vdb` is not displayed in the output, no data disk is attached to the instance. In this case, check whether a data disk is mounted to the instance.

3. Create a single-partition data disk by running the following commands in sequence:
 - a. Run the `fdisk -u /dev/vdb` command to partition the data disk.
 - b. Enter `p` and press Enter to view the partitions of the data disk. In this example, the data disk is not partitioned.
 - c. Enter `n` and press Enter to create a new partition.
 - d. Enter `p` and press Enter to select the primary partition.



Note:

In this example, you are creating a single-partition data disk, so you only need to create one primary partition. If you want to create four or more partitions, you must create at least one extended partition by selecting `e`.

- e. Enter the partition number and press Enter. In this example, enter `1`.
- f. Enter a number for the first available sector, or press Enter to use the default value of `2048`.
- g. Press Enter to use the default number for the last sector.
- h. Enter `p` and press Enter to view the planned partitions of the data disk.
- i. Enter `w` and press Enter to start partitioning and exit after partitioning.

```
[root@ecshost~]# fdisk -u /dev/vdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write
them.
Be careful before using the write command.
Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0x3e60020e.
```

```
Command (m for help): p
Disk /dev/vdb: 21.5 GB, 21474836480 bytes, 41943040 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: 0x3e60020e
Device Boot Start End Blocks Id System
```

```
Command (m for help): n
Partition type:
p primary (0 primary, 0 extended, 4 free)
e extended
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-41943039, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-41943039, default
41943039):
Using default value 41943039
Partition 1 of type Linux and of size 20 GiB is set
```

```
Command (m for help): p

Disk /dev/vdb: 21.5 GB, 21474836480 bytes, 41943040 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: 0x3e60020e
Device Boot Start End Blocks Id System
/dev/vdb1 2048 41943039 20970496 83 Linux
```

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
```

Syncing disks.

4. Run the `fdisk -lu /dev/vdb` command to view the new partition.

If the following information is displayed, the new partition `/dev/vdb1` is created successfully.

```
[ root @ ecshost ~ ]# fdisk -lu /dev/vdb

Disk /dev/vdb : 21.5 GB, 21474836480 bytes,
41943040 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: 0x3e60020e

Device Boot Start End Blocks Id System
/dev/vdb1 2048 41943039 20970496 83 Linux
```

5. Run the `mkfs . ext4 /dev/vdb1` command to create an ext4 file system on the new partition.



Note:

You can also create other file systems as needed. For example, if you need to share files among different operating systems, such as Linux, Windows, and macOS, you can run the `mkfs . vfat` command to create a VFAT file system. The time required to create a file system depends on the data disk size.

```
[ root @ ecshost ~ ]# mkfs . ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label =
OS type: Linux
Block size = 4096 (log = 2)
Fragment size = 4096 (log = 2)
Stride = 0 blocks, Stripe width = 0 blocks
1310720 inodes, 5242624 blocks
262131 blocks (5.00%) reserved for the super user
First data block = 0
Maximum filesystem blocks = 2153775104
160 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736,
1605632, 2654208,
4096000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
```

```
Writing superblock s and filesystem accounting
informatio n : done
```

6. (Recommended) Run the command `cp / etc / fstab / etc / fstab . bak` to back up `etc / fstab` .

7. Run the command `echo / dev / vdb1 / mnt ext4 defaults 0 0`
`>> / etc / fstab` to write the new partition information to `/ etc / fstab` .



Note:

- Ubuntu 12.04 does not support barrier. Therefore, the correct command for this system is `echo '/ dev / vdb1 / mnt ext4 barrier = 0 0`
`0 ' >> / etc / fstab` .
- If you need to mount the data disk to a folder to store web pages separately, replace `/ mnt` with the desired mount point path.

8. Run the `cat / etc / fstab` command to view the new partition information in `/ etc / fstab` .

```
[ root @ ecshost ~ ]# cat / etc / fstab
#
# / etc / fstab
# Created by anaconda on Wed Dec 12 07 : 53 : 08
2018
#
# Accessible filesystem s , by reference , are maintained
under '/ dev / disk '
# See man pages fstab ( 5 ), findfs ( 8 ), mount ( 8 ) and
/ or blkid ( 8 ) for more info
#
UUID = d67c3b17 - 255b - 4687 - be04 - f29190d373 96 / ext4
defaults 1 1
/ dev / vdb1 / mnt ext4 defaults 0 0
```

9. Run the `mount / dev / vdb1 / mnt` command to mount the file system.

10.Run the `df - h` command to view the disk space and usage.



Note:

If the new file system information is displayed in the response message, the file system is successfully mounted, and you can use the new file system without restarting the instance.

```
[ root @ ecshost ~ ]# df - h
Filesystem Size Used Avail Use % Mounted on
/ dev / vda1 40G 1 . 6G 36G 5 % /
devtmpfs 234M 0 234M 0 % / dev
tmpfs 244M 0 244M 0 % / dev / shm
tmpfs 244M 484K 244M 1 % / run
```

```
tmpfs    244M    0    244M    0 % / sys / fs / cgroup
tmpfs    49M    0    49M    0 % / run / user / 0
/ dev / vdb1    20G    45M    19G    1 % / mnt
```

3.7.3 Partition and format data disks larger than 2 TiB

This topic describes how to partition and format data disks larger than 2 TiB in different operating systems.

Precautions

- The time required for creating a snapshot of a data disk is proportional to the volume of data on the data disk. The larger the volume of data, the longer time it takes to create a snapshot.
- Alibaba Cloud Block Storage supports Master Boot Record (MBR) and GUID Partition Table (GPT) partition formats. MBR is applicable to data disks no larger than 2 TiB, and allows you to create up to four primary partitions. To partition a data disk larger than 2 TiB, use the GPT format.



Note:

Conversion between MBR and GPT may cause data loss. If the resulting disk size when you [create a disk by using a snapshot](#) or [resize a disk](#) exceeds 2 TiB, we recommend that you first check whether the disk uses the MBR partition format. If the MBR partition format is used and you want to retain disk data on your instance, we recommend that you create another data disk and attach it to the instance. Then format a GPT partition and copy the data from the MBR partition to the GPT partition.

- For data disks larger than 2 TiB, use the following partition tools, partition formats, and file systems.

Operating system	Partition tool	Partition format	File system
Windows	Disk Management	GPT	NTFS
Linux	parted	GPT	Ext4 or XFS

Prerequisites

1. The data disk has been attached to your instance. For more information, see [#unique_18](#).
2. You have established a remote connection to the ECS instance. For more information about how to remotely connect to an ECS instance, see [#unique_70](#).

Partition and format a data disk on a Windows instance

The section describes how to partition and format a data disk larger than 2 TiB on a Windows instance running the Windows Server 2008 R2 64-bit operating system.

1. On the taskbar, click Server Manager.
2. In the left-side navigation pane of Server Manager, choose Storage > Disk Management.
3. Find the disk to be partitioned and formatted. This example uses Disk 4. The disk is in the Offline state.
4. Right-click the space next to Disk 4, and click Online.

After it comes online, Disk 4 enters the Not Initialized state.

5. Right-click the space next to Disk 4, and choose Initialize Disk from the shortcut menu.
6. In the Initialize Disk dialog box, select Disk 4 and select GPT as the disk partition method.
7. In the Disk Management window, right-click the Unallocated section of Disk 4, and then choose New Simple Volume from the shortcut menu to create a 4 TiB volume in the NTFS format.
8. In the New Simple Volume Wizard window, click Next, and follow these steps:
 - a. Specify Volume Size: Specify the size of the simple volume to create. If you want to create only one primary partition, use the default value. Click Next. You can also divide Disk 4 into multiple partitions.



Note:

Theoretically, the maximum NTFS volume is the maximum volume of NTFS containing $2^{32} - 1$ clusters. However, in Windows XP Pro, the maximum volume of NTFS is $2^{24} - 1$ clusters. For example, NTFS can support a volume up to 256 TiB when the cluster size is 64 KiB. If the cluster size is 4 KiB, then the maximum volume is 16 TiB. NTFS selects the size of a cluster automatically based on the disk capacity.

- b. Assign Drive Letter or Path: Select a drive letter. This example uses G. Click Next.
- c. Format Partition: Select the formatting settings including file system, allocation unit size, and volume label, and then select Perform a quick format and Enable

file and folder compression as needed. In this example, Perform a quick format is selected. Click Next.

- d. After the new simple volume is created, click Finish to close the New Simple Volume Wizard window.

After the partition is formatted, the status of Disk 4 in the Disk Management window is shown in the following figure.

Convert the partition format of a data disk on a Windows instance



Note:

Converting between partition formats may cause data loss. Ensure that you have backed up the data on the disk before you convert to a different partition format.

This section describes how to convert the partition format on a 3 TiB data disk on a Windows instance running the Windows Server 2012 R2 64-bit operating system.

1. On the Windows Server desktop, right-click the Start icon, and select Disk Management.
2. Find the disk to be partitioned and formatted. This example uses Disk 2.
3. Right-click a simple volume, and then choose Delete Volume from the shortcut menu.
4. Right-click the space next to Disk 2, and then choose Convert to GPT Disk from the shortcut menu.
5. In the Disk Management window, right-click the Unallocated section of Disk 2, and then choose New Simple Volume from the shortcut menu to create a 3 TiB volume in the NTFS format.
6. In the New Simple Volume Wizard window, click Next and follow these steps:
 - a. Specify Volume Size: Specify the size of the simple volume to create. If you want to create only one primary partition, use the default value. Click Next. You can also divide Disk 2 into multiple partitions.



Note:

Theoretically, the maximum NTFS volume is the maximum volume of NTFS containing $2^{32} - 1$ clusters. However, in Windows XP Pro, the maximum volume of NTFS is $2^{31} - 1$ clusters. For example, NTFS can support a volume up to 256 TiB when the cluster size is 64 KiB. If the cluster size is 4 KiB, then the

maximum volume is 16 TiB. NTFS selects the size of a cluster automatically based on the disk capacity.

- b. **Assign Drive Letter or Path:** Select a drive letter. This example uses `F` . Click **Next**.
- c. **Format Partition:** Select the formatting settings including file system, allocation unit size, and volume label, and then select **Perform a quick format** and **Enable file and folder compression** as needed. In this example, **Perform a quick format** is selected. Click **Next**.
- d. After a new simple volume is created, click **Finish** to close the **New Simple Volume Wizard** window.

After the partition is formatted, the status of Disk 2 in the Disk Management window is shown in the following figure.

Partition and format a data disk on a Linux instance

This section describes how to use the `parted` and `e2fsprogs` tools to partition and format a data disk larger than 2 TiB on a Linux instance running the CentOS 7.4 64-bit operating system. In the example, the data disk to be processed is a newly-created 3 TiB empty disk and its device name is `/dev/vdd` .

Prerequisites

The `parted` and `e2fsprogs` tools have been installed on your Linux instance.

```
[ root @ ecshost ~ ]# yum install -y parted
[ root @ ecshost ~ ]# yum install -y e2fsprogs
```

Procedure

To partition and format a data disk larger than 2 TiB and mount the file system, follow these steps:

1. Run the `fdisk -l` command to check whether the data disk exists. The expected command output is as follows. If different information is returned, the data disk is not mounted to the instance.

```
[root@ecshost~]# fdisk -l
Disk /dev/vdd: 3221.2 GB, 3221225472000 bytes, 6291456000 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

2. Run the `parted / dev / vdd` command to start partitioning.
 - a. Run the `mklabel gpt` command to convert the partition format from MBR to GPT.
 - b. Run the `mkpart primary 1 100 %` command to create a primary partition, and specify the starting and ending sectors for the partition.
 - c. Run the `align - check optimal 1` command to check the partition alignment.



Note:

If `1 not aligned` is returned, the partition is not aligned. We recommend that you run the following commands and use the formula $(\text{optimal_io_size} + \text{alignment_offset}) / \text{physical_block_size}$ to obtain the starting sector number to align partitions for optimal performance. For example, if the starting sector number is 1024, you can then run the `mkpart primary 1024s 100 %` command to create a new primary partition.

```
[ root @ ecshost ~ ]# cat / sys / block / vdd / queue /
    optimal_io _size
[ root @ ecshost ~ ]# cat / sys / block / vdd / queue /
    minimum_io _size
[ root @ ecshost ~ ]# cat / sys / block / vdd / alignment_
    offset
[ root @ ecshost ~ ]# cat / sys / block / vdd / queue /
    physical_b lock_size
```

- d. Run the `print` command to view the partition table.

```
(parted) mklabel gpt
(parted) mkpart primary 1 100%
(parted) align-check optimal 1
1 aligned
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdd: 3221GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
Number Start End Size File system Name Flags
1 17.4kB 3221GB 3221GB primary
```

- e. Run the `quit` command to exit the `parted` tool.

3. Run the `partprobe` command to make the system re-read the partition table.

4. Run one of the following commands to create a file system for the `/dev/vdd1` partition.

- Create an Ext4 file system.

```
[ root @ ecshost ~ ]# mkfs -t ext4 /dev/vdd1
```

- Create an XFS file system.

```
[ root @ ecshost ~ ]# mkfs -t xfs /dev/vdd1
```



Note:

- If the capacity of the data disk is 16 TiB, you must format it by using the correct version of `e2fsprogs`. For more information, see [Appendix 1: Update e2fsprogs on a Linux instance](#).
- If you want to disable the lazy init function of an Ext4 file system to avoid its impact on data disk I/O performance, see [Appendix 2: Disable the lazy init function on a Linux instance](#).

5. Run the `mkdir /test` command to create a mount point named `/test`.

6. Run the `mount /dev/vdd1 /test` command to mount partition `/dev/vdd1` to mount point `/test`.

7. Run the `df -h` command to view the current disk space and usage.

If the command output shows information about the newly created file system, the mount operation was successful, and the new file system can be used.

```
[root@ecshost~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/vda1 40G 6.4G 31G 18% /
devtmpfs 487M 0 487M 0% /dev
tmpfs 497M 0 497M 0% /dev/shm
tmpfs 497M 364K 496M 1% /run
tmpfs 497M 0 497M 0% /sys/fs/cgroup
tmpfs 100M 0 100M 0% /run/user/0
```

```
/dev/vdd1 2.9T 89M 2.8T 1% /test
```

8. (Optional) Write new partition information to `/etc/fstab` to enable automatic partition mounting when the instance is started.

a. (Optional) Run the `cp /etc/fstab /etc/fstab.bak` command to back up `etc/fstab`.

b. Run the `echo /dev/vdd1 /test ext4 defaults 0 0 >> /etc/fstab` command to write new partition information to `/etc/fstab`.

c. Run the `cat /etc/fstab` command to check `/etc/fstab` information.

If the new partition information appears in the command output, the write operation was successful.

You have now partitioned and formatted a 3 TiB data disk.

Appendix 1: Update e2fsprogs on a Linux instance

If the disk capacity is 16 TiB, you must use e2fsprogs 1.42 or later to format its partitions to an Ext4 file system. If e2fsprogs of a version earlier than 1.42 is used, the following error occurs:

```
mkfs . ext4 : Size of device / dev / vdd too big to be
expressed in 32 bits using a blocksize of 4096 .
```

To install a later version of e2fsprogs, such as 1.42.8, follow these steps:

1. Run the `rpm -qa | grep e2fsprogs` command to check the current e2fsprogs version.

```
$sudo rpm -qa | grep e2fsprogs
e2fsprogs-libs-1.41.12-3
e2fsprogs-1.41.12-3
e2fsprogs-libs-1.39-33.1.ali5
e2fsprogs-devel-1.39-33.1.ali5
```

If the version is earlier than 1.42, perform the following steps to update the software.

2. Run the following command to download e2fsprogs 1.42.8. You can go to [e2fsprogs](#) to obtain the latest software package.

```
wget https://www.kernel.org/pub/linux/kernel/people/tytso/e2fsprogs/v1.42.8/e2fsprogs-1.42.8.tar.gz
```

3. Run the following commands to compile the tool of a later version.

```
tar xvfz e2fsprogs-1.42.8.tar.gz
cd e2fsprogs-1.42.8
./configure
make
make install
```

4. Run the following command to check whether e2fsprogs is updated.

```
rpm -qa | grep e2fsprogs
```

Appendix 2: Disable the lazy init function on a Linux instance

The lazy init function of an Ext4 file system is enabled by default. While this function is enabled, the instance will initiate a thread to continuously initialize the metadata of the Ext4 file system. Therefore, right after you partition and format a data disk, the test of disk IOPS performance will be affected, resulting in a lower IOPS.

If you need to test the data disk performance immediately after partitioning and formatting the disk, run the following command to disable the lazy init function when you initialize the file system.

```
mke2fs -O 64bit,has_journal,extents,huge_file,flex_bg,uninit_bg,dir_nlink,extra_isize -E lazy_itabls_e_init=0,lazy_journal_init=0 /dev/vdd1
```



Note:

If the lazy init function is disabled, it may take a longer time to initialize the file system. For example, it may take 10 to 30 minutes to initialize the file system of a 32 TiB data disk. Enable or disable the lazy init function based on your needs.

3.7.4 Create a file system on a raw data disk

This topic uses an ECS instance running the Ubuntu operating system to show how to create a file system on a raw data disk. You can skip the steps to create a new partition, such as /dev/vdb1 or /dev/vdb2, and directly create a file system if no

partition is required. This method is only applicable to ECS instances running a Linux operating system.

Prerequisites

A cloud disk has been created and attached to an ECS instance. For more information, see [Create a Pay-As-You-Go cloud disk](#) and [#unique_18](#).

Procedure

1. Remotely connect to an ECS instance as a root user. For more information, see [#unique_70](#).
2. Run the following command to view the name of the attached cloud disk.

```
fdisk -l
```

If the following output is displayed, it indicates that the ECS instance has two cloud disks: /dev/vda as the system disk and /dev/vdb as a data disk.

```
root@iZuf6em...dzdZ:~# fdisk -l
Disk /dev/vda: 80 GiB, 85899345920 bytes, 167772160 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
Disklabel type: dos
Disk identifier: 0x6c740fc2

Device      Boot Start      End  Sectors  Size Id Type
/dev/vda1   *      2048 83886046 83883999   40G 83 Linux

Disk /dev/vdb: 100 GiB, 107374182400 bytes, 209715200 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
```

3. Create a file system for the /dev/vdb data disk. Example:

- Create an ext4 file system

```
mkfs . ext4 / dev / vdb
```

- Create an ext3 file system

```
mkfs . ext3 / dev / vdb
```

- Create an xfs file system

```
mkfs . xfs / dev / vdb
```

- Create a btrfs file system

```
mkfs . btrfs / dev / vdb
```

4. Optional: Create a mount directory, such as /media/vdb.

If this step is omitted, you can also attach the disk to an existing directory.

```
mkdir /media/vdb
```

5. Attach the disk to the mount directory.

```
mount /dev/vdb /media/vdb
```

6. Run the `df` command to view the data disk information.

The mount directory information of the disk is displayed, indicating that the operation was successful.

```
[root@ecshost ~]# df -h
Filesystem Size Used Avail Use% Mounted on
udev 3.9G 0 3.9G 0% /dev
tmpfs 798M 2.9M 795M 1% /run
/dev/vda1 40G 3.2G 35G 9% /
tmpfs 3.9G 0 3.9G 0% /dev/shm
tmpfs 5.0M 0 5.0G 0% /run/lock
tmpfs 3.9G 0 3.9G 0% /sys/fs/cgroup
tmpfs 798M 0 798M 0% /run/user/0
/dev/vdb 98G 61M 93G 1% /media/vdb
```

More information

[#unique_51](#)

[#unique_19](#)

3.8 Resize cloud disks

3.8.1 Overview

You can resize cloud disks as your business and application data grows.

Scenarios

You can resize the storage capacity of a single instance in the following ways:

- [Resize an existing cloud disk](#). You need to resize an existing partition or a new partition.
- [Create a new cloud disk](#), [attach](#) it to the instance as a data disk, and then [partition or format](#) the disk.
- [Change a system disk](#) and specify a higher system disk capacity.

This topic describes the thresholds of extended disks and how to resize an existing cloud disk.

Thresholds of extended system disks

The new capacity value must be greater than the existing capacity of the system disk , but equal to or less than 500 GiB. The following table describes the thresholds of extended system disks for different images.

Image	Maximum capacity (GiB)
Linux (excluding CoreOS) and FreeBSD	Max {20, current capacity of the system disk} to 500
CoreOS	Max {30, current capacity of the system disk} to 500
Windows	Max {40, current capacity of the system disk} to 500

For example, the current capacity of the system disk of a CentOS instance is 35 GiB. After you resize the system disk, its capacity must be equal to or greater than 35 GiB, but equal to or less than 500 GiB.

Thresholds of extended data disks


The new value must be greater than the existing capacity of the data disk. The following table lists the data disk resizing limits for different cloud disk categories.

Disk type	Current capacity (GiB)	Maximum capacity (GiB)
Basic disk	< 2,000	2,000

Disk type	Current capacity (GiB)	Maximum capacity (GiB)
SSD disk or ultra disk	< 6,144	6,144
SSD disk or ultra disk	\geq 6,144	Cannot be extended
ESSD disk	< 32,768	32,768

Resize a cloud disk

1. Log on to the console or use the API (ResizeDisk) to [resize a cloud disk](#).
2. Log on to the console or use the API (RebootInstance) to [restart an instance](#).
3. Remote access the instance and resize the partition and file system:

Before you resize the cloud disk	After you resize the cloud disk (GiB)	Resize the partition and file system
Not partitioned	< 2,048	<ul style="list-style-type: none"> • Partition or format a Windows data disk • Partition or format a Linux data disk
	\geq 2,048	Partition or format a data disk of more than 2 TiB
Partitioned	< 2,048	<ul style="list-style-type: none"> • Resize a file system of a Windows disk • Resize a file system of a Linux system disk • Resize a file system of a Linux data disk
	\geq 2,048	<ul style="list-style-type: none"> • For the GPT partition format: <ul style="list-style-type: none"> - Resize a file system of a Windows disk - Resize a file system of a Linux system disk - Resize a file system of a Linux data disk • For the MBR partition format: You cannot resize the partition <div>  <p>Note: If you want to resize a cloud disk to more than 2,048 GiB, the cloud disk cannot use the MBR partition format. In this case, you must first query the partition format of the cloud disk. If the MBR partition format is used, we recommend that you create and attach another data disk. Format a GPT partition and copy the data in the MBR partition to the GPT partition.</p> </div>

3.8.2 Resize cloud disks online

You can resize Alibaba Cloud disks online. In this mode, you do not need to restart instances. You can resize both system disks and data disks online as described in this topic. The size of the disks is expected to increase as your business and application data grows.

What is online resizing?

In the online mode, you can see that cloud disks are extended without restarting instances by using the console or API ([RebootInstance](#)). The difference between the online and offline modes:

- **Online mode:** You do not need to restart instances and instances must be in the `running` (`Running`) state.
- **Offline mode:** You must restart instances and instances can be in the `running` (`Running`) or `stopped` (`Stopped`) state.

Notes

System limits

- For more information about thresholds of extended system disks and data disks, see [Overview](#).
- The operating system must meet the following conditions when you resize system disks. Otherwise, you must select the [offline mode](#).
 - Windows operating system: cannot be Windows Server 2003.
 - Linux operating system: The kernel version must be later than 3.6 when you run the `uname -a` command.
- When you resize cloud disks, only storage capacity is extended, but not file systems. You can allocate your own storage space after resizing. For more information, see [Next operations](#).

What is not supported

- You cannot resize a cloud disk for which a snapshot is being created.
- After you click [Renew for Configuration Downgrade](#) for the subscription instance, you cannot resize subscription cloud disks of the instance during the remaining period of the current billing cycle.
- If a data disk adopts the MBR partition format, you cannot resize the data disk to more than 2 TiB. If you want to resize a data disk to 2 TiB while the MBR partition

format is used, we recommend that you create and attach another data disk.

Format a GPT partition and copy the data in the MBR partition to the GPT partition.

What is supported


- You can resize cloud disks that are in the in use state. This occurs when instances to which the disks are attached are in the Running(Running) state.
- You can resize cloud disks of I/O-optimized instances.
- You can resize ultra disks and SSD disks.
- For Windows instances, you can only resize the NTFS file system.

Preparations

1. [Create a snapshot](#) for data backup, to prevent data loss caused by any errors.
2. Check whether you need to [update the RedHat VirtIO SICI driver](#) for the instance if your Windows instance was created earlier than March 30, 2019.

Procedure

Perform the following steps to resize a cloud disk in the ECS console:

1. Locate the cloud disk to be extended. Choose More > Resize Disk in the Actions column.
2. Select  Resize Online. You do not need to restart the instance to validate new capacity changes (resize and partition operations).
3. Set the extended capacity, which must be greater than the current capacity.
4. Confirm the fee, read ECS Service Terms and Product Terms of Service, and then click Confirm to Resize.
5. Complete the payment.



Note:

If you do not select Resize Online or ECS does not meet the conditions for Resize Online, you must log on to the console or use the API ([RebootInstance](#)) to restart the instance for the changes to take effect.

Related API: You can also call [ResizeDisk](#) to perform the resize operation.

Next operations

The following table lists the available operations after you resize a cloud disk.

However, this depends on whether the cloud disk is attached and partitioned.

Disk status	A disk that is not attached or partitioned	A disk that is attached but not partitioned	A disk that is attached and partitioned
Next	<p>If a data disk is in the available (Available) state, the extended capacity takes effect after you complete the payment. Then you can perform the following steps:</p> <ol style="list-style-type: none"> Log on to the console or use the API (AttachDisk) to attach a cloud disk. Resize or format a partition: <ul style="list-style-type: none"> • Partition or format a Linux data disk of less than 2 TiB • Partition or format a Windows data disk of less than 2 TiB • Partition or format a data disk of more than 2 TiB 	<p>You can perform the following steps:</p> <ol style="list-style-type: none"> Resize or format a partition: <ul style="list-style-type: none"> • Partition or format a Linux data disk of less than 2 TiB • Partition or format a Windows data disk of less than 2 TiB • Partition or format a data disk of more than 2 TiB Resize a file system: <ul style="list-style-type: none"> • For a Linux instance, see Resize a partition and file system of a Linux data disk or Resize a partition and file system of a Linux system disk • For a Windows instance, see Resize a file system of a Windows disk 	<p>Resize a partition and file system:</p> <ul style="list-style-type: none"> • For a Linux instance, see Resize a partition and file system of a Linux data disk or Resize a partition and file system of a Linux system disk • For a Windows instance, see Resize a file system of a Windows disk

3.8.3 Resize cloud disks offline

You can resize both system disks and data disks as described in this topic. The size of the disks is expected to increase as your business and application data grows.

What is offline resizing?

In the offline mode, you can see that cloud disks are extended only after you [restart instances](#) by using the console or API ([RebootInstance](#)). Instances can be in the `running` (`Running`) or `stopped` (`Stopped`) state.

Notes

System limits

- For more information about thresholds of extended system disks and data disks, see [Overview](#).
- When you resize cloud disks, only storage capacity is extended, but not file systems. You can allocate your own storage space after resizing. For more information, see [Next operations](#).

What is not supported

- You cannot resize a cloud disk for which a snapshot is being created.
- After you click [Renew for Configuration Downgrade](#) for a subscription instance, you cannot resize subscription cloud disks of the instance during the remaining period of the current billing cycle.
- You cannot resize system disks of Windows Server 2003 instances.
- If a data disk adopts the MBR partition format, you cannot resize the data disk to more than 2 TiB. If you want to resize a data disk to 2 TiB while the MBR partition format is used, we recommend that you create and attach another data disk. Format a GPT partition and copy the data in the MBR partition to the GPT partition.

What is supported

- You can resize cloud disks which are in the `in use` (`In Use`) state. This occurs when instances to which the disks are attached are in the `running` (`Running`) or `stopped` (`Stopped`) state.
- You can resize basic disks, ultra disks, SSD disks, and ESSD disks.
- For Windows instances, you can only resize the NTFS file system.

Preparations

Create a [snapshot](#) for data backup, to prevent data loss caused by any errors.

Procedure

Perform the following steps to resize a cloud disk in the ECS console:

1. Locate the cloud disk to be extended. Choose More > Resize Disk in the Actions column.
2. Set the extended capacity, which must be greater than the current capacity.
3. Confirm the fee, read ECS Service Terms and Product Terms of Service, and then click Confirm to Resize.
4. Complete the payment.
5. Restart the instance for the changes to take effect.



Note:

You can see that cloud disks are extended only after you [restart instances](#) by using the console or API ([RebootInstance](#)). It does not work if you perform the restart operation in the instance operating system.

Related API: You can also call [ResizeDisk](#) to perform the resize operation.

Next operations

The following table lists the available operations after you resize a cloud disk. However, this depends on whether the cloud disk is attached and partitioned.

Disk status	A disk that is not attached or partitioned	A disk that is attached but not partitioned	A disk that is attached and partitioned
-------------	--	---	---

Next	<p>If a data disk is in the available (Available) state, the extended capacity takes effect after you complete the payment. Then you can perform the following steps:</p> <ol style="list-style-type: none"> 1. Log on to the console or use the API (AttachDisk) to attach a cloud disk. 2. Resize or format a partition: <ul style="list-style-type: none"> • Partition or format a Linux data disk of less than 2 TiB • Partition or format a Windows data disk of less than 2 TiB • Partition or format a data disk of more than 2 TiB 	<p>You can perform the following steps:</p> <ol style="list-style-type: none"> 1. Resize or format a partition: <ul style="list-style-type: none"> • Partition or format a Linux data disk of less than 2 TiB • Partition or format a Windows data disk of less than 2 TiB • Partition or format a data disk of more than 2 TiB 2. Resize a file system: <ul style="list-style-type: none"> • For a Linux instance, see Resize a partition and file system of a Linux data disk or Resize a partition and file system of a Linux system disk • For a Windows instance, see Resize a file system of a Windows disk 	<p>Resize a partition and file system:</p> <ul style="list-style-type: none"> • For a Linux instance, see Resize a partition and file system of a Linux data disk or Resize a partition and file system of a Linux system disk • For a Windows instance, see Resize a file system of a Windows disk
------	--	--	---

3.8.4 Extend a Windows file system

This topic describes how to extend a Windows file system. Resizing a cloud disk does not extend the file system; it only extends the storage capacity. Therefore, you need to format the new storage capacity after you resize a cloud disk.

Limits

The information contained in this topic applies only to disks that are in use and are attached to Running instances. For information on how to attach, partition, or format

disks that are in the Available state, see [Attach a cloud disk](#) and [Partition and format a data disk](#).


Preparations

1. Use the ECS console or call the API to resize the cloud disk.
2. [Create a snapshot](#) to back up your data.
3. Attach the cloud disk to the instance and make sure that the instance is in the Running state. For information about the connection methods, see [Overview](#).
4. Format and partition the data disk. For more information, see [Partition and format a data disk](#).

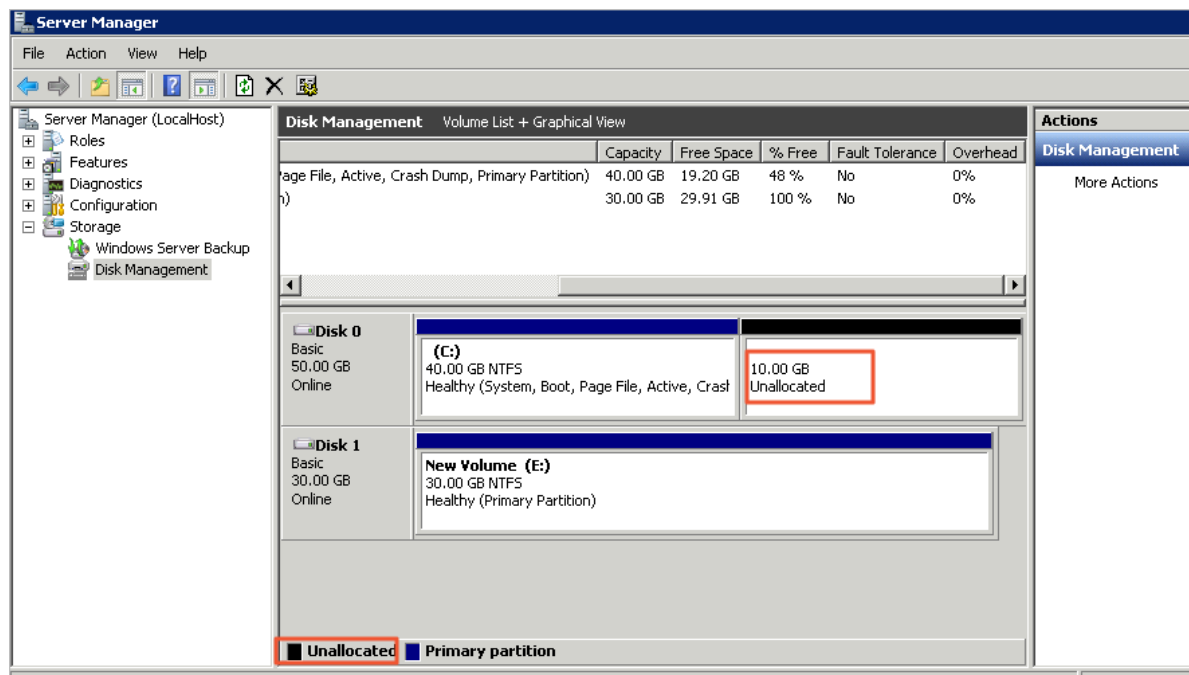
Extend the system disk partition

After you resize a system disk in the ECS console, you need to connect to the instance to extend the file system of the corresponding system disk partition. In this example, the operating system is Windows Server 2008 R2 Enterprise Edition 64-bit and the system disk is resized from 50 GiB to 72 GiB.

1. [Connect to the Windows instance](#).
2. Open Server Manager.
3. In the left-side navigation pane, choose Storage > Disk Management.
4. Choose Action > Refresh or Action > Rescan Disks.

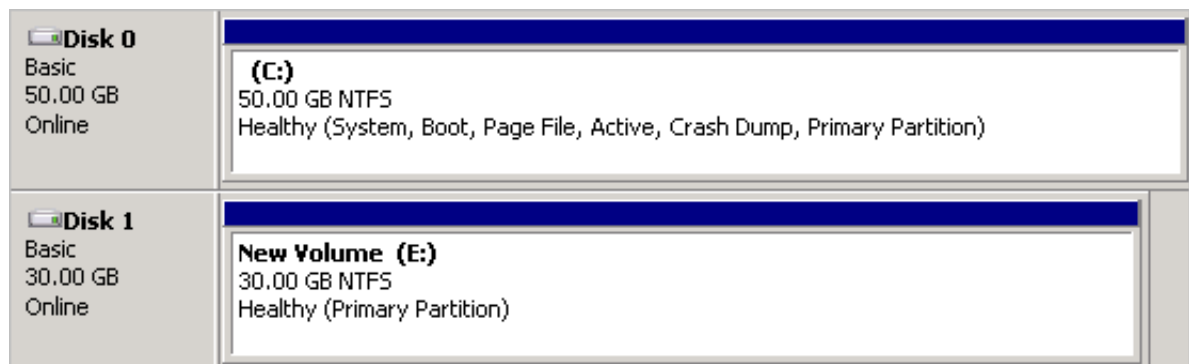


5. In the Disk Management area, view the unallocated capacity. In this example, Disk 0 is the resized system disk.



6. Right-click the blank space in the Disk 0 area, and then select Extend Volume.
7. Follow the instructions provided by the Extend Volume Wizard to extend the volume.

The new disk space is automatically added to the original volume.



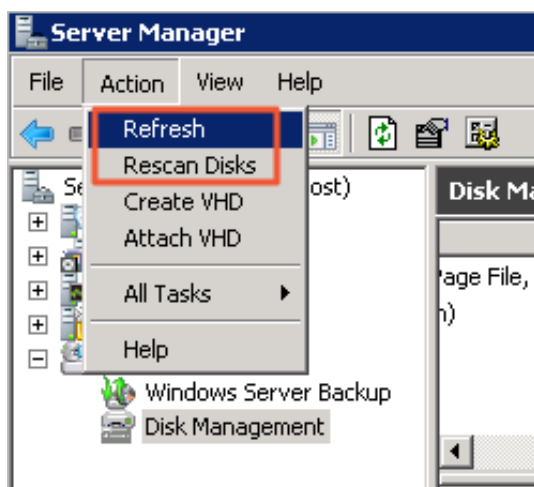
Extend a data disk partition

After you resize a data disk in the ECS console, you need to connect to the instance to extend the file system of the corresponding data disk partitions. In this example, the operating system is Windows Server 2008 R2 Enterprise Edition 64-bit and the data disk is resized from 20 GiB to 30 GiB.


1. [Connect to the Windows instance.](#)
2. Open Server Manager.

3. In the left-side navigation pane, choose Storage > Disk Management.

4. Choose Action > Refresh or Action > Rescan Disks.




5. In the Disk Management area, view the unallocated capacity. In this example, Disk 1 is the resized data disk.




6. Extend Disk 1.

- To use the new disk space to extend the existing partition, follow these steps:
 - a. Right-click the blank space in the Disk 1 area, and then select **Extend Volume**.




- b. Follow the instructions provided by the Extend Volume Wizard to extend the volume.

The new disk space is automatically added to the original volume.




- To use the new disk space to add a new partition, follow these steps:
 - a. Right-click the bank space in the Disk 1 area, and then select **New Simple Volume**.



- b. Follow the instructions provided by the New Simple Volume Wizard to extend the volume.

The new disk space is added to a new partition.



3.8.5 Resize partitions and file systems of Linux system disks

This topic describes how to use the growpart and resize2fs tools to resize the partitions and extend the file systems of Linux system disks.

Application scope

The procedures in this topic apply to system disks with the following partition formats and file system formats:

- Partitions: MBR and GPT
- File systems: ext*, XFS, and Btrfs

Preparations

1. [Create a snapshot](#) for data backup, to prevent data loss caused by misoperations.
2. Use the ECS console or call the API operation to [Resize cloud disks offline](#).
3. Connect to the ECS instance remotely. For more information about connection methods, see [Methods to connect to an ECS instance](#).
4. Install the growpart or xfsprogs tool based on your operating system.

- In CentOS 7 and Aliyun Linux, run the following commands:

```
yum install cloud-utils-growpart
yum install xfsprogs
```

- In Ubuntu 14, Ubuntu 16, Ubuntu 18 and Debian 9, run the following commands:

```
apt install cloud-guest-utils
apt install xfsprogs
```

- In Debian 8, openSUSE 42.3, openSUSE 13.1, and SUSE Linux Enterprise Server 12 SP2, use an upstream version of growpart or xfsprogs.

5. Check the kernel version of your instance, for example, by running the `uname -a` command.

- For kernels 3.6.0 or later, see [Procedure for instances with kernels 3.6.0 or later](#).
- For kernels earlier than 3.6.0, see [Procedure for instances with kernels earlier than 3.6.0](#). If your instance runs a Linux distribution such as CentOS 6, Debian 7, or SUSE Linux Enterprise Server 11 SP4, the instance must be restarted by using the console or an API operation before resizing can be completed.

Procedure for instances with kernels 3.6.0 or later

This procedure uses an instance running CentOS 7 to describe how to resize a system disk partition.

1. Run the `fdisk -l` command to check the size of the disk.

In this example, the `/dev/vda` disk size is 100 GiB.

```
[root@ecshost ~]# fdisk -l
Disk /dev/vda: 107.4 GB, 107374182400 bytes, 209715200 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: 0x0008d73a
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vda1	*	2048	41943039	20970496	83	Linux

2. Run the `df -h` command to check the disk partition size.

In this example, the `/dev/vda1` partition size is 20 GiB.

```
[root@ecshost ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/vda1       20G   1.5G   18G   8% /
devtmpfs        7.8G     0   7.8G   0% /dev
tmpfs           7.8G     0   7.8G   0% /dev/shm
tmpfs           7.8G  344K   7.8G   1% /run
tmpfs           7.8G     0   7.8G   0% /sys/fs/cgroup
tmpfs           1.6G     0   1.6G   0% /run/user/0
```

3. Run the `growpart <DeviceName> <PartionNumber>` command to use the `growpart` tool to resize the specified system disk and partition.

In this example, the first partition of the system disk is resized.

```
[root@ecshost ~]# growpart /dev/vda 1
```

```
CHANGED: partition=1 start=2048 old: size=41940992 end=41943040 new:
size=209710462,end=209712510
```

4. Run the `resize2fs <PartitionName>` command to use the `resize2fs` tool to extend the file system.

In this example, the file system of the `/dev/vda1` partition is extended.

```
[root@ecshost ~]# resize2fs /dev/vda1
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/vda1 is mounted on /; on-line resizing required
old_desc_blocks = 2, new_desc_blocks = 7
The filesystem on /dev/vda1 is now 26213807 blocks long.
```



Note:

If an XFS file system is used, run the `xfs_growfs / dev / vda1` command.

5. Run the `df -h` command to check the size of the disk partition.

In this example, the `/dev/vda1` partition size is 100 GiB, which means that the partition is resized.

```
[root@ecshost ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/vda1       99G   1.6G   93G   2% /
devtmpfs        7.8G     0   7.8G   0% /dev
tmpfs           7.8G     0   7.8G   0% /dev/shm
tmpfs           7.8G  500K   7.8G   1% /run
tmpfs           7.8G     0   7.8G   0% /sys/fs/cgroup
tmpfs           1.6G     0   1.6G   0% /run/user/0
```

Procedure for instances with kernels earlier than 3.6.0

This procedure uses an instance running CentOS 6 to describe how to resize a system disk partition.

1. Install the `dracut-modules-growroot` tool.

```
[root@ecshost ~]# yum install -y dracut-modules-growroot
```

If a package manager other than `yum` is used, change `yum` to the corresponding command.

2. Run the following command to overwrite the existing `initramfs` file.

```
[root@ecshost ~]# dracut -f
```

3. Run the `fdisk -l` command to check the current disk size.

In this example, the `/dev/vda` disk size is 100 GiB.

```
[root@ecshost ~]# fdisk -l
```

```
Disk /dev/vda: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x0003a7b4
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vda1	*	1	2611	20970496	83	Linux

4. Run the `df -h` command to check the disk partition size.

In this example, the `/dev/vda1` partition size is 20 GiB.

```
[root@ecshost ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/vda1        20G   1.1G   18G   6% /
tmpfs            7.8G     0   7.8G   0% /dev/shm
```

5. Run the `growpart <DeviceName> <PartionNumber>` command to use the `growpart` tool to resize the specified system disk and partition.

In this example, the first partition of the system disk is resized.

```
[root@ecshost ~]# growpart /dev/vda 1
CHANGED: partition=1 start=2048 old: size=41940992 end=41943040 new:
size=209710462,end=209712510
```

6. [Reboot the instance](#) in the console or call the [RebootInstance](#) operation to restart the instance.
7. Remotely connect to the instance.
8. Run the `resize2fs <PartitionName>` command to use the `resize2fs` tool to extend the file system.

In this example, the file system of the `/dev/vda1` partition is extended.

```
[root@ecshost ~]# resize2fs /dev/vda1
resize2fs 1.41.12 (17-May-2010)
Filesystem at /dev/vda1 is mounted on /; on-line resizing required
old desc_blocks = 2, new_desc_blocks = 7
Performing an on-line resize of /dev/vda1 to 26213807 (4k) blocks.
The filesystem on /dev/vda1 is now 26213807 blocks long.
```



Note:

If an XFS file system is used, run the `xfs_growfs / dev / vda1` command.

9. Run the `df -h` command to check the size of the disk partition.

In this example, the `/dev/vda1` partition size is 100 GiB, which means that the partition is resized.

```
[root@ecshost ~]# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/vda1	99G	1.1G	93G	2%	/
tmpfs	7.8G	0	7.8G	0%	/dev/shm

Related operations

[Resize partitions and file systems of Linux data disks](#)

3.8.6 Resize partitions and file systems of Linux data disks

When cloud disks are resized, the storage capacity increases but the file system is not affected. To resize the file system and increase the storage capacity of ECS instances, follow the steps in this topic.

Preparations

Before you resize partitions and file systems of a data disk, you must complete the following operations.

1. [Create a snapshot](#) for data backup, to prevent data loss caused by misoperations.
2. Use the ECS console or call the API operation to [Resize cloud disks offline](#).
3. Connect to the ECS instance remotely. For more information about connection methods, see [Methods to connect to an ECS instance](#).

Check the partition format and the file system type

In the following example, the data disk is an ultra disk, the ECS instance operating system is CentOS 7.5 64-bit, and the data disk name is /dev/vdb.

1. Run the `fdisk -lu <DeviceName>` command to check whether the disk is partitioned.

In the example, the disk has a partition called `/ dev / vdb1`. " System "="

Linux " indicates that the partition scheme of the disk is MBR. If " System "="

GPT " is displayed, it indicates that the partition scheme of the disk is GPT.

```
[root@ecshost ~]# fdisk -lu /dev/vdb
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 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: 0x9277b47b

Device Boot Start End Blocks Id System
```

```
/dev/vdb1 2048 41943039 20970496 83 Linux
```

2. Run the `blkid <PartitionName>` command to check the file system type.

In the example, the file system of `/ dev / vdb1` is `ext4` .

```
[root@ecshost ~]# blkid /dev/vdb1
/dev/vdb1: UUID="e97bf1e2-fc84-4c11-9652-73*****24" TYPE="ext4"
```



Note:

No results are returned if a data disk does not have partitions or file systems, or if a data disk has partitions, but not file systems.

3. Run the following command to check the status of the file system.

- For the `ext*` file system: `e2fsck -n <dst_dev_part_path>`
- For the `xfs` file system: `xfs_repair -n <dst_dev_part_path>`



Warning:

In the example, the file system is in the `clean` state. If the file system is not in the `clean` state, troubleshoot the file system.

```
[root@ecshost ~]# e2fsck -n /dev/vdb1
e2fsck 1.42.9 (28-Dec-2013)
Warning! /dev/vdb1 is mounted.
Warning: skipping journal recovery because doing a read-only
filesystem check.
/dev/vdb1: clean, 11/1310720 files, 126322/5242624 blocks
```

Select a method to resize partitions or file systems

Select proper operations based on the partition format and file system conditions.

Scenario	Operation
The data disk has partitions and file systems.	<ul style="list-style-type: none"> • To resize the existing MBR partitions of the data disk, see Option 1: Resize existing MBR partitions. • If the new disk space is used to add new MBR partitions, see Option 2: Add and format MBR partitions. • To resize the existing GPT partitions of the data disk, see Option 3: Resize existing GPT partitions. • If the new disk space is used to add new GPT partitions, see Option 4: Add and format GPT partitions.
The new data disk does not have partitions or file systems.	Partition and format a data disk or <code>#unique_56</code> after you resize the disk space in the console.

Scenario	Operation
The raw data disk has file systems, but not partitions.	After you resize the disk space in the console, see Option 5: Resize the file system of a raw data disk .
The data disk is not attached to any instance.	Attach the data disk to an instance and then perform the steps in this topic to resize the data disk.

**Note:**

If the data disk contains an MBR partition, you cannot resize the data disk to greater than or equal to 2 TiB. To prevent data loss, we recommend that you create a cloud disk larger than 2 TiB. Format a GPT partition as specified in [#unique_56](#) and copy the data from the MBR partition to the GPT partition.

Option 1: Resize existing MBR partitions**Note:**

To prevent data loss, we do not recommend that you resize partitions and file systems when they are attached to ECS instances. Detach the attached partitions (`umount`) first. After you resize the partitions and they can be normally used, attach the partitions (`mount`) again. The following operation methods are recommended for different Linux kernel versions:

- If the instance kernel is earlier than version 3.6, detach the partition, modify the partition table, and then resize the file system.
- If the instance kernel is version 3.6 or later, modify the partition table, notify the kernel of updating the partition table, and then resize the file system.

To resize an existing MBR partition, perform the following steps:

Step 1: Modify the partition table

1. Run the `fdisk -lu /dev/vdb` command and record the start and end sectors of the existing partition.

In this example, the start sector of `/dev/vdb1` is 2048 and the end sector is 41943039 .

```
[root@ecshost ~]# fdisk -lu /dev/vdb
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 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: 0x9277b47b
```

```
Device Boot Start End Blocks Id System
/dev/vdb1 2048 41943039 20970496 83 Linux
```

2. View the mount path of the data disk. Detach the partition based on the returned file path and wait until the partition is fully detached.

```
[root@ecshost ~]# mount | grep "/dev/vdb"
/dev/vdb1 on /mnt type ext4 (rw,relatime,data=ordered)
[root@ecshost ~]# umount /dev/vdb1
[root@ecshost ~]# mount | grep "/dev/vdb"
```

3. Run the `fdisk` command to delete the existing partition.

- a. Run the `fdisk -u /dev/vdb` command to partition the data disk.
- b. Enter `p` to display the partition table.
- c. Enter `d` to delete the partition.
- d. Enter `p` to confirm that the partition has been deleted.
- e. Enter `w` to save changes and exit.

```
[root@ecshost ~]# fdisk -u /dev/vdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Command (m for help): p
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 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: 0x9277b47b
Device Boot Start End Blocks Id System
/dev/vdb1 2048 41943039 20970496 83 Linux
Command (m for help): d
Selected partition 1
Partition 1 is deleted
Command (m for help): p
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 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: 0x9277b47b
Device Boot Start End Blocks Id System
Command (m for help): w
The partition table has been altered!
Calling ioctl() to re-read partition table.
WARNING: Re-reading the partition table failed with error 16: Device
or resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
```

Syncing disks.

4. Run the `fdisk` command to create a partition.
 - a. Run the `fdisk -u /dev/vdb` command to partition the data disk.
 - b. Enter `p` to display the partition table.
 - c. Enter `n` to create a partition.
 - d. Enter `p` to set the partition as the primary partition.
 - e. Enter `<partition number>` to select a partition number. 1 is selected in the example.



Warning:

The start sector of the new partition must be equal to that of the existing partition. The end sector must be greater than that of the existing partition. Otherwise, the resize operation may fail.

- f. Enter `w` to save changes and exit.

In the example, `/dev/vdb1` is extended from 20 GiB to 40 GiB.

```
[root@ecshost ~]# fdisk -u /dev/vdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Command (m for help): p
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 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: 0x9277b47b
Device Boot Start End Blocks Id System
Command (m for help): n
Partition type:
p primary (0 primary, 0 extended, 4 free)
e extended
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-83886079, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-83886079, default
83886079):
Partition 1 of type Linux and of size 30 GiB is set
Command (m for help): p
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 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: 0x9277b47b
Device Boot Start End Blocks Id System
/dev/vdb1 2048 62916607 31457280 83 Linux
Command (m for help): w
```

```
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
```

5. Run the `lsblk /dev/vdb` command to confirm that the partition table has been added.
6. Run the `e2fsck -n /dev/vdb1` command to confirm that the file system after the resize operation is in the `clean` state.

Step 2: Notify the kernel of updating the partition table

Run the `partprobe < dst_dev_path >` or `partx -u < dst_dev_path >` command to notify kernel that the partition table has been modified and needs to be synchronized.

Step 3: Resize the file system

- For the `ext*` file system (such as `ext3` and `ext4`), run the `resize2fs / dev / vdb1` command and attach the partition again.

```
[root@ecshost ~]# resize2fs /dev/vdb1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/vdb1 to 7864320 (4k) blocks.
The filesystem on /dev/vdb1 is now 7864320 blocks long.
[root@ecshost ~]# mount /dev/vdb1 /mnt
```

- For the `xfs` file system, run the `mount / dev / vdb1 / mnt /` command first, and then the `xfs_growfs / dev / vdb1` command.

```
[root@ecshost ~]# mount /dev/vdb1 /mnt/
[root@ecshost ~]# xfs_growfs /dev/vdb1
meta-data=/dev/vdb1          isize=512    agcount=4,
                           agsize=1310720 blks
                           =                               sectsz=512    attr=2, projid32bit=1
                           =                               crc=1        finobt=0 spinodes=0
data      =                               bsize=4096   blocks=5242880,
                           imaxpct=25
                           =                               sunit=0      swidth=0 blks
naming    =version 2           bsize=4096   ascii-ci=0 ftype=1
log       =internal           bsize=4096   blocks=2560, version=2
                           =                               sectsz=512   sunit=0 blks, lazy-
count=1
realtime  =none               extsz=4096   blocks=0, rtextents=0
data blocks changed from 5242880 to 7864320
```

Option 2: Add and format MBR partitions

If the new disk space is used to add a new MBR partition, perform the following steps:

1. Run the `fdisk -u /dev/vdb` command to create a partition.

In the example, partition `/dev/vdb2` is created for the newly added 20 GiB disk space.

```
[root@ecshost ~]# fdisk -u /dev/vdb
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): p

Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 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: 0x2b31a2a3

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1             2048     41943039     20970496    83   Linux

Command (m for help): n
Partition type:
   p   primary (1 primary, 0 extended, 3 free)
   e   extended
Select (default p): p
Partition number (2-4, default 2): 2
First sector (41943040-83886079, default 41943040):
Using default value 41943040
Last sector, +sectors or +size{K,M,G} (41943040-83886079, default
83886079):
Using default value 83886079
Partition 2 of type Linux and of size 20 GiB is set

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

2. Run the `lsblk /dev/vdb` command to view the partition.

```
[root@ecshost ~]# lsblk /dev/vdb
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vdb   253:16  0  40G  0 disk
├─vdb1 253:17  0   20G  0 part
└─vdb2 253:18  0   20G  0 part
```

3. Format the new partition.

- To create an ext4 file system, run the `mkfs -t ext4 /dev/vdb2` command.

```
[root@ecshost ~]# mkfs.ext4 /dev/vdb2
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
```

```

Stride=0 blocks, Stripe width=0 blocks
1310720 inodes, 5242880 blocks
262144 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2153775104
160 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736,
    1605632, 2654208,
    4096000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
[root@ecshost ~]# blkid /dev/vdb2
/dev/vdb2: UUID="e3f336dc-d534-4fdd-****-b6ff1a55bdbb" TYPE="ext4"

```

- To create an ext3 file system, run the `mkfs . ext3 / dev / vdb2` command.

```

[root@ecshost ~]# mkfs.ext3 /dev/vdb2
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1310720 inodes, 5242880 blocks
262144 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
160 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736,
    1605632, 2654208,
    4096000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
[root@ecshost ~]# blkid /dev/vdb2
/dev/vdb2: UUID="dd5be97d-a630-4593-****-5056def914ea"
SEC_TYPE="ext2" TYPE="ext3"

```

- To create an xfs file system, run the `mkfs . xfs - f / dev / vdb2` command.

```

[root@ecshost ~]# mkfs.xfs -f /dev/vdb2
meta-data=/dev/vdb2            isize=512    agcount=4,
  agsize=1310720 blks
                               sectsz=512    attr=2,
  projid32bit=1
                               crc=1          finobt=0, sparse=0
data      =                    bsize=4096    blocks=5242880,
  imaxpct=25
                               sunit=0        swidth=0 blks
naming    =version 2           bsize=4096  ascii-ci=0 ftype=1

```

```
log      =internal log      bsize=4096   blocks=2560,
version=2
      =                      sectsz=512   sunit=0 blks, lazy-
count=1
realtime =none              extsz=4096   blocks=0,
      rtextents=0
[root@ecshost ~]# blkid /dev/vdb2
/dev/vdb2: UUID="66251477-3ae4-4b44-****-5604420dbecb" TYPE="xfs"
```

- To create a btrfs file system, run the `mkfs . btrfs / dev / vdb2` command.

```
[root@ecshost ~]# mkfs.btrfs /dev/vdb2
btrfs-progs v4.9.1
See http://btrfs.wiki.kernel.org for more information.

Label:              (null)
UUID:               6fb5779b-57d7-4aaf-bf09-82b46f54a429
Node size:          16384
Sector size:        4096
Filesystem size:    20.00GiB
Block group profiles:
  Data:              single           8.00MiB
  Metadata:          DUP              1.00GiB
  System:            DUP              8.00MiB
SSD detected:       no
Incompat features:  extref, skinny-metadata
Number of devices:  1
Devices:
  ID        SIZE  PATH
  1         20.00GiB /dev/vdb2
[root@ecshost ~]# blkid /dev/vdb2
/dev/vdb2: UUID="6fb5779b-57d7-4aaf-****-82b46f54a429"
          UUID_SUB="9bdd889a-ab69-4653-****-d1b6b8723378" TYPE="btrfs"
```

4. Run the `mount / dev / vdb2 / mnt` command to attach the file system to the data disk.
5. Run the `df -h` command to check the current capacity and usage of the data disk.

If information about the new file system is displayed, it indicates that the attach operation is successful.

```
[root@ecshost ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/vda1 40G 1.6G 36G 5% /
devtmpfs 3.9G 0 3.9G 0% /dev
tmpfs 3.9G 0 3.9G 0% /dev/shm
tmpfs 3.9G 460K 3.9G 1% /run
tmpfs 3.9G 0 3.9G 0% /sys/fs/cgroup
/dev/vdb2 9.8G 37M 9.2G 1% /mnt
```

```
tmpfs 783M 0 783M 0% /run/user/0
```

Option 3: Resize existing GPT partitions

To resize an existing GPT partition, perform the following steps: In the example, a data disk of 1 TiB is extended to 32 TiB. The existing partition is `/dev/vdb1`.

1. Run the `fdisk` command to view information about the partition to be extended.

```
[root@ecshost ~]# fdisk -l
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 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: 0x000b1b45

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1 *        2048     83875364     41936658+   83   Linux
WARNING: fdisk GPT support is currently new, and therefore in an
experimental phase. Use at your own discretion.

Disk /dev/vdb: 35184.4 GB, 35184372088832 bytes, 68719476736 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: gpt
Disk identifier: BCE92401-F427-45CC-8B0D-B30EDF279C2F

#           Start          End         Size      Type           Name
 1          2048      2147483647     1024G   Microsoft basic mnt
```

2. View the mount path of the data disk. Detach the partition based on the returned file path and wait until the partition is fully detached.

```
[root@ecshost ~]# mount | grep "/dev/vdb"
/dev/vdb1 on /mnt type ext4 (rw,relatime,data=ordered)
[root@ecshost ~]# umount /dev/vdb1
```

```
[root@ecshost ~]# mount | grep "/dev/vdb"
```

3. Use the `parted` tool to allocate space for the partition.

- a. Run the `parted / dev / vdb` command to start the `parted` tool.
- b. (Optional) Run the `help` command to view the help information.
- c. Run the `print` command to view the number (Number) and capacity (Size) of the partition to be extended.

In the example, partition 1 is to be extended and its existing capacity is 1100 GiB . The new capacity will be allocated to partition 1.

- d. Run the `resizepart < partition number > < capacity allocation percentage >` command to resize the partition.

`resizepart 1 100 %` is used in the example.

- e. Run the `print` command to check whether the number (Number) and capacity (Size) of the partition has been changed.

```
[root@ecshost ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.

(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 35.2TB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB  1100GB  1100GB  ext4         mnt

(parted) resizepart 1 100%
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 35.2TB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB  35.2TB  35.2TB  ext4         mnt
```

- f. Run the `quit` command to exit the `parted` tool.

4. Run the `fsck -f / dev / vdb1` command to check whether the file system is consistent.

```
[root@ecshost ~]# fsck -f /dev/vdb1
fsck from util-linux 2.23.2
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
```



```

Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/vdb1: 11/67108864 files (0.0% non-contiguous),
4265369/268435200 blocks

```

5. Resize the file system and attach the partition again.

- For the ext* file system (such as ext3 and ext4), run the `resize2fs / dev / vdb1` command and attach the partition again.

```

[root@ecshost ~]# resize2fs /dev/vdb1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/vdb1 to 8589934331 (4k) blocks.
The filesystem on /dev/vdb1 is now 8589934331 blocks long.
[root@ecshost ~]# mount /dev/vdb1 /mnt

```

- For the xfs file system, run the `mount / dev / vdb1 / mnt /` command first, and then the `xfs_growfs / dev / vdb1` command.

```

[root@ecshost ~]# mount /dev/vdb1 /mnt/
[root@ecshost ~]# xfs_growfs /dev/vdb1

```

Option 4: Add and format GPT partitions

If the newly added disk space is used to add a new GPT partition, perform the following steps: In the example, a data disk of 32 TiB is used. The existing partition `/ dev / vdb1` has a 4.8 TiB capacity. The `/ dev / vdb2` partition is to be created.

1. Run the `fdisk` command to view information about the existing partition.

```

[root@ecshost ~]# fdisk -l
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 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: 0x000b1b45

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1 *         2048     83875364     41936658+  83  Linux
WARNING: fdisk GPT support is currently new, and therefore in an
experimental phase. Use at your own discretion.

Disk /dev/vdb: 35184.4 GB, 35184372088832 bytes, 68719476736 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: gpt
Disk identifier: BCE92401-F427-45CC-8B0D-B30EDF279C2F

#           Start          End      Size  Type              Name

```

1	2048	10307921919	4.8T	Microsoft basic mnt
---	------	-------------	------	---------------------

2. Use the `parted` tool to create a new partition and allocate space for it.

- a. Run the `parted / dev / vdb` command to start the `parted` tool.
- b. Run the `print free` command to view the disk space to be allocated.

Record the start and end sectors and capacity of the existing partition.

In the example, the start sector of `/ dev / vdb1` is 1,049 KB, the end sector is 5,278 GB, and the capacity is 5,278 GiB.

```
(parted) print free
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 35.2TB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
	17.4kB	1049kB	1031kB	Free Space		
1	1049kB	5278GB	5278GB	ext4	mnt	
	5278GB	35.2TB	29.9TB	Free Space		

- c. Run the `mkpart < partition name > < start sector > < capacity allocation percentage >` command.

In the example, the `/ dev / vdb2` partition named `test` is created. The start sector of the new partition is the end sector of the existing partition. The new capacity is allocated to the new partition.

- d. Run the `print` command to check whether the capacity (`Size`) of the partition has changed.

```
(parted) mkpart test 5278GB 100%
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 35.2TB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	1049kB	5278GB	5278GB	ext4	mnt	
2	5278GB	35.2TB	29.9TB		test	

- e. Run the `quit` command to exit the `parted` tool.

3. Create a file system for the new partition.

- To create an ext4 file system, run the `mkfs . ext4 / dev / vdb2` command.
- To create an ext3 file system, run the `mkfs . ext3 / dev / vdb2` command.
- To create an xfs file system, run the `mkfs . xfs - f / dev / vdb2` command.
- To create a btrfs file system, run the `mkfs . btrfs / dev / vdb2` command.

In the example, an xfs file system is created.

```
[root@ecshost ~]# mkfs -t xfs /dev/vdb2
meta-data=/dev/vdb2            isize=512    agcount=28,
          =                    sectsz=512    attr=2, projid32bit=1
          =                    crc=1        finobt=0, sparse=0
data      =                    bsize=4096    blocks=7301444096,
          =                    imaxpct=5
naming    =version 2           sunit=0      swidth=0 blks
log        =internal log      bsize=4096    ascii-ci=0 ftype=1
          =                    bsize=4096    blocks=521728,
version=2  =                    sectsz=512    sunit=0 blks, lazy-
count=1    =
realtime   =none              extsz=4096    blocks=0, rtextents=0
```

4. Run the `fdisk -l` command to view partition capacity changes.

```
[root@ecshost ~]# fdisk -l
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 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: 0x000b1b45

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1 *        2048     83875364   41936658+  83  Linux
WARNING: fdisk GPT support is currently new, and therefore in an
experimental phase. Use at your own discretion.

Disk /dev/vdb: 35184.4 GB, 35184372088832 bytes, 68719476736 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: gpt
Disk identifier: BCE92401-F427-45CC-8B0D-B30EDF279C2F

#           Start          End         Size      Type          Name
 1          2048    10307921919    4.8T    Microsoft basic mnt
 2    10307921920    68719474687   27.2T    Microsoft basic test
```

5. Run the `blkid` command to view the file system types.

```
[root@ecshost ~]# blkid
```

```
/dev/vda1: UUID="ed95c595-4813-480e-****-85b1347842e8" TYPE="ext4"
/dev/vdb1: UUID="21e91bbc-7bca-4c08-****-88d5b3a2303d" TYPE="ext4"
PARTLABEL="mnt" PARTUUID="576235e0-5e04-4b76-****-741cbc7e98cb"
/dev/vdb2: UUID="a7dcde59-8f0f-4193-****-362a27192fb1" TYPE="xfs"
PARTLABEL="test" PARTUUID="464a9fa9-3933-4365-****-c42de62d2864"
```

6. Attach the new partition.

```
[root@ecshost ~]# mount /dev/vdb2 /mnt
```

Option 5: Resize the file system of a raw data disk

If a raw data disk contains a file system, but not a partition, you can perform the following steps to directly resize the file system.

1. Run different commands based on the file system type.

- For the ext* file system, run the `resize2fs` command using the root permissions to resize the file system. Example:

```
resize2fs /dev/vdb
```

- For the xfs file system, run the `xfs_growfs` command using the root permissions to resize the file system. Example:

```
xfs_growfs /dev/vdb
```

2. Run the `df -h` command to view the data disk information.

The file system has a larger capacity, indicating that the operation was successful.

```
[root@ecshost ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/vda1 40G 1.6G 36G 5% /
devtmpfs 3.9G 0 3.9G 0% /dev
tmpfs 3.9G 0 3.9G 0% /dev/shm
tmpfs 3.9G 460K 3.9G 1% /run
tmpfs 3.9G 0 3.9G 0% /sys/fs/cgroup
/dev/vdb 98G 37G 61G 37% /mnt
tmpfs 783M 0 783M 0% /run/user/0
```

Related operations

- [#unique_67](#)
- [#unique_61](#)

3.8.7 Update the Red Hat VirtIO SCSI driver for Windows ECS instances

You can resize Alibaba Cloud disks online without having to restart ECS instances. If the ECS instances you want to resize were created before March 30, 2019, follow the

procedures described in this topic to check whether the Red Hat VirtIO SCSI driver for these instances need to be updated.

Instructions

- The Red Hat VirtIO SCSI driver supports only Windows Server 2008 and later.
- If an ECS instance is equipped with multiple data disks, the driver update process may take a few minutes.

Download a driver

Download and decompress the [instdrv driver package](#). The following table describes the relationship between operating systems of ECS instances and extracted driver folders.

Driver folder name	Folder description
win7	Windows Server 2008 R2 and Windows 7
Wlh	Windows Server 2008
Win8	Windows Server 2012 and Windows Server 2012 R2
win10	Windows 10, Windows Server 2016, and Windows Server systems in later versions
amd64	64-bit
x86	32-bit
instdrv.exe	Checks whether ECS instances can be resized online

Check the driver version

Method 1: Use instdrv.exe to check the driver version.

1. [Remotely connect to a Windows ECS instance](#).
2. Open Command Prompt.
3. Go to the directory where `instdrv . exe` resides. For example, run the `cd C:\User\lizr\Desktop\instdrv\` command to go to the directory.
4. (Optional) Run the `instdrv . exe` program to view the instructions on how to use the tool. Run the `instdrv . exe - v` command to view the version of the tool.
5. Run the `instdrv . exe - c` command. You can determine whether the operating system supports online resizing based on the command output. If yes,

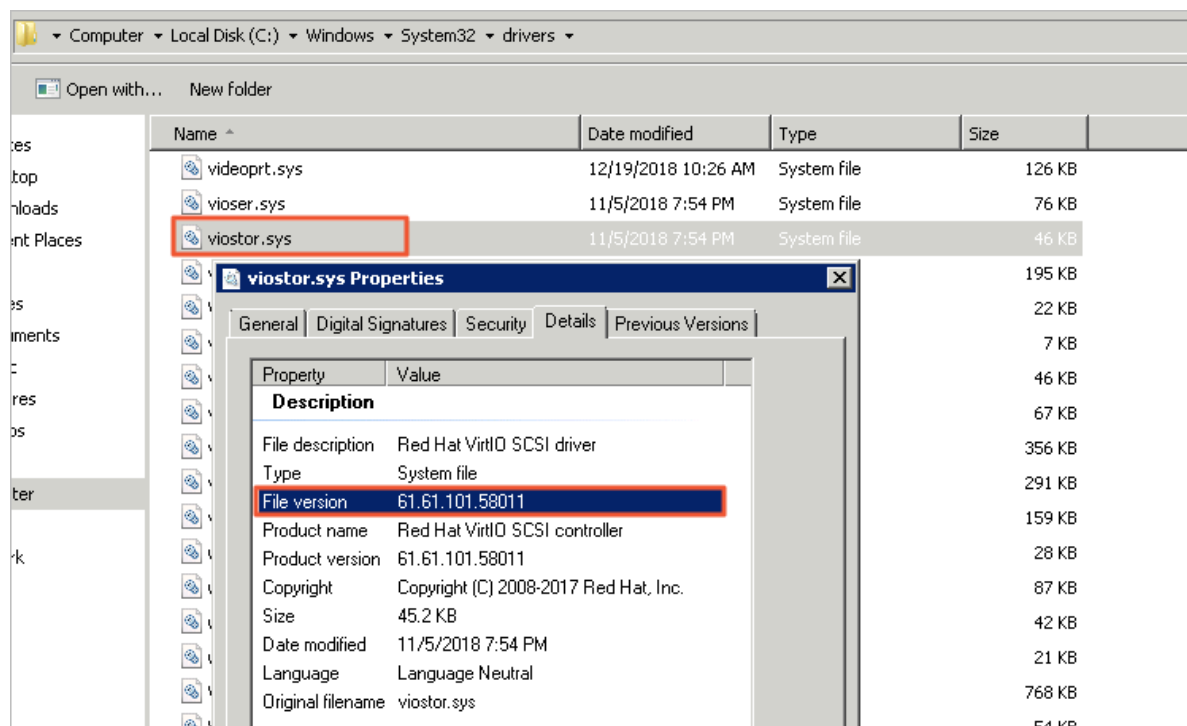
you can directly [resize cloud disks online](#). If not, you need to [update the Red Hat VirtIO SCSI driver](#).

```
PS C:\> [System.Diagnostics.FileVersionInfo]::GetVersionInfo("C:\Windows\System32\drivers\viostor.sys")
```

ProductVersion	FileVersion	FileName
62.61.101.58013	62.61.101.58013	C:\Windows\System32\drivers\viostor.sys

Method 2: Manually check the driver version.

1. Remotely connect to a Windows ECS instance.
2. Go to the system directory `C:\Windows\System32\drivers`.
3. Right-click the `viostor.sys` file and choose Properties from the shortcut menu. On the page that appears, choose Details > File version to view the driver version. If the version of the Red Hat VirtIO SCSI driver for a Windows Server ECS instance is greater than or equal to 58011, you can directly [resize cloud disks online](#). If the driver version is smaller than 58011, [update the Red Hat VirtIO SCSI driver](#).



Update the Red Hat VirtIO SCSI driver

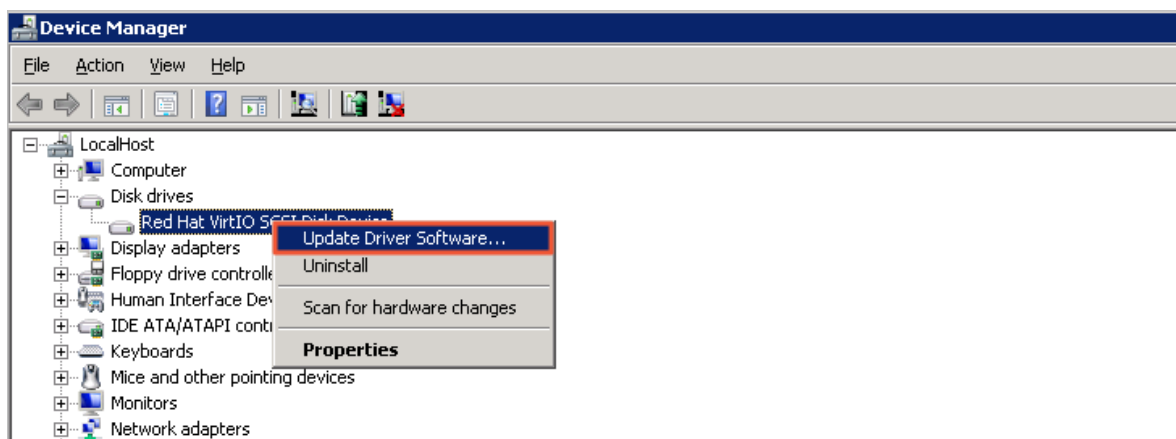
Method 1: Use pnputil to add a driver.

1. Open Command Prompt.
2. Go to the directory where `viostor.inf` resides. For example, run the `cd C:\Users\Administrator\Desktop\amd` command to go to the directory.

3. Run the `pnputil -i -a <directory>` command to add a driver package. Make sure that you have decompressed the target `.inf` file to the directory specified with `<directory>`.
4. Restart the operating system of the ECS instance to make the driver update take effect.

Method 2: Manually add a driver.

1. Open Device Manager.
2. Right-click Red Hat VirtIO SCSI controller under Storage controllers, and select Update Driver Software....



3. Select Browse my computer for driver software.
4. Select Let me pick from a list of device drivers on my computer.
5. Click Have Disk.
6. Select the driver file `viostor` in the corresponding folder, and follow the wizard to update the driver.
7. Restart the operating system of the ECS instance to make the driver update take effect.

Next operations

[Resize cloud disks online](#)

3.9 Change the operating system

3.9.1 Replace the system disk (non-public image)

Replacing the system disk refers to the allocation of a new system disk for the instance. The original system disk is released and the system disk ID is updated. You

can replace the system disk in the following scenarios: 1. You selected an incorrect operating system when you created an ECS instance and need to change the operating system. 2. When your business expands, you want to upgrade the system disk capacity or use another operating system.

Prerequisites

Before you change the image of a system disk to a non-public image, follow these steps:

- For custom images:
 - To use an image captured from an existing ECS instance, you must create a system disk snapshot of this instance and use the snapshot to create a custom image. For more information, see [#unique_55](#) and [#unique_101](#). You must make a copy of the image if the captured image and the instance with the system disk that you want to replace are in different regions. For more information, see [#unique_102](#).
 - To use a local image file, you must import the image in the console or use Packer to create and import the image. The image must be in the same region as your instance. For more information, see [#unique_103](#) and [#unique_104](#).
 - If you want to use an image from another region, you must first make a copy of the image. For more information, see [#unique_102](#).



Note:

Images created by using this method appear in the Custom Image list when you change the system disk.

- To use images owned by another Alibaba Cloud account, you must first let the other account share the images. Then, you can obtain the shared images. For more information, see [#unique_105](#).
- If you want to change to a Linux operating system and use an SSH key pair for authentication, you must first create an SSH key pair. For more information, see [#unique_106](#).
- Replacing a system disk is a high-risk operation that may cause data loss or service interruptions. To minimize the impact of replacing the system disk on your business, we recommend that you create a snapshot of the current system disk before replacing the system disk. For more information, see [#unique_55](#).

- If you want to replace the system disk of a Windows Server instance, make sure that the system disk has sufficient available space. We recommend that you reserve 1 GiB of available space. Otherwise, the system may fail to restart after the system disk is changed.

**Note:**

Create snapshots during off-peak hours to minimize the impact on your business. The first time you create a 40 GiB snapshot requires about 40 minutes. Therefore, you must allocate sufficient time to create the snapshot. Creating a snapshot may temporarily reduce the I/O performance of a block storage device by about 10% and cause a temporary slowdown.

Context

You can select a public image, shared image, custom image, or an image from the marketplace for the ECS instance when you change the system disk. This topic describes how to replace an image of a system disk with a non-public image. For more information about how to replace an image of a system disk with a public image, see [#unique_21](#).

**Note:**

Microsoft no longer offers support for Windows Server 2003. To make sure the safety of your data, we recommend that you stop using Windows Server 2003 for your ECS instances. Alibaba Cloud no longer provides images based on Windows Server 2003. For more information, see [Discontinuation of Windows Server 2003 system image](#).

Replacing a system disk is a high-risk operation that may cause data loss or service interruptions. When you replace a system disk, you must carefully read the following:

- Before you replace the system disk
 - The original system disk is released after you replace the system disk.
We recommend that you create a snapshot to back up the data before the replacement. For more information, see [#unique_55](#).
 - Your workloads are interrupted because you must stop the instance to replace the system disk. For more information, see [#unique_107](#).
 - Make sure that you have enough snapshot quota to use the automatic snapshot policy for the new system disk. You can delete outdated system disk snapshots that you no longer require. For more information, see [#unique_108](#).
 - You cannot change the category of the system disk for your ECS instance.
- After you replace the system disk
 - The IP address and the MAC address of the ECS instance remain unchanged.
 - A new system disk with a new disk ID is allocated to your instance, and the original system disk is released.
 - Snapshots of the outdated system disk cannot be used to restore the new system disk.
 - Manually created snapshots are not released. You can still use the outdated system disk snapshots to create custom images.
 - If the outdated system disk has automatic snapshot policies with Delete Automatic Snapshots While Releasing Disk enabled, the automatic snapshots of the outdated system disk are automatically deleted. The outdated automatic snapshot policies are no longer applicable to new system disks. You must reset the policies.
 - You must redeploy the service environment on the new system disk, which may cause a long interruption to your business.
- Cross-OS system disk replacement

**Note:**

Cross-OS disk replacement refers to switching between Linux and Windows Server systems. After a system disk switches to a different operating system, the ECS instance cannot identify the file system format of the data disk. For regions outside mainland China, system disk replacement across operating systems is

not supported. The system only supports replacement across different Linux distributions or different Windows Server versions.

- If the data disk does not contain important data, you can reinitialize the data disk and create a file system supported by the corresponding system for the data disk. For more information, see [Reinitialize a data disk](#).
- If important data exists on the data disk, follow these steps:
 - For replacing Windows Server with Linux: Manually install the required file system driver, such as NTFS-3G. Linux is unable to recognize the NTFS file system.
 - For replacing Linux with Windows Server: Manually install the required file system driver, such as Ext2Read and Ext2Fsd. Windows Server is unable to recognize some file systems, such as ext3, ext4, XFS.
 - For replacing Windows Server with Linux: You can use either password authentication or SSH key pair authentication.

Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Instances & Images > Instances.
3. In the top navigation bar, select a region.
4. Find the instance with the system disk that you want to replace. In the Actions column, choose More > Instance Status > Stop.

When the instance status changes to Stopped, the instance is successfully stopped.

5. In the Actions column, choose More > Disk and Image > Replace System Disk.
6. In the dialog box that appears, read the precautions for replacing the system disk, and click OK.
7. On the Change System Disk page, configure the following parameters.
 - a) Image Type: Select Custom Image, Shared Image, or Marketplace Image, and select the image that you want.
 - b) System Disk: You cannot change the type of the system disk. However, you can resize the system disk based on your business needs and the new image. The maximum capacity of the new system disk is 500 GiB. The minimum capacity

depends on the current capacity of the system disk and the size of the image.

The following table lists the rules.

Image	Resized capacity (GiB)
Linux (excluding CoreOS) + FreeBSD	From Max{20 or current system disk capacity} (the greater of the two) to 500
CoreOS	From Max{30 or current system disk capacity} (the greater of the two) to 500
Windows Server	From Max{40 or current system disk capacity} (the greater of the two) to 500



Note:

If you have renewed your instance and downgraded the configurations, the system disk capacity cannot be changed until the next billing cycle starts.

c) Security:

- If the new operating system is Windows Server, you can only use password authentication.
- If your instance is an I/O optimized instance and the new operating system is Linux, you can use either password authentication or SSH key pair authentication. You must set a logon password or bind an SSH key pair.

d) Confirm the Instance Cost. The cost includes both the image fee and the system disk fee. For more information about product pricing, see [ECS pricing](#).

e) Click Confirm to Change.

Result

Log on to the ECS console to monitor the instance status. The replacement of the system disk requires about 10 minutes. After the system disk is replaced, the instance automatically restarts.

What's next

After changing the system disk, follow these steps:

- Optional. The automatic snapshot policies of the outdated system disk are now invalid. You can set new automatic snapshot policies for the new system disk. For more information, see [#unique_109](#).

- If both the new and the outdated operating systems are Linux, the outdated instance has data disks mounted and automatic mounting of the partition at startup enabled: After the system disk is changed, the configuration for the mounting of the data disk partitions from the outdated system disk will be lost. You must configure the new partition in the `/etc/fstab` file and mount the partition. You do not need to format and partition the data disk. You can perform the following steps. For more information, see [#unique_19](#).
1. We recommend that you back up the `/etc/fstab` file.
 2. Write the new partition information to the `/etc/fstab` file.
 3. View the new partition information in the `/etc/fstab` file.
 4. Run the `mount` command to mount the partition.
 5. Run the `df -h` command to check the current disk capacity and utilization.

After the partition is mounted, you can use the new data disk without restarting the ECS instance.

More information
[#unique_110](#)

3.9.2 Replace the system disk by using a public image

This topic describes how to replace the system disk by using a public image.

Scenarios

We recommend that you replace the system disk in the following scenarios:

- You selected an incorrect OS when creating an ECS instance.
- You need to replace the current OS.

To replace the system disk, you can replace the disk image with a public image, shared image, custom image, or any other image found on the Alibaba Cloud Marketplace. This topic uses a public image as an example. If you need to use an image that is not a public image, see [Replace the system disk \(non-public image\)](#).

After you replace the system disk:

- A new system disk with a new ID is allocated to your instance, and the original system disk is released.
- The cloud disk type remains unchanged.
- The IP address and the MAC address remain unchanged.

- We recommend that you [delete snapshots or automatic snapshot policies](#) to ensure that you have sufficient snapshots available for the automatic snapshot policies of the new system disk.

Precautions



Note:

Microsoft has ended extended technical support for Windows Server 2003. For data security purposes, we recommend that you discontinue running Windows Server 2003 on your ECS instances, and update the OS running on your instances. Alibaba Cloud no longer provides an image of this OS. For more information, see [Offline announcement of Windows Server 2003 system image](#).

Risks

Replacing the system disk may involve the following risks:

- The original system disk will be released. Therefore, we recommend that you [create a snapshot](#) to back up your data before replacing the system disk.
- Your instance will be stopped and your services interrupted.
- You must redeploy the service environment on the new system disk. However, this may result in an interruption to your services.
- The disk ID will be changed. Therefore, snapshots of the original system disk cannot be used to roll back the new system disk.



Note:

After you replace the system disk, the snapshots you have manually created are not affected. You can still use them to create custom images.

If you have configured automatic snapshot policies for the original system disk to allow automatic snapshots to be released along with the disk, the snapshot policies will no longer apply. All automatic snapshots of the original system disk will be automatically deleted.

Precautions during cross-OS disk replacements

Cross-OS disk replacements refer to replacing the system disk from one OS to another , specifically a switch between Linux and Windows.



Note:

Regions outside Mainland China do not support disk replacements between Linux and Windows, but they do support disk replacements between different Linux or Windows editions.

During cross-OS disk replacements, the file format of the data disk may be unidentifiable.

- In case no important data exists on the data disk, we recommend that you [reinitialize the disk](#) and format it to the default file system of your OS.
- In case important data exists in your data disk, perform the following operations as required:
 - For switches from Windows to Linux, you must install a software application, such as NTFS-3G, because Linux cannot identify NTFS.
 - For switches from Linux to Windows, you also must install a software application, such as Ext2Read or Ext2Fsd, because Windows cannot identify ext3, ext4, or XFS.

Precautions during Windows OS replacements

- Windows OS only supports password authentication.



Note:

Linux OS supports password authentication and SSH key pair authentication. As a result, when you replace Windows OS with Linux OS on your instance, you will have more optional authentication methods.

- If you are using a non-I/O-optimized instance, you can only replace your OS with the following Windows Server OSs by calling [ReplaceSystemDisk](#).

OS version	Image ID
Windows Server 2008 R2 Enterprise Edition (English)	win2008r2_64_ent_sp1_en-us_40G_alibase_20170915.vhd
Windows Server 2008 R2 Enterprise Edition (Chinese)	win2008r2_64_ent_sp1_zh-cn_40G_alibase_20170915.vhd
Windows Server 2012 R2 Datacenter Edition (Chinese)	win2012r2_64_dtc_17196_zh-cn_40G_alibase_20170915.vhd
Windows Server 2012 R2 Datacenter Edition (English)	win2012r2_64_dtc_17196_en-us_40G_alibase_20170915.vhd

OS version	Image ID
Windows Server 2016 Datacenter Edition (Chinese)	win2016_64_dtc_1607_zh-cn_40G_alibase_20170915.vhd
Windows Server 2016 Datacenter Edition (English)	win2016_64_dtc_1607_en-us_40G_alibase_20170915.vhd

Preparations

- Make sure that there is sufficient system disk space. We recommend that you reserve 1 GiB. Otherwise, the OS may not properly start after system disk replacement.
- If you want to replace the OS to Linux and use an SSH key pair for authentication, you need to [create an SSH key pair](#) first.
- System disk replacement may lead to data loss or service interruption. To minimize the impact on your services, we recommend that you [create a snapshot](#) for the original system disk during off-peak hours before starting the replacement process.



Note:

Creating a snapshot of 40 GiB takes about 40 minutes the first time. Creating a snapshot may also reduce I/O performance of a block storage device by up to 10%.

Procedure

To replace the system disk, follow these steps:

1. In the Actions column of the target instance, choose More > Instance Status > Stop and stop the instance as prompted.



Note:

If the instance is a Pay-As-You-Go VPC instance with the No Fees for Stopped Instances function enabled, the instance may not be properly started after system disk replacement. We recommend that you disable this function when you stop the instance.

2. From the Actions column, choose More > Disk and Image > Replace System Disk.
3. In the displayed dialog box, read the precautionary statement about system disk replacement and then click OK.

4. On the Replace System Disk page, set the following parameters:
- a. Image Type: Select Public Image and then select the image version.
 - b. System Disk: Unchangeable. However, you can expand the disk space to meet the requirements of your system disk and services. The maximum disk space is 500 GiB. The minimum disk space you can configure depends on the current disk space and image type.

Image	Space range (GiB)
Linux (excluding CoreOS) + FreeBSD	[Max{20, current space of the system disk}, 500]
CoreOS	[Max{30, current space of the system disk}, 500]
Windows	[Max{40, current space of the system disk}, 500]



Note:

If you have renewed your instance and downgraded the configurations, the system disk space cannot be changed until the next billing cycle starts.

c. Security enhancement:

- If the new OS is Windows, you can only use password authentication.

Image Type:

Public Image Custom Image Shared Image Marketplace Image

Public Image:

Windows Server Version 1709 DataCenter Edition 64bit Chinese ... Selection advice >

☒ Security enhancement ⓘ

System Disk:

Ultra Cloud Disk 40 GB 2120 IOPS The default system disk device name : /dev/xvda

To learn how to select SSD cloud disks, ultra cloud disks, and basic cloud disks, [Learn More >](#)

Login name:
administrator

Login password: It must be 8 - 30 characters long and contain three types of

Confirm password:

- If the instance is an I/O optimized instance and the new OS is Linux, you can use a password or an SSH key pair for authentication. In this case, we recommend you set a logon password or bind an SSH key pair.

Image Type:

Public Image Custom Image Shared Image Marketplace Image

Public Image:

CentOS 7.4 64bit Selection advice >

☒ Security enhancement ?

System Disk:

Ultra Cloud Disk 40 GB 2120 IOPS The default system disk device name : /dev/xvda

To learn how to select SSD cloud disks, ultra cloud disks, and basic cloud disks, [Learn More >](#)

Security:

Key Pair Password Set Now

A key pair includes a public key and a private key. Currently only I/O-optimized instances support the use of key pair. Using a key pair, you cannot log on with a user name and password.

Key Pair:

Select the Key Pair ↕

Also, you can go to the console to [create an access key >](#)

d. Confirm Instance Cost, which includes the image fee and system disk fee. For more information, see the [pricing](#) page of ECS.

e. Check your settings and click Confirm to change.

Replacing the OS and updating the system status take about 10 minutes. After the OS is replaced, the instance automatically starts.

What to do next

After replacing the system disk, perform the following operations as needed:

- (Optional) [Apply automatic snapshot policies to disks](#). Automatic snapshot policies are bound to the disk ID. After the system disk is replaced, automatic snapshot policies applied on the original disk will fail automatically. Therefore, you need to configure automatic snapshot policies for the new system disk.

- Write the new partition information to the `/etc/fstab` file of the new system disk and mount the partition as follows:



Note:

If the OS before and after disk replacement is Linux, and if a data disk is mounted to the instance and the partition is set to be mounted automatically at instance startup, then all mounting information will be lost. For more information, see [Format and mount data disks for Linux instances](#).

1. Back up the `/etc/fstab` file.
2. Write information about the new partition to the `/etc/fstab` file.
3. Check the information in the `/etc/fstab` file.
4. Run the `mount` command to mount the partition.
5. Run the `df -h -h` command to check the file system space and usage.

After the data partition is mounted, the data disk is ready for use without the need for instance restart.

Related API

[ReplaceSystemDisk](#)

3.10 Roll back a cloud disk

If you have [created snapshots](#) for a cloud disk, you can roll the disk back to restore a cloud disk to a specific snapshot status at a given time point.

Limits

Before you roll back a cloud disk, note the following:

- Rolling back a cloud disk is an irreversible action. Once rollback is complete, data cannot be restored. Exercise caution when performing this action.
- After the disk is rolled back, data from the creation date of the snapshot to the rollback date is lost.
- After a system disk is restored, the logon password or the SSH key pair of the ECS instance is retained.

Prerequisites

Before rolling back a cloud disk, check that:

- You have [created a snapshot](#) for the cloud disk, and no snapshot creation is in progress.
- You have not released the cloud disk.
- The cloud disk has been [attached to an ECS instance](#) and the instance is in the [Stopped](#) state.

**Note:**

For a Pay-As-You-Go VPC-Connected ECS instance, if the [No fees for stopped instances \(VPC-Connected\)](#) feature is enabled, to stop an instance, in the Notice dialog box, click OK. Then in the Stop dialog box, select Keep Instance with Fees, and click OK. If you use the No fees for stopped instances (VPC-Connected) feature, you may not be able to start the instance successfully after changing the system disk.

Stop Instance

The operation will be performed on the selected 1 Instances. Are you sure you want to proceed?

Stop Mode:

- ☒ Stop
- ☐ Force Stop

Stopped By:

- ☒ Retain Instance and Continue Charging After Instance Is Stopped
- ☐ No Charges After Instance Is Stopped

After a subscription instance is stopped, its expiration time does not change.

If you need to stop an instance for system disk replacement, disk reinitialization, instance upgrade, or private IP address modification, we recommend that you select Keep Stopped Instances and Continue Billing to avoid startup failure.

OK Cancel

Procedure

To roll back a cloud disk, follow these steps:

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Instances & Images > Instances.
3. In the top navigation bar, select a region.
4. Find the target instance and click its ID to go to the Instance Details page.
5. In the left-side navigation pane, click Instance Snapshots.
6. Find the target snapshot and then in the Actions column, click Roll Back Disk.
7. In the dialog box, click OK.

**Note:**

If you select Start Instance after Rollback, the instance starts automatically after the disk is restored.

Related API

[ResetDisk](#)

What to do next

If you resize a cloud disk after creating a snapshot, you can connect to the instance to resize its file system. For more information, see:

- [Linux _ Resize a data disk](#)
- [Windows _ Resize a data disk](#)

3.11 Reinitialize a cloud disk

3.11.1 Reinitialize a cloud disk

When a cloud disk is attached to an ECS instance, you can reinitialize the disk to restore the system disk or the data disks to the status when they were created. After a cloud disk is reinitialized:

- The system disk is restored to the initial status when it was created. For example, if you select a public image to create an ECS instance, after the system disk is reinitialized, the operating system is retained, but all other applications that were installed after the instance creation are deleted.

**Note:**

After you change the operating system or resize the system disk, the instance is not fully restored to the status at which it was created, but only to the status of the new system disk when it was created.

- Depending on how the data disk was created, it is restored to the following initial status:
 - Restored to an empty disk if it was an empty disk
 - Restored to a disk with only the data of the source snapshot if it was [created from a snapshot](#)
- If an automatic snapshot policy is applied to a cloud disk, the policy is retained and does not need to be applied again after reinitialization.
- After a cloud disk is reinitialized, all the snapshots, both automatically and manually created, are retained. You can use them to [roll back a cloud disk](#).



Warning:

- Because you must stop your ECS instance to reinitialize a cloud disk, your business services may be disrupted. Exercise caution when performing this action
- After a cloud disk is reinitialized, its data is lost. Therefore, we recommend you back up the data. To do so, you can [create snapshots](#).

Reinitialize a system disk

Prerequisites

If an SSH key pair is used as the authentication method, check that you have [created an SSH key pair](#) or [imported an SSH key pair](#).

Procedure

To reinitialize a system disk, follow these steps:

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Instances & Images > Instances.
3. In the top navigation bar, select a region.
4. Find the target ECS instance and click its ID to go to the Instance Details page.
5. Click Stop.



Note:

For a Pay-As-You-Go VPC-Connected ECS instance, if the [#unique_117](#) feature is enabled, in the Notice dialog box, click OK, and then in the Stop Instance dialog box, select Retain Instance and Continue Charging After Instance Is Stopped. If you select the No Charges After Instance Is Stopped, you may not be able to start the instance successfully after you reinitialize the system disk.

Stop Instance

The operation will be performed on the selected 1 Instances. Are you sure you want to proceed?

Stop Mode:

- ☒ Stop
- ☐ Force Stop

Stopped By:

- ☒ Retain Instance and Continue Charging After Instance Is Stopped
- ☐ No Charges After Instance Is Stopped *i*

After a subscription instance is stopped, its expiration time does not change.

If you need to stop an instance for system disk replacement, disk reinitialization, instance upgrade, or private IP address modification, we recommend that you select Keep Stopped Instances and Continue Billing to avoid startup failure.

OK Cancel

6. After the instance is stopped, click Disks in the left-side navigation pane.
7. Find the system disk and then, in the Actions column, click Reinitialize Disk.

8. In the Reinitialize Disk dialog box, complete the following configuration:

a. Authentication method:

- For a Windows instance, you must specify a logon password. You can either use a previous password or specify a new one.

Reinitialize Disk

Are you sure you want to reinitialize the following disks:

System Disk: [REDACTED]

Security: ☐ Set SSH Key ☒ Set Password

*Logon Password:

The password can be 8 to 30 characters in length and must contain three types of the following characters: Uppercase letters, lowercase letters, numbers, and special characters. Special characters include parentheses (()), graves (`), tildes (~), exclamation points (!), at signs (@), number signs (#), dollar signs (\$), percent signs (%), carets (^), ampersands (&), asterisks (*), hyphens (-), plus signs (+), equal signs (=), vertical bars (|), curly braces ({ }), braces ([]), colons (:), semicolons (;), apostrophes ('), angle brackets (< >), commas (,), periods (.), question marks (?), and forward slashes (/).

*Confirm Password:

Security Enhancement: ☒ Activate ⓘ

Instance Startup Policy: ☒ Start Instance after Resetting Disk

Note: 1. After reinitialization, the selected system disk will be restored to its image.
2. You must reset the password for logging on to the instance when reinitializing the disk.

Note: Your automatic snapshot policy will become invalid soon and must be reconfigured. Alibaba Cloud is not responsible for any data loss caused by failure to reconfigure the policy.

Confirm Cancel

- For a Linux instance, select Set SSH Key or Set Password as the security setting. If Key Pair is selected, bind a key pair. If Password is selected, specify a logon password.

Reinitialize Disk

Are you sure you want to reinitialize the following disks:

System Disk: [Redacted]

Security: ☒ Set SSH Key ☐ Set Password

SSH Key Pair: Search by SSH key pair name

Security Enhancement: ☒ Activate

Instance Startup Policy: ☒ Start Instance after Resetting Disk

Note: 1. After reinitialization, the selected system disk will be restored to its image.
2. You must reset the password for logging on to the instance when reinitializing the disk.

Note: Your automatic snapshot policy will become invalid soon and must be reconfigured. Alibaba Cloud is not responsible for any data loss caused by failure to reconfigure the policy.

Confirm Cancel

- (Optional) Security Enhancement: Select Activate. After the security enhancement feature is enabled, ECS security components are loaded. These components provide security features such as backdoor detection, remote logon reminders, brute-force cracking prevention mechanisms, and more.
 - (Optional) Instance Startup: Select Start Instance Resetting Disk. .
 - Click Confirm.
9. For Linux instances: If you have attached a data disk to the instance, connect to the instance and [create a mounting point for the partitions of data disks](#), because the mounting points are lost after the system disk is reinitialized.



Note:

For a Windows instance, both the system disk and the data disks are ready for use. No additional operations are needed.

After the system disk is reinitialized, you must deploy all applications to restore your business operations.

Reinitialize a data disk

Once reinitialized, a data disk is in a different status according to its original status and the operating system of the instance:

- For a Windows instance, the data disk is ready to use without any additional operations required.
- For a Linux instance:
 - If the data disk was an empty disk after it was created, then all the data and partitions on the disk are lost. You must partition and format the disk, and mount the partitions again.



Note:

If you configured the `/etc/fstab` file to automatically mount the disk partitions at startup of the instance, you must comment out the lines from the `/etc/fstab` file before reinitializing a data disk. Otherwise, your instance will fail to start.

- If the data disk was created from a snapshot, then the data disk is recovered to the point in time at which the snapshot was generated. You do not have to mount the partitions again, but all the data generated after the disk creation is lost.

In this section, `/dev/vdb1` is the example partition and `/InitTest` is the example mounting point. Replace these details with your actual information.

Prerequisites

The data disk to be reinitialized must be attached to an ECS instance. For more information, see [Attach a cloud disk](#).

Procedure

To reinitialize a data disk, follow these steps:

1. For Linux instances: If the data disk was an empty disk after it was created, and the mounting configuration was added to the `/etc/fstab` file, you must comment

out the mounting configuration from the `/ etc / fstab` file. To do so, follow these steps:

- a. [Connect to the Linux instance](#).
- b. Run `vim / etc / fstab`.
- c. Press the `i` key to enter the Insert mode.
- d. Locate the mounting configuration lines and type `#` before the lines. For example:

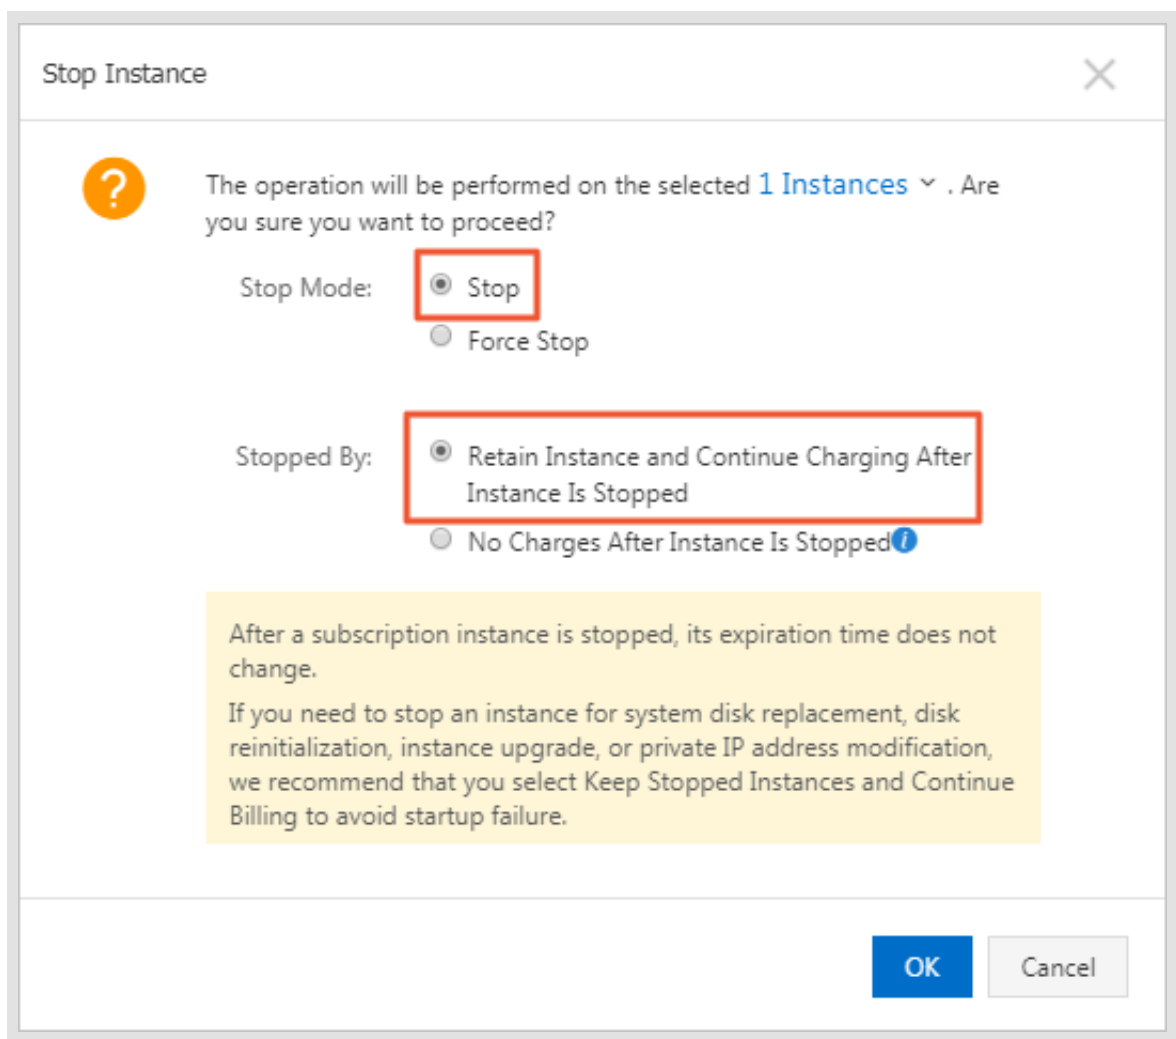
```
# / dev / vdb1 / InitTest ext3 defaults 0 0
```

- e. Press the `Esc` key to exit the Insert mode, and then run `:wq` to save and exit.
2. Log on to the [ECS console](#).
3. In the left-side navigation pane, choose Instances & Images > Instances.
4. In the top navigation bar, select a region.
5. Find the target ECS instance and click its ID to go to the Instance Details page.
6. Click Stop.



Note:

For a Pay-As-You-Go VPC-Connected ECS instance, if the [#unique_117](#) feature is enabled, in the Notice dialog box, click OK, and then in the Stop Instance dialog box, select Retain Instance and Continue Charging After Instance Is Stopped. If you select the No Charges After Instance Is Stopped, you may not be able to start the instance successfully after you reinitialize the system disk.



7. After the instance is Stopped, click Disks in the left-side navigation pane.
8. Find the target data disk and in the Actions column, click Reinitialize Disk.
9. In the Reinitialize Disk dialog box, read the notes and click Confirm.
10. In the left-side navigation pane, click Instance Details.
11. Click Start.
12. For Linux instances: If the data disk was an empty disk after it was created, [format and mount data disks for Linux instances](#).

After the data disk is reinitialized, you may need to deploy applications to restore your business operations.

Related API

[ReInitDisk](#)

3.11.2 Reinitialize a data disk

When a data disk is attached to an Elastic Compute Service (ECS) instance, you can reinitialize the disk to restore it to the status when it was created.

Prerequisites

- The data disk has been attached to an instance. For more information about how to attach a data disk to an instance, see [#unique_18](#).
- For an instance that runs Linux, you can add a command in the `/etc/fstab` file to mount partitions of a data disk at the startup of the instance. If the data disk was empty when it was created, after you reinitialize the data disk, the command will not be executed and the instance cannot start up as expected. We recommend that you comment out the command in the `/etc/fstab` file. The procedure is as follows:

1. [Remotely connect to an instance that runs Linux](#).
2. Run the `vim /etc/fstab` command.
3. Press the `i` key to enter the editing mode.
4. Find the command used to mount data disk partitions and comment it out by using `#`, as shown in the following line.

```
# / dev / vdb1 / InitTest ext3 defaults 0 0
```



Note:

`/ dev / vdb1` is an example partition and `/ InitTest` is an example mount point. You can replace them based on your business requirements.

5. Press the `Esc` key to exit the editing mode. Then, enter `:wq` to save any changes and exit the vim editor.

Context

The status of a data disk after being reinitialized varies depending on its original status when it was created and the operating system the instance runs:

- The data disk is restored to the initial status when it was created:
 - It becomes an empty disk if it was originally an empty disk.
 - It stores the data recorded in the source snapshot if it was created from a snapshot.

- For an instance that runs Windows, after you reinitialize its data disk, the data disk is ready for use without any additional operations regardless of its original status.
- For an instance that runs Linux:
 - If the data disk was created from a snapshot, it stores only the data recorded in the source snapshot after being reinitialized. You do not need to mount the partitions again, but all the data generated after the disk creation is lost.
 - If the data disk was created as an empty disk, all the data and partitions on the disk are lost. You must partition and format the disk, and mount the partitions again.

Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Instances & Images > Instances.
3. In the top navigation bar, select a region.
4. Find the ECS instance where you need to reinitialize a data disk, and click the instance ID to go to the Instance Details page.
5. Click Stop in the upper-right corner of the page to stop the instance.
6. After the instance is Stopped, click Disks in the left-side navigation pane.
7. Find the target data disk, and click Reinitialize Disk in the Actions column.
8. In the Reinitialize Disk dialog box, read the notes and click Confirm.
9. In the left-side navigation pane, click Instance Details.
10. Click Start in the upper-right corner of the page to start the instance and complete the reinitialization of the data disk.

What's next

- If a data disk is mounted to an instance running Linux and the data disk was created as an empty disk, you must format the disk after you reinitialize it. For more information, see [#unique_19](#).
- After the data disk is reinitialized, you need to deploy and configure applications to restore your business as soon as possible.

More information

[Reinitialize a cloud disk](#)

More information

[#unique_123](#)

3.12 Convert billing methods of cloud disks

The billing method of a cloud disk depends on how it is created:

- For cloud disks created with Subscription instances, prepayment of the service fee is required for it to be available for use. For more information, see [Subscription](#).
- For cloud disks created jointly with Pay-As-You-Go instances, or created separately the billing is on a Pay-As-You-Go basis. For more information, see [Pay-As-You-Go](#).

You can change the billing method of a cloud disk, as shown in the following table.

Conversion of billing methods	Conversion method	Suitable for	Effective date
Subscription —> Pay-As-You-Go	Renew for configuration downgrade	Subscription cloud disks attached to Subscription instances. The billing method of the system disk cannot be changed.	Effective from the next billing cycle
Pay-As-You-Go —> Subscription	Upgrade configurations	Pay-As-You-Go data disks attached to Subscription instances. The billing method of the system disk cannot be changed.	Effective immediately
	Switch from Pay-As-You-Go to subscription	The system disks and data disks attached to the Pay-As-You-Go instances.	

3.13 Modify the performance level of an ESSD cloud disk

This topic describes how to modify the performance level of a running ESSD cloud disk in the ECS console.

Prerequisites

You can follow the steps in this guide if your ESSD cloud disk meets the following requirements:

- Your account does not have overdue payment.
- If you attach an ESSD cloud disk to a Pay-As-You-Go ECS instance, make sure that the instance is not in the Expired state.
- You can change the performance level of a new ESSD cloud disk only after it enters the Unattached (Available) state.

Context

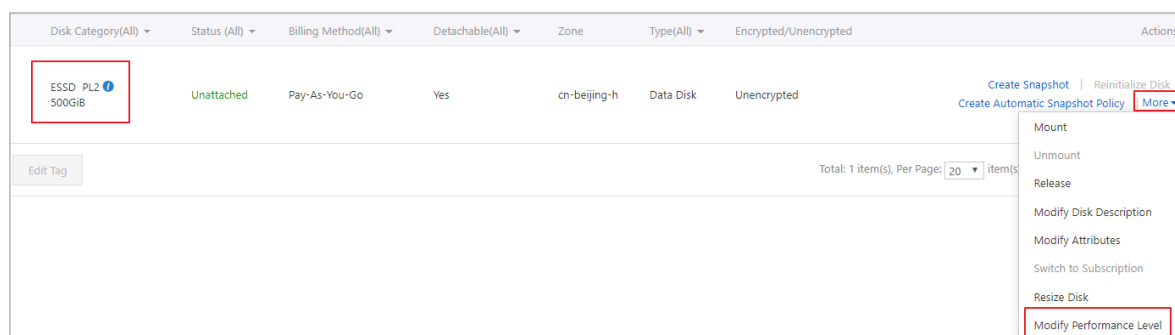
When you create an ECS instance, you can set the ESSD cloud disk as the system disk or data disk. You can also create a separate ESSD cloud disk. For information about how to create an ESSD cloud disk, see [Create an instance by using the wizard](#) and [Create a Pay-As-You-Go cloud disk](#) . For information about ESSD cloud disks, see [ESSD cloud disk](#).

After the performance level is modified, the ESSD cloud disk is billed according to the new performance level.

You can also call the API action [#unique_126](#) to complete this task.

Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
3. In the top navigation bar, select a region.
4. Find the target ESSD cloud disk. In the Actions column, choose More > Modify Performance Level.



5. In the Modify Performance Level dialog box, select a higher performance level and click OK.

Modify Performance Level

Disk ID: d-2zeill1kchndcr9fo7agpb

Name: -

*Performance Level: PL2: indicates that a single ESSD can deliver up to 100,000 random read/write IOPS ▼

Price:

PL1: indicates that a single ESSD can deliver up to 50,000 random read/write IOPS

PL2: indicates that a single ESSD can deliver up to 100,000 random read/write IOPS

PL3: indicates that a single ESSD can deliver up to 1,000,000 random read/write IOPS

OK Cancel


The performance level you can select for an ESSD cloud disk is determined by its storage capacity. If the performance level of an ESSD cloud disk cannot be upgraded, [resize the ESSD cloud disk](#) and then modify its performance level.


3.14 Monitor a cloud disk


You can monitor the IOPS and throughput of a cloud disk in the ECS console or, if you have installed the CloudMonitor agent, you can monitor the disk in the CloudMonitor console.



To monitor the IOPS and throughput of a cloud disk in the ECS console, follow these steps:


1. Find a cloud disk and click its ID to go to the Details page.
2. In the left-side navigation pane, click Disk Monitoring.



3. On the Monitoring Information page, click the  icon and set the start time and end time for monitoring information. You can check the monitoring information of a cloud disk for up to 15 days.

2018-06-05 22:38 - 2018-06-05 23:38 

Start Time : 2018-06-05 

22 : 38 


End Time : 2018-06-05 

23 : 38 


1Hour(s) 6Hour(s) 1Day(s) 7Day(s)

Maximum interval of 15 days


OK

4. View the IOPS and throughput of the cloud disk.



Note:

Click a legend in the chart to view a single performance index of a cloud disk.



3.15 Detach a cloud disk

If a Pay-As-You-Go cloud disk is attached to an ECS instance as a data disk, you can detach it from the instance and release it. However, if the disk is used as a system disk, you cannot detach it.

When detaching a cloud disk, consider the following:

- Only the Pay-As-You-Go cloud disks in the In Use state and used as a Data Disk can be detached.
- You cannot detach a local disk.

- For a Windows instance:
 - To guarantee data integrity, we recommend that you stop writing or reading the files on the cloud disk. Otherwise, data may be lost.
 - Before detaching a cloud disk in the ECS console, you must [connect to the instance](#) and set its status to Offline in Disk Management.
- For a Linux instance:
 - Before detaching a cloud disk in the ECS console, you must [connect to the instance](#) and run `umount` to unmount the partitions.
 - If you have configured the `/etc/fstab` file to automatically mount the partitions at the startup of the instance, before detaching it, you must delete the configurations from the `/etc/fstab` file. Otherwise, you cannot connect to the instance after the instance is restarted.

The following table describes the actions available for you to detach a cloud disk in the ECS console.

Scenario	Action
You want to detach one or more cloud disks from one instance.	Detach cloud disks on the Instance Disk page.
You want to detach one specific cloud disk.	Detach a cloud disks on the Disk List page.

Detach cloud disks on the Instance Disk page

On the Instance Disks page, you can delete one or more cloud disks that are attached to the instance.

Prerequisites

The cloud disks are [attached to the instance](#) and its status is In Use.

If you are detaching a cloud disk from a Linux instance, and you have configured the `/etc/fstab` file to mount the partitions at the startup of the instance, you must first delete the configurations.

Procedure

To detach a cloud disk from the Instance Disks page, follow these steps:

1. Connect to the instance and unmount its partitions. According to the operating system, follow the recommended steps detailed in the following table.

Operating system	Steps
Linux	Run <code>umount [partition]</code> . For example, <code>umount / dev / vdb1</code> .
Windows	Start Disk Management, right-click the disk name (for example, Disk 2) and then click Offline.

2. Log on to the [ECS console](#).
3. In the left-side navigation pane, choose Instances & Images > Instances.
4. In the top navigation bar, select a region.
5. Find the target instance and click its ID to go to the Instance Details page.
6. In the left-side navigation pane, click Disks.
7. Find the target cloud disk and then, in the Actions column, select More > Unmount.

Only cloud disks that have the following attributes can be detached:

- Status must be In Use.
 - Detachable must be Yes.
 - Type must be Data Disk.
8. In the dialog box, click Confirm.
 9. Optional. If you want to detach multiple cloud disks, repeat steps 7 and 8 as required.

When the status of the cloud disk becomes Unattached, the disk is detached.

Detach a cloud disks on the Disks page

On the Disk List page, you can detach a specific cloud disk from an ECS instance.

Prerequisites

The cloud disk is [attached to the instance](#) and is in the In Use state.

If you are detaching a cloud disk from a Linux instance, and you have configured the `/ etc / fstab` file to mount the partitions at the startup of the instance, delete the configurations.

Procedure

To detach a cloud disk on the Disk List page, follow these steps:

1. Connect to the instance and unmount the partitions. According to the operating system, follow the recommended steps detailed in the following table.

Operating system	Steps
Linux	Run <code>umount [partition]</code> . For example, <code>umount / dev / vdb1</code> .
Windows	Start Disk Management, right-click the disk name (for example, Disk 2) and then click Offline.

2. Log on to the [ECS console](#).
3. In the left-side navigation pane, choose Storage & Snapshots > Disks.
4. In the top navigation bar, select a region.
5. Find the target cloud disk and then, in the Actions column, select More > Unmount.

Only cloud disks that have the following attributes can be detached:

- Status must be In Use.
- Detachable must be Yes.
- Type must be Data Disk.

6. In the dialog box, click Confirm.

When the status of the cloud disk becomes Unattached, the disk is detached.

Related API

[DetachDisk](#)

What to do next

If you no longer need the disk, you can [release it](#).

3.16 Release a data disk

We recommend you release a data disk when you no longer require it to avoid incurring excess fees. Releasing a data disk is a permanent, irreversible action and is irreversible. After its release, data on the cloud disk cannot be restored, exercise caution when performing this action.

Note

When releasing a cloud disk, note that:

- Only cloud disks that are in the Available state can be released independently. Other cloud disks, such as those used as system disks, or Subscription billed cloud disks used as data disks, can only be released together with ECS instances. If a cloud disk is in the In Use state, you must first detach it from the instance.
- By default, automatic snapshots are released together with their cloud disks. However, manually created snapshots are not. You can change the snapshot release configuration when attaching a cloud disk.

**Note:**

To make sure you have sufficient storage space for the automatic snapshots, we recommend that you release automatically or manually created snapshots that you no longer require. For more information about snapshot quota, see [#unique_41](#).

- You can have data backed up before releasing a cloud disk. For example, by creating a snapshot.

Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
3. In the top navigation bar, select a region.
4. Select the disk that you want to release and make sure that it is in the Unattached state. Then, in the Actions column, choose More > Release.
5. In the Release Disk dialog box, read the note and click Confirm Release.

Related API

[DeleteDisk](#)

4 Shared Block Storage

Shared Block Storage is a block-level data storage service that features high concurrency, high performance, and high reliability. It supports concurrent reads and writes on multiple ECS instances, and provides data reliability of up to 99.9999999%.

Benefits

Shared Block Storage stores three copies of each piece of data, and distributes the copies across different servers. This achieves 99.9999999% data reliability for ECS instances. Shared Block Storage automatically copies your data within a zone to ensure data availability and prevent service disruptions, for example, due to a hardware failure.

Scenarios

In a traditional cluster architecture, multiple computing nodes access the same copy of data to provide services. To prevent service disruptions due to single point of failures, you can use Shared Block Storage to ensure access to the data, achieving high availability. We recommend that you store business data in Shared Block Storage devices and use a cluster file system to manage these devices. Data consistency can be guaranteed between multiple front-end computing nodes during concurrent read/write operations.

The high availability architecture of Shared Block Storage is suitable for enterprise-level applications. It provides shared access to block storage devices in a shared-everything architecture, such as the high availability server cluster architecture and Oracle Real Application Clusters (RAC). Oracle RAC is common among government departments, enterprises, and financial customers.

Types

Shared Block Storage can be divided into the following types based on performance:

- Shared SSD Block Storage

Shared SSD Block Storage uses SSDs as storage media to provide stable and high-performance storage with enhanced random I/O and data reliability.

- Shared Ultra Block Storage

Shared Ultra Block Storage uses a combination of SSDs and HDDs as storage media.

Performance

For more information about the performance of Shared Block Storage, see [Block storage performance](#).

Pricing

Shared Block Storage is free and under public preview.

Operations

- You can create, attach, detach, and release Shared Block Storage devices in a similar manner as disks.
- Shared Block Storage devices can only be created separately and used only as data disks.
- A single Shared Block Storage device can be attached to a maximum of eight ECS instances in the same zone and region at the same time.
- When attached to an ECS instance, Shared Block Storage devices share the data disk quota with disks, that is, up to 16 data disks can be attached to each ECS instance.
- You can allocate the capacity of Shared Block Storage devices in a similar manner as hard disks. You can format and partition Shared Block Storage devices that are attached to an ECS instance and create a file system.

File system

Shared Block Storage devices are not pre-installed with a cluster file system. However, you can install cluster file systems such as Google File System (GFS) and General Parallel File System (GPFS) to manage these devices. If Oracle RAC is deployed, we recommend that you use Automatic Storage Management (ASM) to manage storage volumes and file systems.

We recommend that you do not use a traditional file system to manage Shared Block Storage devices because the following two errors may occur:

- Disk space allocation conflict

If a Shared Block Storage device is attached to multiple instances and one of these instances (for example, instance A) writes data to a file, instance A checks the file

system and available disk space. After the write operation is complete, the space allocation record is changed in instance A, but the records in other instances are not updated. In this case, when another instance (for example, instance B) tries to write data to the file, this instance may allocate the disk space that has been already allocated by instance A, resulting in a disk space allocation conflict.

- Data file inconsistency

After an instance (for example, instance A) reads and caches data, another process that accesses the same data will directly read the data from the cache. However, if a copy of the same data that is stored on another instance (for example, instance B) is modified during this period, and instance A does not update according to the latest data change, instance A still reads the data from the cache. Data inconsistency occurs as a result.

5 Local disks

This topic provides information about local disks that you can purchase along with instances.

Background information

Local disks are disks attached to a machine on which ECS instances are hosted, and are designed to provide local storage and access for instances. Features of local disks include low latency, high random IOPS, high throughput, and cost-effective performance. For more information, see [Instance type families](#). If you are using phased-out local SSDs, see [Ephemeral SSDs](#).

Disk types

Currently, Alibaba Cloud provides the following two local disks:

- **NVMe SSD:** A local disk type that can be used by such instance type families as i2, i1, and gn5. NVMe SSDs are suitable for the following scenarios:
 - Online services such as online gaming, e-commerce, live video, and media, where instance type families such as i1 and i2 meet low latency and high I/O performance required by I/O intensive applications in block storage.
 - Service scenarios with high requirements for I/O performance and high-availability architecture at the application layer, such as NoSQL non-relational databases, MPP data warehouses, and distributed file systems.
- **SATA HDD:** A local disk type that can be used by such instance type families as d1ne and d1, which are applicable to multiple industries with high requirements for big data computing, storage analysis, massive data storage, and offline computing (for example, e-commerce and financial industries). SATA HDDs are designed to meet the requirements of distributed computing service types (for example, Hadoop) for instance storage performance, storage capacity, and intranet bandwidth.

NVMe SSD performance

The following table describes the performance indexes of the NVMe SSD disk type.

NVMe SSD index	Single disk performance	Overall instance performance
Maximum capacity	1456 GiB	2,912 GiB
Maximum IOPS	240,000	480,000
Write IOPS *	min{165 * capacity, 240,000}	2 * min{165 * capacity, 240,000}
Read IOPS *		
Maximum read throughput	2 Gbit/s	4 Gbit/s
Read throughput *	min{1.4 * capacity, 2,000} Mbit/s	2 * min{1.4 * capacity, 2,000} Mbit/s
Maximum write throughput	1.2 Gbit/s	2.4 Gbit/s
Write throughput *	min{0.85 * capacity, 1,200} Mbit/s	2 * min{0.85 * capacity, 1,200} Mbit/s
Access latency	Microsecond (μs)	

* The performance of a single disk is calculated as follows:

- Write IOPS: 165 IOPS for each GiB, up to 240,000 IOPS
- Write throughput: 0.85 Mbit/s for each GiB, up to 1,200 Mbit/s

To obtain the standard performance data and measure the Quality of Service (QoS) of Alibaba Cloud local disks, you can test the bandwidth, IOPS, and latency of the NVMe SSD by using the method described in the [NVMe SSD performance test](#).

SATA HDD performance

The following table describes the performance indexes of the SATA HDD disk type.

SATA HDD index	Single disk performance	Overall instance performance
Maximum capacity	5,500 GiB	154,000 GiB
Maximum throughput	190 Mbit/s	5,320 Mbit/s
Access latency	Millisecond (ms)	

Billing

The fees billed for local disks are included in the fees billed for the instances to which local disks are attached. For more information about instance billing methods, see [Subscription](#) and [Pay-As-You-Go](#).

Limits

- A single point of failure (SPOF) may be evident at the application layer due to each local disk being attached to a single host machine, meaning data reliability is dependent on the reliability of the host machine. We recommend that you implement data redundancy at the application layer to guarantee data availability.



Warning:

When you use a local disk to store data, there is a risk of data loss (for example, if a hardware fault occurs to the host machine). Therefore, we recommend you do not use a local disk to store any data that you need to keep for a long term. If your applications have no data reliability architecture, we recommend that you use [cloud disks or Shared Block Storage](#) in your instances to increase data reliability.

- After you purchase an instance along with a local disk, you need to log on to the instance to [partition and format the local disk](#).
- Regarding disk operations, you are not able to:
 - Create a separate local disk.
 - Use a snapshot to create a local disk.
 - Attach a local disk.
 - Detach and release a local disk separately.
 - Resize a local disk.
 - Reinitialize a local disk.
 - Create a snapshot for a local disk.
 - Use a snapshot to roll back a local disk.

Disk initialization sequence


When you create an instance along with a local disk, all disks are initialized according to the following rules:

- Rule 1: If the specified image does not have a data disk snapshot, the local disk is initialized before the cloud disks created along with the instance.


- Rule 2: If the specified image has a data disk snapshot, the sequence of the data disks recorded in the snapshot is retained, subject to Rule 1.

For example, for a Linux image that includes a snapshot of two data disks, the disks are initialized according to the following sequence.

- If the original device names of the two data disks are `/dev/xvdb` and `/dev/xvdc` respectively, these names are allocated to the specified data disks in the image. The disk initialization starts from the system disk, then proceeds to data disk 1 specified in the image, data disk 2 specified in the image, local disk 1, local disk 2, cloud disk 1, cloud 2, and in such a sequence until initialization is completed.



- If the original device names of the two data disks are `/dev/xvdc` and `/dev/xvdd` respectively, these names are allocated to the specified data disks in the image. Then, the remaining device names are allocated to the local disks. The disk initialization starts from the system disk, then proceeds to local disk 1, data disk 1 specified in the image, data disk 2 specified in the image, local disk 2, cloud disk 1, cloud 2, and in such a sequence until initialization is completed.



Lifecycle

The lifecycle of a local disk is the same as that of the instance to which the local disk is attached. For more information, see [ECS instance lifecycle](#).

Impact of instance operations on the data stored on local disks

The following table describes the impact of instance operations on the data stored on local disks.

Instance operation	Data on the local disk	Local disk
Restart the operating system, restart the instance by using the ECS console , or forcibly restart the instance.	Retained	Retained
Shut down the operating system, stop the instance by using the ECS console , or forcibly stop the instance.	Retained	Retained
Automatically restore the instance.	Erased	Released
Release the instance.	Erased	Released
A Subscription instance is stopped upon expiry or your account has an overdue payment, but the instance is not released.	Retained	Retained
A Subscription instance is stopped upon expiry or your account has an overdue payment, and the instance is released.	Erased	Released

NVMe SSD performance test

We recommend that you use the fio tool to test the performance of a local disk for Linux instances and Windows instances. The following example describes how to test the block storage performance of a bare disk / dev / vdx . For a description of the test command parameters, see [performance indexes of block storage](#).



Warning:

A direct bare disk test destroys the file system structure. Therefore, you must back up your data before the test. Additionally, the write operation overwrites the data on local disks. We recommend that you only test local disk performance on newly purchased ECS instances. If you decide to test a bare disk, we recommend that you exercise caution when performing this action.

- NVMe SSD bandwidth performance test

- To test the read bandwidth, run the following command:

```
fio - direct = 1 - iodepth = 128 - rw = read - ioengine =  
libaio - bs = 128k - numjobs = 1 - time_based = 1 - runtime =  
1000 - group_repo rting - filename = / dev / vdx
```

- To test the write bandwidth, run the following command:

```
fio - direct = 1 - iodepth = 128 - rw = write - ioengine =  
libaio - bs = 128k - numjobs = 1 - time_based = 1 - runtime =  
1000 - group_repo rting - filename = / dev / vdx
```

- NVMe SSD IOPS performance test

- To test the random read IOPS, run the following command:

```
fio - direct = 1 - iodepth = 32 - rw = randread - ioengine =  
libaio - bs = 4k - numjobs = 4 - time_based = 1 - runtime =  
1000 - group_repo rting - filename = / dev / vdx
```

- To test the random write IOPS, run the following command:

```
fio - direct = 1 - iodepth = 32 - rw = randwrite - ioengine =  
libaio - bs = 4k - numjobs = 4 - time_based = 1 - runtime =  
1000 - group_repo rting - filename = / dev / vdx
```

- NVMe SSD latency performance test

- To test the random read latency, run the following command:

```
fio - direct = 1 - iodepth = 1 - rw = randread - ioengine =  
libaio - bs = 4k - numjobs = 1 - time_based = 1 - runtime =  
1000 - group_repo rting - filename = / dev / vdx
```

- To test the random write latency, run the following command:

```
fio - direct = 1 - iodepth = 1 - rw = randwrite - ioengine =  
libaio - bs = 4k - numjobs = 1 - time_based = 1 - runtime =  
1000 - group_repo rting - filename = / dev / vdx
```

- To test the sequential read latency, run the following command:

```
fio - direct = 1 - iodepth = 1 - rw = read - ioengine =  
libaio - bs = 4k - numjobs = 1 - time_based = 1 - runtime =  
1000 - group_repo rting - filename = / dev / vdx
```

- To test the sequential write latency, run the following command:

```
fio - direct = 1 - iodepth = 1 - rw = write - ioengine =  
libaio - bs = 4k - numjobs = 1 - time_based = 1 - runtime =  
1000 - group_repo rting - filename = / dev / vdx
```

6 Block Storage FAQ

- Enhanced SSD FAQ
 - [What is an enhanced SSD?](#)
 - [When are enhanced SSDs available for purchase?](#)
 - [What are the differences between the enhanced SSDs offered during and after the public preview?](#)
 - [What are the similarities and differences among enhanced SSDs, standard SSDs, and ultra disks?](#)
 - [How is the performance level of an enhanced SSD measured?](#)
 - [How do I test the performance of an enhanced SSD and obtain the one million IOPS result?](#)
 - [What is the relationship between the storage performance of an enhanced SSD and storage performance of the instance to which the enhanced SSD is attached?](#)
 - [How are enhanced SSDs billed?](#)
 - [In which regions and zones are enhanced SSDs available?](#)
 - [To which instance families can enhanced SSDs be attached?](#)
- Common FAQ
 - [What must I consider when I select zones to create disks and then attach the disks to ECS instances?](#)
 - [What are the common operations that can be performed on a disk?](#)
 - [What is I/O optimization? Can I upgrade an existing ECS instance to an I/O optimized instance?](#)
- Performance test FAQ
 - [What tools can I use to test the performance of Block Storage?](#)
 - [Why does my instance shut down when I use the FIO tool to test I/O performance?](#)

- **Standard SSD FAQ**

- [What is the I/O performance of a standard SSD?](#)
- [What scenarios are suitable for standard SSDs?](#)
- [Can I replace a basic disk with a standard SSD?](#)
- [How do I purchase a standard SSD? What are the pricing options for I/O optimized instances and standard SSDs?](#)
- [Can I upgrade a standard SSD after I purchase it?](#)
- [An error is reported when I attempt to attach a standard SSD to my I/O optimized Linux instance. Why?](#)
- [For Linux instances, what must I be aware of before I add mount information for a basic disk or a standard SSD?](#)

- **Shared Block Storage FAQ**

- [What is Shared Block Storage?](#)
- [When can I purchase Shared Block Storage?](#)
- [What scenarios are suitable for Shared Block Storage?](#)
- [Can I attach a Shared Block Storage device to ECS instances in different regions?](#)
- [How many Shared Block Storage devices can be attached to one ECS instance?](#)
- [What are the specifications and performance of Shared Block Storage?](#)
- [How do I test the performance of Shared Block Storage?](#)

- **FAQ about attaching and detaching disks**

- [What is a device name \(mount point\)?](#)
- [What is an independent disk?](#)
- [Can I attach one disk to multiple ECS instances?](#)
- [Do I need to partition and format a pay-as-you-go disk after I purchase it and attach it to an ECS instance?](#)
- [Why am I unable to find the data disk that I purchased for a Linux instance?](#)
- [How many disks can be attached to one ECS instance?](#)
- [Why am I unable to find my ECS instance when I attempt to attach a disk to it?](#)
- [Can I attach a disk to an ECS instance that is in a different zone from the disk?](#)
- [Will data in my data disk be lost when I detach the disk?](#)
- [Can I detach a system disk?](#)

- **Independent disk FAQ**
 - [How is an independent pay-as-you-go data disk billed?](#)
 - [A disk that I created separately was attached to my ECS instance. Why was the disk released with the instance?](#)
 - [Can I attach an independent pay-as-you-go data disk to a subscription ECS instance?](#)
 - [Can I detach a data disk from a subscription ECS instance?](#)
 - [I have upgraded or downgraded an instance during renewal. Can I switch a subscription disk to a pay-as-you-go disk in the instance within the instance renewal period?](#)
- **Disk snapshot FAQ**
 - [When I delete a disk, will its snapshots also be deleted?](#)
 - [Why are some automatic snapshots on my disk missing?](#)
 - [Can I use a snapshot to create an independent disk?](#)
- **FAQ about reinitializing disks**
 - [What can I do if I cannot access the data in a Linux data disk because an error occurred while attaching the data disk?](#)
 - [Are my snapshots retained if I reinitialize a disk?](#)
 - [What can I do if the data disk of my ECS Linux instance cannot be found after I restart the instance or reinitialize the system disk?](#)
 - [How do I re-mount data disks after I reinitialize the system disk of my ECS Linux instance?](#)

- **FAQ about resizing disks**
 - Are my snapshots retained if I replace a system disk?
 - What must I be aware of before I replace a system disk?
 - How do I resize a system disk?
 - Can I scale down a system disk after I scale it up?
 - What Block Storage devices can be resized when they are used as system disks, and do any regional limits apply to this operation?
 - Can the system disks of both subscription ECS instances and pay-as-you-go ECS instances be resized?
 - What is the storage capacity range of a system disk?
 - I changed the specifications of my ECS instance when renewing it. Can I specify a new system disk capacity when I replace the system disk?
 - How I can use a temporary pay-as-you-go disk created from a snapshot of a data disk to resize the data disk without data loss?
 - What can I do if the error message "Bad magic number in super-block while trying to open /dev/xvdb1" is returned when I resize a disk of an ECS Linux instance?
- **Partition FAQ**
 - Can I partition a data disk for data storage?
 - For a disk with multiple partitions, are snapshots created for the entire disk or only for a specific partition?
 - What must I be aware of before I re-partition a disk?
- **FAQ about rolling back disks**
 - I rolled back a data disk by using a snapshot after I re-partitioned the disk. How many partitions are available in the disk?
 - An error message similar to the following is returned when I attempt to roll back a disk: "A disk can be rolled back only when the associated instance has been stopped and the disk has no snapshots being created. If the operating system of an ECS instance has been replaced, the snapshot taken before the replacement cannot be used to roll back the system disk." What can I do?
- **Other FAQ**
 - How do I copy data across instances?
 - How do I test the performance of an enhanced SSD?

What is an enhanced SSD?

An enhanced SSD is an ultra-high performance disk provided by Alibaba Cloud. Enhanced SSDs use 25 GE networks and the remote direct memory access (RDMA) technology to deliver up to 1 million random IOPS with low one-way latency. For more information, see [ESSD cloud disk](#).

When are enhanced SSDs available for purchase?

Enhanced SSDs are commercially available since the end of June 2019. You can log on to the [ECS console](#) or call the [CreateDisk](#) operation to create enhanced SSDs.

What are the differences between the enhanced SSDs offered during and after the public preview?

Different specifications of enhanced SSDs are offered during and after the public preview. The commercially available specifications are described in the following table.

For ESSD cloud disks, larger capacity positively corresponds with higher data-processing performance. The capacity, IOPS, and throughput provided by ESSD cloud disks vary depending on their performance levels. The specific capacity and performance are detailed in the following table.

Performance level	Capacity (GiB)	Maximum IOPS	Maximum throughput (Mbit/s)
PL1	20-32,768	50,000	350
PL2	461-32,768	100,000	750
PL3	1,261-32,768	1,000,000	4,000

What are the similarities and differences among enhanced SSDs, standard SSDs, and ultra disks?

- **Similarities:** These three types of disks are all based on a distributed Block Storage architecture, provide high reliability and scalability, and support snapshots and data encryption.
- **Differences:** SSDs have the best performance of the three types of disks. For more information, see [ESSD cloud disk](#) and [#unique_6](#).

How is the performance level of an enhanced SSD measured?

The performance level of an enhanced SSD is proportional to its capacity. An enhanced SSD that has a larger capacity delivers higher performance. Compared with

standard SSDs, enhanced SSDs have much better performance. For more information, see [ESSD cloud disk](#).

How do I test the performance of an enhanced SSD and obtain the one million IOPS result?

We recommend that you use the FIO tool to test the performance of enhanced SSDs. For more information, see [How do I test the performance of an enhanced SSD?](#).

What is the relationship between the storage performance of an enhanced SSD and storage performance of the instance to which the enhanced SSD is attached?

Part of instance-level storage performance is proportional to the specifications that instance types have. The higher specifications an instance type has, the higher IOPS and throughput it can deliver.

For example, you created an ECS instance of storage optimized instance family g5se and attached enhanced SSDs to the instance:

- If the total storage performance of the enhanced SSDs does not exceed the maximum storage performance that the instance type can deliver, the total storage performance of the enhanced SSDs takes precedence.
- If the total storage performance of the enhanced SSDs exceeds the maximum storage performance that the instance type can deliver, the maximum storage performance that the instance type can deliver takes precedence.

For example, you created a 16-GiB instance of the ecs.g5se.xlarge type that can deliver up to 60,000 IOPS. If you attach an enhanced SSD to the instance and the enhanced SSD has a capacity of 2 TiB with a maximum IOPS of 101,800, the maximum IOPS of the instance is 60,000, instead of 101,800.

The following table describes the performance and specifications of the g5se instance family.

Instance type	vCPUs	Memory (GiB)	IOPS (thousand)	Throughput (Mbit/s)
ecs.g5se.large	2	8	30	118
ecs.g5se.xlarge	4	16	60	235
ecs.g5se.2xlarge	8	32	120	470
ecs.g5se.4xlarge	16	64	230	900

Instance type	vCPUs	Memory (GiB)	IOPS (thousand)	Throughput (Mbit/s)
ecs.g5se.6xlarge	24	96	340	1,350
ecs.g5se.8xlarge	32	128	450	1,800
ecs.g5se.16xlarge	64	256	900	3,600
ecs.g5se.18xlarge	70	336	1,000	4,000

How are enhanced SSDs billed?

Enhanced SSDs support both the subscription and pay-as-you-go billing methods. For more information, see [Prices](#).

In which regions and zones are enhanced SSDs available?

Enhanced SSDs are available in the following regions and zones:

- China Hangzhou Zone H, Zone I, and Zone G
- China Shanghai Zone E, Zone F, and Zone G
- China Beijing Zone H, Zone G, and Zone F
- China Zhangjiakou Zone A and Zone B
- China Shenzhen Zone E and Zone D
- China Chengdu Zone A
- China Hongkong Zone C
- India Mumbai B
- UK London Zone B
- Australia Sydney Zone B

To which instance families can enhanced SSDs be attached?

Enhanced SSDs can be attached to instances from the instance families that support 25 GE networks (g6, g5, ic5, c6, c5, r6, r5, and g5se), ECS Bare Metal Instance families (ebmhfg5, ebmc4, ebmg5, ebmg6v, ebmg6i, ebmc5s, ebmg5s, ebmr5s, and sccgn6), and enterprise-level heterogeneous computing instance families (vgn5i, gn6i, gn6v, gn5, gn5i, gn4, ga1, f1, and f3).

What tools can I use to test the performance of Block Storage?

For details, see [#unique_6](#).

Why does my instance shut down when I use the FIO tool to test I/O performance?

You can use the FIO tool to test I/O performance by bare disk partition or file system. However, if you test bare disk partitions, the metadata of file systems in the bare disk partitions may be damaged, leaving you unable to access the files in the partitions. This then causes the instance to shut down. This problem does not occur when you use the FIO tool to test I/O performance by file system.

What must I consider when I select zones to create disks and then attach the disks to ECS instances?

A pay-as-you-go disk can only be attached to an ECS instance in the same zone.

- For high-availability applications, we recommend that you create data disks and attach them to ECS instances in different zones.
- For low-latency applications, we recommend that you create data disks and attach them to ECS instances in the same zone.

What are the common operations that can be performed on a disk?

The following topics describe the common operations that you can perform on a disk:

- [#unique_18](#)
- [#unique_19](#)
- [#unique_51](#)
- [#unique_43](#)
- [Reinitialize a cloud disk](#)
- [#unique_112](#)
- [#unique_21](#)
- [Resize cloud disks online](#)
- [Resize cloud disks offline](#)
- [Resize partitions and file systems of Linux data disks](#)
- [#unique_61](#)

What is I/O optimization? Can I upgrade an existing ECS instance to an I/O optimized instance?

I/O optimization provides better storage performance for instances and disks. For example, you can optimize the storage performance of a standard SSD by attaching it to an I/O optimized instance.

You can call the [#unique_142](#) or [#unique_143](#) operation to upgrade your non-I/O optimized instances to I/O optimized instances.

What is the I/O performance of a standard SSD?

For details, see [#unique_6](#).

What scenarios are suitable for standard SSDs?

Standard SSDs offer high performance and high reliability. They are applicable to I/O-intensive applications, such as MySQL, SQL Server, Oracle, PostgreSQL, and other small and medium-sized relational databases. They are also applicable to small and medium-sized development and testing environments.

Can I replace a basic disk with a standard SSD?

No, standard SSDs cannot be used to replace basic disks, because standard SSDs use all-SSD storage media.

How do I purchase a standard SSD? What are the pricing options for I/O optimized instances and standard SSDs?

For the pricing details, see [Pricing](#).

Can I upgrade a standard SSD after I purchase it?

Yes, you can upgrade and resize it. For more information, see [#unique_26](#).

An error is reported when I attempt to attach a standard SSD to my I/O optimized Linux instance. Why?

In the Linux operating system, the mount point for a standard SSD is `/dev/vd*`, and the mount point for a basic disk is `/dev/xvd*`. If you specify the mount point in the format of `/dev/xvd*` in a command to mount a standard SSD, an error occurs. Specify the mount point in the format of `/dev/vd*` in the command to mount the standard SSD.

For Linux instances, what must I be aware of before I add mount information for a basic disk or a standard SSD?

When you attach a data disk to a Linux instance, you need to format and partition the disk. During this process, note that `/ dev / xvdb1` is the mount point for a basic disk, and `/ dev / vdb1` is the mount point for a standard or enhanced SSD. If you add invalid information, you will fail to mount the disk by using the `mount - a` command. To avoid this problem, perform the following steps:

1. Run the `fdisk - l` command to view the data disk information.
2. Check whether the information added to `/ etc / fstab` is valid.



Note:

Do not add duplicate mount information because this will cause a system startup failure.

3. Run the `vim` command to modify the `/ etc / fstab` file.
4. Comment out or delete invalid information, and add valid mount information.
5. Run the `mount - a` command to check whether the disk is mounted.

For more information, see [#unique_19](#).

What is Shared Block Storage?

Shared Block Storage is a block-level data storage service that features high concurrency, high performance, and high reliability. It supports concurrent reads from and writes to multiple ECS instances. For more information, see [Shared Block Storage](#).

Each Shared Block Storage device can be attached to a maximum of eight ECS instances in the same zone and region at the same time. If you want to attach a Shared Block Storage device to more than eight ECS instances, [submit a ticket](#) to raise the limit.

When can I purchase Shared Block Storage?

Shared Block Storage is under public preview. During the public preview, Shared Block Storage is provided for free in all regions. After the public preview ends, Shared Block Storage supports both the subscription and pay-as-you-go billing methods.

What scenarios are suitable for Shared Block Storage?

The high availability architecture of Shared Block Storage is suitable for enterprise-level applications. It provides shared access to block storage devices in a shared-everything architecture, such as the high availability server cluster architecture and Oracle Real Application Clusters (RAC). Oracle RAC is common among government departments, enterprises, and financial customers.

Can I attach a Shared Block Storage device to ECS instances in different regions?

No, you can attach a Shared Block Storage device only to ECS instances in the same zone and region.

How many Shared Block Storage devices can be attached to one ECS instance?

When used as data disks, Shared Block Storage devices share the data disk quota with disks. That is, a maximum of 16 data disks can be attached to an instance.

What are the specifications and performance of Shared Block Storage?

Shared SSD Block Storage and Shared Ultra Block Storage are available. For more information about their specifications and performance, see [Block Storage performance](#).

How do I test the performance of Shared Block Storage?

You can use the FIO tool to perform the performance stress test. During the test, the total iodepth value of all clients cannot exceed 384. For example, if the test is conducted simultaneously on four instances, we recommend that the iodepth value of each client does not exceed 96.

- When you perform the performance stress test on two ECS instances:

- Run the following command to test random write IOPS:

```
FIO - direct = 1 - iodepth = 128 - rw = randwrite - ioengine = libaio - bs = 4k - size = 1G - numjobs = 1 - runtime = 1000 - group_repo rting - filename = iotest - name = Rand_Write_Testing
```

- Run the following command to test random read IOPS:

```
FIO - direct = 1 - iodepth = 128 - rw = randread - ioengine = libaio - bs = 4k - size = 1G - numjobs = 1 - runtime = 1000
```

```
- group_repo rting - filename = iotest - name = Rand_Read_
Testing
```

- Run the following command to test write throughput:

```
FIO - direct = 1 - iodepth = 64 - rw = write - ioengine =
libaio - bs = 64k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Write_PPS_
Testing
```

- Run the following command to test read throughput:

```
FIO - direct = 1 - iodepth = 64 - rw = read - ioengine =
libaio - bs = 64k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Read_PPS_T
esting
```

- When you perform the performance stress test on four ECS instances:

- Run the following command to test random write IOPS:

```
FIO - direct = 1 - iodepth = 96 - rw = randwrite - ioengine =
libaio - bs = 4k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Rand_Write
_Testing
```

- Run the following command to test random read IOPS:

```
FIO - direct = 1 - iodepth = 96 - rw = randread - ioengine =
libaio - bs = 4k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Rand_Read_
Testing
```

- Run the following command to test write throughput:

```
FIO - direct = 1 - iodepth = 64 - rw = write - ioengine =
libaio - bs = 64k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Write_PPS_
Testing
```

- Run the following command to test read throughput:

```
FIO - direct = 1 - iodepth = 64 - rw = read - ioengine =
libaio - bs = 64k - size = 1G - numjobs = 1 - runtime = 1000
- group_repo rting - filename = iotest - name = Read_PPS_T
esting
```

What is a device name (mount point)?

A device name (mount point) is the location of an ECS disk on the disk controller bus. The selected device name matches the disk device number in Linux and matches the disk sequence number in the disk manager in Windows.

What is an independent disk?

An independent disk is a disk that you purchase separately. It can be attached to or detached from any ECS instance in the same zone. Before using an independent disk, you must attach it to an instance and partition and format the disk. For more information, see [Create a Pay-As-You-Go cloud disk](#).

Can I attach one disk to multiple ECS instances?

No, a disk can be attached only to one ECS instance in the same zone.

Do I need to partition and format a pay-as-you-go disk after I purchase it and attach it to an ECS instance?

If you purchase a pay-as-you-go disk separately, you must attach it to an ECS instance and then partition and format it. For more information, see [#unique_19](#) and [#unique_51](#).

Why am I unable to find the data disk that I purchased for a Linux instance?

If you purchase a pay-as-you-go disk separately, you must attach it to an ECS Linux instance and then partition and format the disk before using it. For more information, see [#unique_19](#) and [#unique_18](#).

How many disks can be attached to one ECS instance?

When used as data disks, Shared Block Storage devices and disks share the data disk quota. That is, up to 16 data disks can be attached to one ECS instance.

Why am I unable to find my ECS instance when I attempt to attach a disk to it?

Check whether your ECS instance has been released. If no, make sure that it is in the same zone and region as the disk.

Can I attach a disk to an ECS instance that is in a different zone from the disk?

No, you can attach a pay-as-you-go disk only to an ECS instance that is in the same zone as the disk.

Will data in my data disk be lost when I detach the disk?

- In Windows, we recommend that you stop the read and write operations on all file systems of the disk to ensure data integrity. Otherwise, the data being read or written will be lost.

- In Linux, you must log on to the ECS instance and run the `umount` command on the disk. After the command is executed, log on to the ECS console to detach the disk.

Can I detach a system disk?

No, you cannot detach a system disk but you can replace it. For more information, see [#unique_112](#).

How is an independent pay-as-you-go data disk billed?

A pay-as-you-go data disk is billed based on the space usage. Note that if your account balance is insufficient, the services of the data disk will be suspended.

A disk that I created separately was attached to my ECS instance. Why was the disk released with the instance?

This is because you have set the disk to be released along with the instance when you attach the disk to the instance. You can change this setting at any time in the ECS console or by calling the related API operation. For more information, see [#unique_18](#).

Can I attach an independent pay-as-you-go data disk to a subscription ECS instance?

Yes.

Can I detach a data disk from a subscription ECS instance?

No, you cannot detach a data disk from a subscription ECS instance. The data disk will be released along with the instance when the instance expires. If you want to release the data disk, switch it to a pay-as-you-go data disk, and then detach and release it. For information about how to switch the billing method of disks, see [#unique_29](#).

I have upgraded or downgraded an instance during renewal. Can I switch a subscription disk to a pay-as-you-go disk in the instance within the instance renewal period?

No, you can switch a subscription disk to a pay-as-you-go disk only after the instance renewal period.

When I delete a disk, will its snapshots also be deleted?

If you have enabled Delete Automatic Snapshots While Releasing Disk, the automatic snapshots of the disk will be deleted when you delete the disk. However, the manual snapshots are retained. You can change this setting at any time. For more information, see [#unique_109](#).

Why are some automatic snapshots on my disk missing?

When the number of snapshots reaches the upper limit, the earliest automatic snapshot are automatically deleted but manual snapshots are not affected.



Note:

The automatic snapshot policy that is applied to a disk can be executed only after the disk is attached to an instance.

Can I use a snapshot to create an independent disk?

Yes, you can use an existing snapshot to create an independent pay-as-you-go disk.

For more information, see [Create a cloud disk by using a snapshot](#).

What can I do if I cannot access the data in a Linux data disk because an error occurred while attaching the data disk?

Perform these steps to fix the error:

1. Find the data disk and check whether it is attached to the corresponding ECS instance by using either of the following methods:
 - Perform the check in the ECS console. For more information, see [#unique_145](#).
 - Log on to the instance, and run the `fdisk -l` command to check whether the data disk partition information is correct. Run the `df -h` and `mount | grep "< devpath >"` commands to view the mount information.
2. Run the `cat` command to view the `/etc/fstab` file and check whether you have mounted two disks to the same directory.
 - If two disks are mounted to the same directory, the second disk will replace the first disk, which causes data to become inaccessible. We recommend that you mount either of the disks to a different directory.
 - If they are mounted to different directories but the mount information shows that they are in the same directory, run the `ll` command to check whether a link exists between the two directories. If yes, we recommend that you run the `mkdir` command to create a new directory, and mount either of the disks to the new directory. Then check whether the data can be accessed.

What can I do if data is lost after I restart my ECS Linux instance?

- **Problem description:** All data in a directory such as `/alidata` is lost after you restart your ECS Linux instance.

- Cause: After you run the `df -h` command, the command output shows that no data disk partitions are mounted to the directory.
- Solution: This solution uses an I/O optimized instance as an example. If your ECS instance is a non-I/O optimized instance, enter the device name of a disk partition in the format of `/dev/xvd * 1` in the mount command.
 1. Run the `fdisk -l` command to view the data disk partitions that are not mounted.
 2. Run the `mount /dev/vdb1 /alidata` command to mount a data disk partition to the directory.
 3. Run the `df -h` command to check whether the data disk partition is mounted to the directory.
 4. (Optional) You can configure the disk partition to be automatically mounted during system startup in the `/etc/fstab` file to avoid this problem.

Are my snapshots retained if I reinitialize a disk?

Yes, both the manual snapshots and automatic snapshots of the disk will be retained.

What can I do if the data disk of my ECS Linux instance cannot be found after I restart the instance or reinitialize the system disk?

- Problem description: After restarting an ECS Linux instance or reinitializing the system disk, you log on to the instance and run the `df -h` command. The command output shows that no data disks are found.
- Cause:
 - Restarting an instance: Mount information was not written to the `/etc/fstab` file when data disks were mounted. As a result, data disks are not automatically mounted after the instance restarts.
 - Reinitializing the system disk: The `/etc/fstab` file is reset after the system disk is initialized. As a result, data disks are not automatically mounted during system startup.

• **Solution:**

1. Run the `mount / dev / xvdb1` command to re-mount the original data disk partition of the instance.
2. Run the `mount` command to check the format of files in the `/ dev / xvdb1` partition.
3. Assume that the format of files in the `/ dev / xvdb1` partition is `ext3`. Run the following command to write the partition mount information to the `/ etc / fstab` file:

```
echo '/ dev / xvdb1 / data ext3 defaults 0 0 ' >> / etc / fstab
```

4. Restart the instance in the ECS console.

How do I re-mount data disks after I reinitialize the system disk of my ECS Linux instance?

After you reinitialize the system disk of your ECS Linux instance, data in the data disks that are attached to the instance remains unchanged, but the mount information of the data disks is lost. As a result, all the data disks are unmounted.

Assume that the mount point is `/ dev / vdb1` , and the mount point is named `/ InitTest` before the system disk is reinitialized. Run the following commands to re-mount data disk partition `/dev/vdb1`:

1. Run the `mount` command to view the mount information of data disks.

The command output does not contain information of `/ dev / vdb1` .

2. Run the `fdisk - l` command to view data disk partitions.
3. Run the `cat / etc / fstab` command to view the original mount point name of `/ dev / vdb1` .
4. Run the `mkdir / InitTest` command to re-create the mount point for the data disk partition.

The new mount point name must be the same as the original name of mount point `/ dev / vdb1` .

5. Run the `mount / dev / vdb1 / InitTest` command to re-mount the data disk partition.
6. Run the `df - h` command to check whether the data disk partition is mounted.

7. Perform the following steps to check whether `/ dev / vdb1` can be mounted automatically:

- a. Run the `umount / dev / vdb1` command to unmount `/ dev / vdb1` .
- b. Run the `mount` command to check whether the partition is unmounted.

If yes, the command output does not contain information of `/ dev / vdb1` .

- c. Run the `mount - a` command to automatically mount `/ dev / vdb1` .
- d. Run the `mount` command to check whether the partition is automatically mounted.

If yes, the command output shows information of `/ dev / vdb1` .

Are my snapshots retained if I replace a system disk?

This depends on the creation method of snapshots. Manual snapshots will be retained , but automatic snapshots will be deleted if Delete Automatic Snapshots While Releasing Disk is enabled.



Note:

After a system disk is replaced, the disk ID changes. You cannot use the snapshots of the original system disk to roll back the new system disk.

What must I be aware of before I replace a system disk?

We recommend that you create snapshots of the current system disk before you replace it. Make sure that the current system disk has enough space available. The recommended available disk space is 1 GiB or larger. If disk space is insufficient, the instance may not start properly after you replace the system disk.

How do I resize a system disk?

You can resize a system disk in the ECS console or by calling the [#unique_146](#) operation.

Can I scale down a system disk after I scale it up?

No, you cannot scale down a system disk after you scale it up. We recommend that you scale the system disk up properly.

What Block Storage devices can be resized when they are used as system disks, and do any regional limits apply to this operation?

Ultra disks, standard SSDs, and enhanced SSDs can be resized when they are used as system disks. You can resize system disks in all regions.

Can the system disks of both subscription ECS instances and pay-as-you-go ECS instances be resized?

Yes, the system disks of both subscription ECS instances and pay-as-you-go ECS instances can be resized.

What is the storage capacity range of a system disk?

The capacity range of a system disk varies depending on the operating system. For more information, see [#unique_26](#).

I changed the specifications of my ECS instance when renewing it. Can I specify a new system disk capacity when I replace the system disk?

After you perform the [#unique_15](#) operation on your ECS instance, you can resize the system disk of the instance only after the next billing cycle starts.

How I can use a temporary pay-as-you-go disk created from a snapshot of a data disk to resize the data disk without data loss?

If a data disk cannot be resized without data loss due to a disk error, you can purchase a temporary pay-as-you-go disk to store data and then format the data disk. To do so, perform the following steps:

1. Create a snapshot of the current data disk (source data disk). For more information, see [#unique_55](#).
2. Go to the [disk purchase](#) page. Select the same region and zone as the ECS instance to purchase a pay-as-you-go disk. Click Create from Snapshot. In the dialog box that appears, select the snapshot created in the previous step.
3. Log on to the ECS console and then attach the new data disk you purchased to the ECS instance.
4. Log on to the ECS instance and run the `mount` command to mount the new data disk. For more information about how to mount a disk created from a snapshot, see [Create a cloud disk by using a snapshot](#).
5. Check whether files in the new data disk are the same as those in the source data disk.

6. Run the `fdisk` command to delete the original partition table. Then run commands such as `fdisk` and `mkfs . ext3` to re-partition and re-format the source data disk, resizing it to the target capacity. For more information, see [Resize partitions and file systems of Linux data disks](#).
7. Run the `cp -R` command to copy all the data in the new data disk back to the source data disk.

You can add the `--preserve = all` parameter to preserve the files' properties when copying the files.
8. Run the `umount` command to unmount the new data disk.
9. Log on to the ECS console, detach the new data disk from the ECS instance, and release the disk.

What can I do if the error message "Bad magic number in super-block while trying to open /dev/xvdb1" is returned when I resize a disk of an ECS Linux instance?

- **Problem description:** When you run the `e2fsck -f /dev/xvdb` command to resize a disk of an ECS Linux instance, the error message "Bad magic number in super-block while trying to open /dev/xvdb1" is returned.
- **Cause:** The disk to be resized is not partitioned.
- **Solution:** Run the `e2fsck -f /dev/xvdb` and `resize2fs /dev/xvdb` commands to resize the disk. Then run the `mount` command to mount the disk.

Can I partition a data disk for data storage?

Yes, you can divide a data disk into multiple partitions based on your needs. We recommend that you use the system tool for partitioning.

For a disk with multiple partitions, are snapshots created for the entire disk or only for a specific partition?

Snapshots are created for the entire disk, instead of for a specific partition.

What must I be aware of before I re-partition a disk?

To ensure data security, we recommend that you create a snapshot of a disk before re-partitioning the disk. This way, you can roll back the disk if an exception occurs. For more information, see [Create a snapshot](#) and [#unique_147](#).

I rolled back a data disk by using a snapshot after I re-partitioned the disk. How many partitions are available in the disk?

When you roll a data disk back by using a snapshot, the disk returns to the state it was in when the snapshot was taken. If the disk has not been re-partitioned at the time when the snapshot was taken, only one partition is available in the disk.

An error message similar to the following is returned when I attempt to roll back a disk: "A disk can be rolled back only when the associated instance has been stopped and the disk has no snapshots being created. If the operating system of an ECS instance has been replaced, the snapshot taken before the replacement cannot be used to roll back the system disk." What can I do?

- **Problem description:** When you attempt to roll back a disk by using a snapshot, an error message similar to the following is returned: "A disk can be rolled back only when the associated instance has been stopped and the disk has no snapshots being created. If the operating system of an ECS instance has been replaced, the snapshot taken before the replacement cannot be used to roll back the system disk ."
- **Cause:** The problem may be caused by incorrect disk properties or disk states.
- **Solution:** You can troubleshoot the problem based on the instance state or snapshot state.

- Check whether the instance to which the disk is attached has stopped.

You can only roll back disks attached to instances in the stopped state. You can log on to the ECS console and check the state of the instance on the Instances page.

- Check whether the system disk of the instance associated with the snapshot has been replaced.

After you select a new image for the instance, a new system disk will be automatically re-created from the new image for the instance and the system disk ID will change. Therefore, you cannot use the snapshots taken before the system disk is replaced to roll back the system disk. However, you can create a custom image from one of these snapshots and then use this image to replace the system disk of the instance. In this way, the instance returns to the state it

was in when the snapshot was taken. For more information, see [#unique_148](#) and [#unique_112](#).

- Check whether the disk to be rolled back has a snapshot being created.

To ensure data consistency, do not roll back a disk that has a snapshot being created. You can go to the Instance Details page. Click Instance Snapshots in the left-side navigation pane to go to the Snapshots page. Check the states of snapshots. A snapshot is being created if the Progress value is not 100% and the Status value is Progressing.

If you need to forcibly terminate the creation process of a snapshot to roll back the disk, select the snapshot and click Delete.

How do I copy data across instances?

You can choose data copy methods based on the operating system:

- Copy data between Linux instances

- Use the lrzsz tool

Log on to the Linux instances, install the lrzsz tool, run the `rz` command to upload files to one Linux instance, and run the `sz` command to download these files to the other Linux instance.

You can also first run the `sz` command to download files to your local PC and then run the `rz` command to upload these files to the other Linux instance.

- Use the FTP tool

If you use the SFTP tool, we recommend that you use the root account for instance logon and file upload and download.

- Use the wget command

After compressing a file or a folder, save it to the web directory to generate a download URL. Then run the `wget` command on the other Linux instance to download the file or folder.

- Copy data between a Linux instance and a Windows instance

We recommend that you use the SFTP tool to download files from the Linux instance to your local PC and then use the FTP tool to upload the files to the Windows instance.

- Copy data between Windows instances
 - Use the FTP tool
 - Use TradeManager

How do I test the performance of an enhanced SSD?

This topic describes how to configure the environment to perform a performance stress test on an enhanced SSD and obtain the one million IOPS result. Note that the disk type and actual environment used will impact the test result.



Warning:

You can obtain accurate test results by testing bare disk partitions. However, you may destroy the file system structure in a bare disk partitions if you test the partition directly. Before you perform such a test, you must [create a snapshot of the disk](#) to back up your data. We recommend that you use newly purchased ECS instances that contain no data for the test to prevent data loss.

We recommend that you use the latest versions of Linux images provided by Alibaba Cloud, such as CentOS 7.4 64-bit, CentOS 7.3 64-bit, CentOS 7.2 64-bit, and Aliyun Linux 17.1 64-bit. This is because the drivers provided by Linux images of earlier versions and Windows images may not be applicable to enhanced SSDs.

We also recommend that you use the [FIO](#) tool for the performance stress test.

This example describes how to test the random write (`randwrite`) performance of an enhanced SSD. Assume that the instance type is `ecs.g5se.18xlarge` and the device name of the enhanced SSD is `/dev/vdb`.

1. Connect to the Linux instance.
2. Run the following commands to install the `libaio` and `FIO` tools:

```
sudo yum install libaio - y
sudo yum install libaio - devel - y
sudo yum install fio - y
```

3. Run the `cd / tmp` command to switch to the `/tmp` directory.
4. Run the `vim test100w . sh` command to create script `test100w.sh` and paste the following code to the script to test the random write (`randwrite`) IOPS.

```
function RunFio
{
    numjobs =$ 1 # The number of the threads to be
    tested . In this example , the value is 8 .
```



```

iodepth=$ 2 # The maximum number of concurrent I / O
requests . In this example , the value is 64 .
bs=$ 3 # The size of the data block per I /
O . In this example , the value is 4k .
rw=$ 4 # The read and write policy . In this
example , the value is randwrite .
filename =$ 5 # The name of the file to be tested
. In this example , the value is / dev / vdb .
nr_cpus=` cat / proc / cpuinfo | grep " processor " | wc - l

if [ $ nr_cpus - lt $ numjobs ]; then
    echo " Numjobs is more than cpu cores , exit !"
    exit - 1
fi
let nu =$ numjobs + 1
cpulist=""
for (( i = 1 ; i < 10 ; i ++))
do
    list=` cat / sys / block / vdb / mq /*/ cpu_list | awk '{
if ( i <= NF ) print $ i ;}' i ="$ i " | tr - d ',' | tr '\
n ' ','`
    if [ - z $ list ]; then
        break
    fi
    cpulist=${ cpulist }${ list }
done
spincpu=` echo $ cpulist | cut - d ',' - f 2 -${ nu }`
echo $ spincpu
fio -- ioengine = libaio -- runtime = 30s -- numjobs=${
numjobs } -- iodepth=${ iodepth } -- bs=${ bs } -- rw=${ rw }
-- filename=${ filename } -- time_based = 1 -- direct = 1 --
name = test -- group_repo rting -- cpus_allow ed=$ spincpu --
cpus_allow ed_policy = split
}
echo 2 > / sys / block / vdb / queue / rq_affinit y
sleep 5
RunFio 10 64 4k randwrite / dev / vdb

```

**Note:**

- You need to modify the parameter settings in the following commands based on your test environment:

```

- list=` cat / sys / block / vdb / mq /*/ cpu_list | awk '{
if ( i <= NF ) print $ i ;}' i ="$ i " | tr - d ',' | tr
'\ n ' ','` command: vdb .

```

```

- RunFio 10 64 4k randwrite / dev / vdb command: 10,64,
4k, randwrite ,and / dev / vdb .

```

- The file system structure may be damaged if you test the bare disk partition directly. If you choose to proceed with this operation, set `filename` to a device name such as `/dev/vdb`. If you do not want to risk data loss, set `filename` to a file path such as `/mnt/test.image`.

5. Run the `sh test100w . sh` command to start testing the performance of the enhanced SSD.

```
[root@ ~]# sh test100w.sh
4,8,12,16,20,24,28,1
test: (g=0): rw=randwrite, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B, ioengine=libaio, iodepth=64
...
fio-3.0
Starting 8 processes
Jobs: 8 (f=8): [w(8)][100.0%][r=0KiB/s,w=3965MiB/s][r=0,w=1015k IOPS][eta 00m:00s]
test: (groupid=0, jobs=8): err= 0: pid=9565: Thu Jan 11 17:37:00 2018
write: IOPS=1006k, BW=3931MiB/s (4122MB/s)(115GiB/30047msec)
    slat (usec): min=2, max=52179, avg= 4.22, stdev= 9.78
    clat (usec): min=82, max=624296, avg=500.05, stdev=814.37
    lat (usec): min=86, max=624300, avg=504.46, stdev=814.43
    clat percentiles (usec):
      1.00th=[ 174],  5.00th=[ 210], 10.00th=[ 235], 20.00th=[ 273],
     30.00th=[ 306], 40.00th=[ 338], 50.00th=[ 383], 60.00th=[ 441],
     70.00th=[ 523], 80.00th=[ 660], 90.00th=[ 930], 95.00th=[ 1139],
     99.00th=[ 1582], 99.50th=[ 1909], 99.90th=[ 3458], 99.95th=[ 4424],
     99.99th=[17957]
    bw (  KiB/s): min=256936, max=526637, per=12.53%, avg=504445.98, stdev=33299.37, samples=480
```

Script details:

- Block device parameters

In the [script](#), the `echo 2 > / sys / block / vdb / queue / rq_affinit y` command is used to set the `rq_affinit y` parameter to 2.

- If `rq_affinit y` is set to 1, the block device forwards the I/O Completion events it receives to the vCPU groups that submit the I/O requests. When multiple threads concurrently processes I/O requests, the I/O Completion events may run on the same vCPU, which results in a performance bottleneck.
- If `rq_affinit y` is set to 2, the block device forwards the I/O Completion events it receives to the vCPU that submits the I/O requests. The performance of each vCPU is fully delivered when multiple threads concurrently process I/O requests.
- Binding threads to corresponding vCPUs
 - Typically, a device has only one Request-Queue. This unique Request-Queue becomes a performance bottleneck when multiple threads concurrently process I/O requests.
 - In the latest Multi-Queue mode, one device can have multiple Request-Queues that process I/O requests, which deliver full performance of the back-end storage. If you have four I/O threads, you must bind them to the CPU cores that correspond to different Request-Queues. This allows you to use the full capabilities of Multi-Queue mode to improve the storage performance.

To ensure that the device delivers its full performance, you must assign the I/O requests to different Request-Queues. The following command in the [script](#) can be

used to bind `jobs` to different CPU cores. `/dev/vd*` is the device name of your enhanced SSD, for example, `/dev/vdb`.

```
fio -ioengine=libaio -runtime=30s -numjobs=${numjobs}
-iodepth=${iodepth} -bs=${bs} -rw=${rw} -filename=${filename}
-time_based=1 -direct=1 -name=test
-group_reporting -cpus_allowed=${spincpu} -cpus_allowed_policy=split
```

FIO provides the `cpus_allowed` and `cpus_allowed_policy` parameters to bind vCPUs. The preceding command runs multiple `jobs`. They are bound to different CPU cores and correspond to different queue IDs.

To view the CPU core ID that corresponds to a queue ID, perform the following steps:

- Run the `ls /sys/block/vd/mq/` command to view the queue ID of the enhanced SSD whose device name is `/dev/vd*`, for example, `/dev/vdb`.
- Run the `cat /sys/block/vd/mq//cpu_list` command to view the CPU core ID corresponding to the queue ID.