

Alibaba Cloud Elastic Compute Service

Block Storage

Issue: 20190606

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 Storage parameters and performance test.....	3
3 Block storage.....	12
3.1 Cloud disks and Shared Block Storage.....	12
3.2 Triplicate technology.....	16
3.3 ECS disk encryption.....	17
3.4 Create a cloud disk.....	21
3.4.1 Create a cloud disk.....	21
3.4.2 Create a Subscription cloud disk.....	23
3.4.3 Create a cloud disk by using a snapshot.....	26
3.5 Attach a cloud disk.....	29
3.6 Format a data disk.....	33
3.6.1 Format a data disk for Windows instances.....	33
3.6.2 Format a data disk of a Linux instance.....	35
3.6.3 Partition and format data disk more than 2 TiB.....	40
3.7 Resize cloud disks.....	48
3.7.1 Disk resizing overview.....	48
3.7.2 Resize a cloud disk.....	50
3.7.3 Extend a Windows file system.....	52
3.7.4 Extend the file system of the Linux system disk.....	57
3.7.5 Extend the file system of a Linux data disk.....	61
3.8 Change the operating system.....	71
3.8.1 Replace the system disk (non-public image).....	71
3.8.2 Replace the system disk (public image).....	76
3.9 Roll back a cloud disk.....	83
3.10 Reinitialize a cloud disk.....	85
3.11 Convert billing methods of cloud disks.....	92
3.12 Monitor a cloud disk.....	93
3.13 Detach a cloud disk.....	95
3.14 Release a cloud disk.....	98
4 Shared Block Storage.....	100
5 Local disks.....	103
6 Block storage FAQ.....	109

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 [Block storage performance](#).

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 Storage parameters and performance test

This document describes the performance index of block storage, performance testing methods, and how to interpret the testing results.

Performance index of block storage

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

- IOPS

IOPS stands for Input/Output Operations per Second, which means the number of write or read operations that can be performed each second. Transaction-intensive applications, such as database applications, are sensitive to IOPS.

The following table lists common performance characteristics that are measured.

IOPS performance characteristics	Description	
Total IOPS	The total number of I/O operations per second	
Random read IOPS	The average number of random read I/O operations per second	Random access to locations on 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	Sequential access to locations on storage devices
Sequential write IOPS	The average number of sequential write I/O operations per second	

- Throughput

Throughput measures the data size successfully transferred per second.

Applications that require mass read or write operations (such as Hadoop offline computing applications) are sensitive to throughput.

- Latency

Latency is the period that is needed to complete an I/O request.

For latency-sensitive applications (such as databases) in which high latency may lead to performance reduction or error reports in applications, we recommend that you use SSD disks, SSD Shared Block Storage, or local SSD disks.

For throughput-sensitive applications (such as Hadoop offline computing) that are less sensitive to latency, we recommend that you use ECS instances with local HDD disks, such as instances of the d1 or d1ne instance type family.

Performance

This section describes the performance of various block storage products.

Block Storage capacity is measured in binary units, such as kibibyte (KiB), mebibyte (MiB), gibibyte (GiB), or Tebibyte (TiB).

**Note:**

1 KiB is 1,024 bytes. 1MiB is 1,024 KiB. 1GiB is 1,024 MiB. 1TiB is 1,024 GiB.

- Cloud disks

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

Parameter	ESSD Cloud Disk	SSD Cloud Disk	Ultra Cloud Disk	Basic Cloud Disk
Capacity of a single disk	32,768 GiB	32,768 GiB	32,768 GiB	2,000 GiB
Max. IOPS	1,000,000	25,000*	5,000	Several hundreds
Max. throughput	4,000 MBps	300 MBps*	140 MBps	30–40 MBps

Parameter	ESSD Cloud Disk	SSD Cloud Disk	Ultra Cloud Disk	Basic Cloud Disk
Formulas to calculate performance of a single disk **	$\text{IOPS} = \min\{1800 + 50 \times \text{capacity}, 1,000,000\}$	$\text{IOPS} = \min\{1800 + 30 \times \text{capacity}, 25,000\}$	$\text{IOPS} = \min\{1800 + 8 \times \text{capacity}, 5,000\}$	N/A
	$\text{Throughput} = \min\{120 + 0.5 \times \text{capacity}, 4,000\} \text{ MBps}$	$\text{Throughput} = \min\{120 + 0.5 \times \text{capacity}, 300\} \text{ MBps}$	$\text{Throughput} = \min\{100 + 0.15 \times \text{capacity}, 140\} \text{ MBps}$	N/A
Data reliability	99.9999999%	99.9999999%	99.9999999%	99.9999999%
API name	cloud_essd	cloud_ssd	cloud_efficiency	cloud

Parameter	ESSD Cloud Disk	SSD Cloud Disk	Ultra Cloud Disk	Basic Cloud Disk
Scenarios	<ul style="list-style-type: none"> - OLTP databases : relational databases such as MySQL, PostgreSQL, Oracle, and SQL Server - NoSQL databases : non-relational databases such as MongoDB, HBase, and Cassandra - ElasticSearch distributed logs: Elasticsearch, Logstash and Kibana (ELK) log analysis 	<ul style="list-style-type: none"> - Large and medium -sized relational databases , such as MySQL, SQL Server, PostgreSQL, and Oracle - Large or medium -sized development or testing applications that require high data reliability 	<ul style="list-style-type: none"> - Small or medium -sized relational databases , such as MySQL, SQL Server, and PostgreSQL - Large or medium -sized development or testing applications that require high data reliability and medium performance 	<ul style="list-style-type: none"> - Applications with infrequent access or low I/O load. If higher I/O performance is needed , we recommend that you use SSD disks. - Applications that require low costs and random read and write I/O operations

* The performance of an SSD Cloud Disk varies with the data block size. Smaller data blocks result in lower throughput and higher IOPS, as shown in the following table. An SSD Cloud Disk can achieve the expected performance only when it is attached to an I/O-optimized instance. In other words, an SSD Cloud Disk cannot achieve the expected performance if it is not attached to an I/O-optimized instance.

Data block size	Maximum IOPS	Throughput
4 KiB	About 25,000	Far smaller than 300 MBps
16 KiB	About 17,200	Close to 300 MBps

Data block size	Maximum IOPS	Throughput
32 KiB	About 9,600	
64 KiB	About 4,800	

** An SSD Cloud Disk is taken as an example to describe the performance of a single disk:

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

The random write latency varies with the disk categories as follows:

- ESSD disks: 0.1–0.2 ms
- SSD disks: 0.5–2 ms
- Ultra Cloud Disks: 1–3 ms
- Basic Cloud Disks: 5–10 ms
- Shared Block Storage

The following table lists the features and typical scenarios of different types of Shared Block Storage.

Parameter	SSD Shared Block Storage	Ultra Shared Block Storage
Capacity	<ul style="list-style-type: none"> - Single disk: 32,768 GiB - Single instance: 128 TiB 	<ul style="list-style-type: none"> - Single disk: 32,768 GiB - Single instance: 128 TiB
Maximum random read/write IOPS*	30,000	5,000
Maximum sequential read/write throughput*	512 MBps	160 MBps
Formulas to calculate performance of a single disk**	$IOPS = \min\{1600 + 40 \times \text{capacity}, 30,000\}$	$IOPS = \min\{1000 + 6 \times \text{capacity}, 5,000\}$
	$\text{Throughput} = \min\{100 + 0.5 \times \text{capacity}, 512\} \text{ MBps}$	$\text{Throughput} = \min\{50 + 0.15 \times \text{capacity}, 160\} \text{ MBps}$

Parameter	SSD Shared Block Storage	Ultra Shared Block Storage
Scenarios	<ul style="list-style-type: none">- Oracle RAC- SQL Server- Failover cluster- High-availability architecture of servers	<ul style="list-style-type: none">- High-availability architecture of servers- High-availability architecture of development and testing databases

* 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.

** An SSD Shared Block Storage is used as an example to describe the performance of a single disk:

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

The latency varies with the shared block storage categories as follows:

- SSD Shared Block Storage: 0.5–2 ms
- Ultra Shared Block Storage: 1–3 ms
- Local disks

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

Test disk performance

fio is recommended to test disk performance.



Note:

The disk benchmark tested by different tools varies with different operating systems. The performance parameters in this article are the results tested by fio with a Linux instance, and are used as the index reference of block storage product performance.

This section describes how to test disk performance, taking the fio tool used with a Linux instance as an example. Before you test the disk, verify that the disk is 4 KiB aligned.

**Warning:**

You can test bare disks to obtain more accurate performance data, but the structure of the file system will be damaged. Make sure that you back up your data before testing. We recommend that you use a new ECS instance without data to test the disks to avoid data loss.

- 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
```

- 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
```

- Test 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
```

- Test 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
```

The command for testing random read IOPS is used as an example to describe the meaning of the parameters of a fio command, as shown in the following table.

Parameter	Meaning
-direct=1	Ignore I/O buffer when testing. Data is written directly.
-iodepth=128	Indicates that when you use AIO, the maximum number of I/O issues at the same time is 128.

Parameter	Meaning
-rw=randwrite	<p>Indicates that the read and write policy is random write. Other options include:</p> <ul style="list-style-type: none">· randread (random read)· read (sequential read)· write (sequential write)· randrw (random read and write)
-ioengine=libaio	<p>Use libaio as the testing method (Linux AIO, Asynchronous I/O). Usually there are two ways for an application to use I/O:</p> <ul style="list-style-type: none">· Synchronous <p>Synchronous I/O only sends out one I/O request at a time, and returns only after the kernel is completed. In this case, the iodepth is always less than 1 for a single job, but can be resolved by multiple concurrent jobs. Usually 16–32 concurrent jobs can fill up the iodepth.</p> <ul style="list-style-type: none">· Asynchronous <p>The asynchronous method uses libaio to submit a batch of I/O requests each time, thus reducing interaction times and making interactions more effective.</p>

Parameter	Meaning
-bs=4k	<p>Indicates the size of each block for one I/O is 4 KiB. If not specified, the default value 4 KiB is used.</p> <p>When IOPS is tested, we recommend that you set <code>bs</code> to a small value, for example, 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 this example command.</p>
-size=1G	Indicates the size of the testing file is 1 GiB.
-numjobs=1	The number of testing jobs is 1.
-runtime=1000	Testing time is 1,000 seconds. If not specified, the test will write data of the file whose size is specified by <code>-size</code> block by block, with the data block size specified by <code>-bs</code> .
-group_reporting	The display mode for showing the testing results. Group_reporting means the statistics of each job are summed up, instead of all statistics of each job being shown.
-filename=iotest	The output path and name of the test files, for example, iotest. You can test bare disks to obtain more accurate performance data, but the test causes damage to the structure of the file system. Make sure that you back up your data before testing.
-name=Rand_Write_Testing	The name of the testing task.

3 Block storage

3.1 Cloud disks and Shared Block Storage

Cloud disks and Shared Block Storage are block-level data storage products provided by Alibaba Cloud for ECS that features low latency, high performance, persistence, and high reliability. They use a [triplicate distributed system](#) to provide 99.9999999% data reliability for ECS instances. Cloud disks and Shared Block Storage can automatically copy your data within the target zone to help you prevent unexpected hardware faults from causing data unavailability or service disruption. Just like what you do with a hard disk, you can partition and format the cloud disks and Shared Block Storage attached to an ECS instance, create a file system, and store data on them.

You can expand the cloud disks and Shared Block Storage as needed at any time. For more information, see [Linux - Resize a data disk](#) and [increase system disk size](#). You can also create snapshots to back up data for the cloud disks and Shared Block Storage. For more information about snapshots, see [what are ECS snapshots](#).

Cloud disks and Shared Block Storage differ in whether they can be simultaneously attached to multiple ECS instances and perform read and write operations. Details are as follows:

- Cloud disks can be attached to only one ECS instance in the same zone of the same region.
- Shared Block Storage devices can be mounted to a maximum of eight ECS instances in the same zone of the same region.



Note:

Shared Block Storage is currently in public beta phase. You can [open a ticket](#) to submit your application for beta testing.

Cloud disks

- Performance-based category
 - ESSD: An ultra-high-performance cloud product based on the next generation distributed block storage architecture. ESSD combines 25 GE networks with RDMA technology, offering the capability of up to 1 million random read/write

operations and a shorter single-link latency. ESSD is currently in public beta phase. For more information, see [FAQ about ESSD cloud disks](#).

- SSD cloud disks: high-performance disks with stable and high random I/O performance and high data reliability
- Ultra cloud disks: with high cost performance, medium random I/O performance, and high data reliability
- Basic cloud disks: with high data reliability and general random I/O performance
- Function-based category
 - System disks: have the same life cycle as the ECS instance to which it is mounted. A system disk is created and released at the same time as the instance. Shared access is not allowed. The available size range of a single system disk varies according to the image, as follows:
 - Linux (excluding CoreOS) and FreeBSD: 20–500 GiB
 - CoreOS: 30–500 GiB
 - Windows: 40–500 GiB
 - Data disks: can be [created separately](#) or at the same time as ECS instances. If the release with instance attribute is set to a data disk, it has the same life cycle as the instance to which it is attached, and is released along with the instance. Data disks created separately can be [released independently](#) or at the same time as the corresponding ECS instances. Shared access is not allowed. The performance of data disks depends on the cloud disk type. For more information, see [storage parameters and performance test](#).

When used as data disks, up to 16 cloud disks can be attached to one ECS instance.

Shared Block Storage

Shared Block Storage is a block-level data storage service with strong concurrency, high performance, and high reliability. It supports concurrent reads from and writes to multiple ECS instances, and provides data reliability of up to 99.9999999%. Shared Block Storage can be mounted to a maximum of 8 ECS instances.

Shared Block Storage can only be used as data disks and can only be created separately. Shared access is allowed. You can set the Shared Block Storage device to be released when the ECS instances are released.

Shared Block Storage can be divided into:

- SSD Shared Block Storage, which uses SSD as the storage medium to provide stable and high-performance storage with enhanced random I/O and data reliability.
- Ultra Shared Block Storage, which uses the hybrid media of SSD and HDD as the storage media.

When used as data disks, Shared Block Storage allows up to 16 data disks to be attached to each ECS instance.

For more information, see [FAQ about Shared Block Storage](#).

Billing

Shared Block Storage is currently in public beta phase free of charge.

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

- Cloud disks created with Subscription instances are billed before the service is ready for use. For more information, see [Subscription](#).
- Cloud disks created at the same time as Pay-As-You-Go instances, or created separately, are billed on a Pay-As-You-Go basis. For more information, see [Pay-As-You-Go](#).

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

Conversion of billing methods	Feature	Effective time	Suitable for
Subscription -> Pay-As-You-Go	Renew for configuration downgrade	Effective from the next billing cycle	Subscription cloud disks mounted to Subscription instances. The billing method of the system disk cannot be changed.
Pay-As-You-Go -> 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.

Conversion of billing methods	Feature	Effective time	Suitable for
	Switch from Pay-As-You-Go to Subscription billing		System disks and data disks mounted to Pay-As-You-Go instances.

Related operations

You can perform the following operations on cloud disks:


- If [a cloud disk or Shared Block Storage device is created separately from a data disk](#), you must [attach a cloud disk](#) in the ECS console, and then connect to the ECS instance to [partition and format the data disk](#).
- If you want to encrypt the data on a cloud disk, [encrypt the disk](#).
- If your system disk capacity is insufficient, you can [increase system disk size](#).
- If you want to expand the data disk capacity, you can [resize the data disk](#).
- If you want to change the OS, you can [change the system disk](#).
- If you want to back up the data of a cloud disk or Shared Block Storage device, you can [manually create snapshots for the cloud disk or Shared Block Storage](#) or [apply an automatic snapshot policy to it](#) to automatically create snapshots on schedule.
- If you want to use the OS and data environment information of one instance on another instance, you can [create a customized image using the system disk snapshots of the latter instance](#).
- If you want to restore a cloud disk or Shared Block Storage device to the status when the snapshot is created, you can [roll back a cloud disk](#) using its snapshot.
- If you want to restore a cloud disk to its status at the time of creation, you can [reinitialize a cloud disk](#).
- If you do not need a cloud disk or Shared Block Storage device, you can [detach a cloud disk](#) and [release a cloud disk](#).
- If you no longer need a Subscription billed cloud disk, you can [convert the billing methods of cloud disks](#), and then [detach a cloud disk](#) and [release a cloud disk](#).

For more information about operations on cloud disks, see [cloud disks](#) in *User Guide* .

3.2 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.

3.3 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 unauthorized 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.

3.4 Create a cloud disk

3.4.1 Create a cloud disk

You can create a cloud disk to work as a data disk in the ECS console or by using the API. This article introduces how to create a new empty cloud disk in the ECS console.

Notes

Before you create a cloud disk, consider the following:

- Only [Pay-As-You-Go](#) cloud disks can be created in this way, and they can be used as data disks only.



Note:

You can create cloud disks as data disks when creating an ECS instance. Those disks have the same billing method of the instance.

- You can create a new empty cloud disk or [create a cloud disk from a snapshot](#).

- The quota of the Pay-As-You-Go cloud disks that are used as data disks of each account in all regions is five times than that of the Pay-As-You-Go instances. For more information, see [limits](#).
- Currently, you cannot merge multiple cloud disks. After cloud disks are created, they are independent from each other, and you cannot merge their space by formatting. We recommend that you determine the number of disks and disk sizes required for your business before you create cloud disks.
- Because you can create a snapshot for a single cloud disk, we do not recommend that you create LVM (Logical Volume Manager) volumes as the volumes may result in data loss if you use the snapshot to roll back the cloud disk.
- You can convert a Pay-As-You-Go billed cloud disk to Subscription as follows:
 - [Upgrade configurations of Subscription instances](#).
 - [Switch from Pay-As-You-Go to subscription](#).
- If a cloud disk is created in this way, and its billing method is retained as Pay-As-You-Go, you can [detach a cloud disk](#) and [release a cloud disk](#) at any time.

Prerequisites

If you want to [attach a cloud disk](#) to an instance, make sure they are in the same region and zone.

Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, select Block Storage > Disks.
3. In the upper-right corner of the Diskslist page, click Create Disk to go to the Create page.
4. 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 and the same region.

5. Select a cloud disk category and specify the disk size and the quantity. You can also select [create a cloud disk from a snapshot](#).
6. Confirm the configuration and the Total cost.
7. Click Preview, confirm you order, and click Create.

After you complete the payment, return to the Disks page and refresh it. The new disk is displayed and its status is Available.

Additional operations

[Attach a cloud disk](#).

Related APIs

To create a disk after creating an instance, see [#unique_35](#).

To create a cloud disk when creating an instance, see [#unique_45](#) or [#unique_34](#).

3.4.2 Create a Subscription cloud disk

To increase the storage space for a Subscription instance, you can create a Subscription cloud disk for that instance in the ECS console. This topic describes how to create a Subscription cloud disk for a Subscription instance in the ECS console.

Precautions

Before you create a cloud disk, note the following:

- If you create a Subscription cloud disk for a Subscription instance on the Instances page in the ECS console, that cloud disk is billed in the [Subscription](#) method and can only work as a data disk.



Note:

You can create a cloud disk as a data disk when creating a Subscription instance. Cloud disks created in this way have the same billing method as the corresponding instance.

- You can create a new empty cloud disk or [create a cloud disk from a snapshot](#).
- Currently, ECS does not allow you to merge multiple cloud disks. Each cloud disk is an independent entity. You cannot merge the disks by formatting them. We recommend that you determine the number and size of cloud disks before you create them.
- For multiple cloud disks that were previously created, we do not recommend that you create logical volumes such as Logical Volume Manager (LVM) volumes. This is because a snapshot is created only for an independent cloud disk. If LVM is used, data loss will occur when you use a snapshot to restore a cloud disk.


- For cloud disks created in this way, you cannot [detach a cloud disk](#). Such cloud disks expire at the same time as the corresponding instances.

**Note:**

To release a Subscription cloud disk, convert its billing method to Pay-As-You-Go, detach it, and then release it.

Procedure

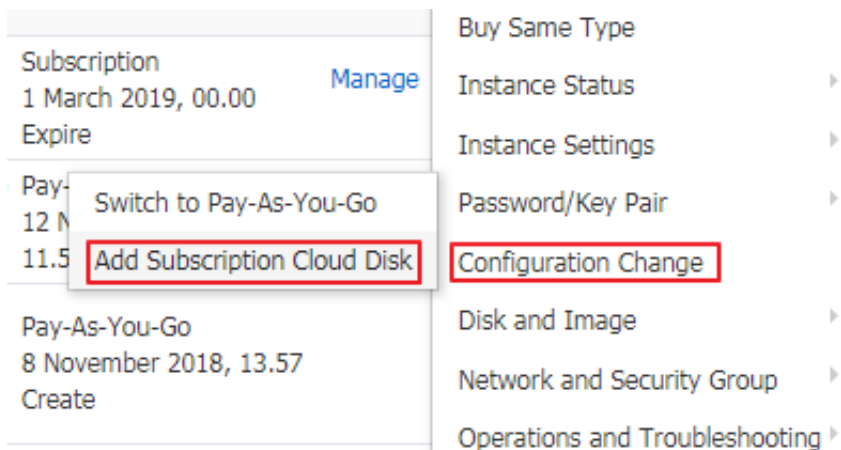
- Log on to the [ECS console](#).
- In the left-side navigation pane, click Instances.



- On the Instances page, find the target Subscription instance, and then click More in the Actions column.

Configuration	Billing Method ▾	Actions	
2 vCPU 8 GB (I/O Optimized) ecs.g5.large 1Mbps	Subscription 7 February 2019, 00.00 Expire	Manage Connect Change Configuration Renew	More ▾

- Click Configuration Change > Add Subscription Cloud Disk.



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

- Cloud disk type: Select a [cloud disk type](#) in the drop-down box.
- 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 [ECS disk encryption](#) is needed, select the Encrypted check box.
- Quantity: Enter the number of cloud disks to add in the text box.



Note:

You can create up to 16 data disks (including cloud disks and shared block storage devices) for a single instance.

- Disk Name: Optional. You can enter a disk name in the text box. A disk name can contain 2 to 128 English letters or Chinese characters in length. It can also contain numbers, periods (.), colons (:), underscores (_), and hyphens (-).
- Description: Optional. You can enter a disk description in the text box. The description information can contain 2 to 256 English letters or Chinese characters. It cannot start with http:// or https://.
- If you need to [create a cloud disk from a snapshot](#), click Create from snapshot.

6. Select the *ECS Service Level Agreement* check box.

7. Click Preview.

8. Click Create Order.

9. Select a payment method and click Confirm to Pay to complete the creation.

10. Click ECS console to return to the Instances page. In the instance list, click the instance for which the Subscription cloud disk has just been added.

11.Click Disks to view the newly added Subscription cloud disk.

Disk									Create Disk	Mount
ID/Name	Tag	Disk Type(All)	Status (All)	Billing Method(All)	Unmountable(All)	Zone	Type(All)	Encrypted/Unencrypted	Actions	
d-4mhnazc2t0lloio1		SSD 20GB	In Use	Subscription	No	Qingdao Zone C	Data Disk	Unencrypted	Create Snapshot	Reinitialize Disk
-									Create Automatic Snapshot Policy	More
d-4smeespcvndzscig1		Ultra Disk 20GB	In Use	Subscription	No	Qingdao Zone C	Data Disk	Unencrypted	Create Snapshot	Reinitialize Disk
-									Create Automatic Snapshot Policy	More
d-4m3ecgkqjxwsmrsl14		Ultra Disk 40GB	In Use	Subscription	No	Qingdao Zone C	System Disk	Unencrypted	Create Snapshot	Reinitialize Disk
-									Create Automatic Snapshot Policy	More

3.4.3 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 [Pay-As-You-Go](#) 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 [What is OSS?](#)
- 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 [Limits](#).

- 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. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
3. In the upper-right corner of the Disks page, click Create Disk.
4. 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.

5. 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 [Partition and format data disk more than 2 TiB](#).

- 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.
6. Confirm your settings and the estimated cost displayed.
 7. 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 [Upgrade configurations of Subscription instances](#).
 - If the cloud disk is attached to a Pay-As-You-Go instance, see [Switch from Pay-As-You-Go to Subscription billing](#).

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.

3.5 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 status, but not in the Locked status.
 - The instance must not have any overdue payments.
- The disk to be attached must be in the Unmounted status.
- 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 [Create an instance by using the wizard](#) and [Create a 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. Log on to the [ECS console](#).
2. In the left-side navigation pane, click Instances.
3. Select the target region.
4. Find the target ECS instance and click its ID to go to its Instance Details page.

5. In the left-side navigation pane, click Disks. Then on the Disks page, click Mount in the upper right corner.

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

- **Target Disk:** Select a cloud disk in the Unmounted status 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.

7. Refresh the Disks page.

When the status of the cloud disk is In Use, the attachment is successful.

8. 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	Format a data disk of a Linux instance . If the cloud disk is larger than 2 TiB, see Partition and format data disk more than 2 TiB .
	Windows	Format a data disk for Windows instances . If the cloud disk is larger than 2 TiB, see Partition and format data disk more than 2 TiB .

Disk content	Operating system of the ECS instance	Follow-up operations
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. Log on to the [ECS console](#).
2. In the left-side navigation pane, select Block Storage > Disks.
3. Select the target region.
4. Find a cloud disk in the Unmounted status. Then in the Actions column, select More > Mount.
5. 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.

6. Refresh the Disks page.

When the status of the cloud disk is In Use, the attachment is successful.

7. 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	Format a data disk of a Linux instance . If the cloud disk is larger than 2 TiB, see Partition and format data disk more than 2 TiB .
	Windows	Format a data disk for Windows instances . If the cloud disk is larger than 2 TiB, see Partition and format data disk more than 2 TiB .
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](#)

3.6 Format a data disk

3.6.1 Format a data disk for Windows instances

If data disks are selected when you create a Windows instance, you need to partition and format them for use.

This article describes how to create a single-partition data disk using a new data disk and how to mount a file system. You can also configure multiple partitions based on business needs. This article applies only to data disks that are not larger than 2 TiB. For those that are larger than 2 TiB, see [Partition and format data disk larger than 2 TiB](#).



Warning:

- Disk partitioning and formatting are high-risk operations. Please proceed with caution. This article describes how to deal with a blank data disk. If you have data on a data disk, be sure to create a snapshot for the data disk to avoid any possible data loss.
- ECS instances only support partitioning data disks, not system disks. If you use a third-party tool to forcibly partition the system disk, unknown risks may occur, for example, system crash and data loss.

Prerequisites

For a separately [purchased data disk](#), you must [attach the data disk to an instance](#) before partitioning and formatting.

A data disk purchased along with the instance can be partitioned and formatted without being attached.

Procedure

This example describes how to partition and format a 20 GiB data disk on the 64-bit Windows Server 2012 R2.

1. [Connect to an instance](#).

2. On Windows Server desktop, right click the Start icon, then select Disk management.

The unformatted data disk (Disk 2) appears as Offline.

3. Right click the blank area around Disk 2, and select Online in the context menu.

After going online, the status of Disk 2 is displayed as Not Initialized.

4. Right click the blank area around Disk 2, and then select Initialize Disk in the context menu.

5. In the Initialize Disk dialog box, select Disk 2 and a partitioning method:

- MBR is still the most common partitioning method. However, this method only supports data disks that no greater than 2 TB and can divide a disk into up to four primary partitions. If you want to divide a disk into more than four partitions, you need to take a primary partition as an extended partition and create logical partitions within it.
- GPT is a new partitioning method, and cannot be recognized by earlier versions of Windows. The size of GPT-partitioned data disk is determined by the operating system and the file system. In the Windows operating system, GPT supports up to 128 primary partitions.


In this example, select the MBR partitioning method, and click OK.

6. In the Disk Management window, right click the Unallocated area for Disk 2 and select New Simple Volume.

7. In the New Simple Volume Wizard, follow these steps:

- a. Click Next.
- b. Specify Volume Size: Specify the size of the simple volume to create. If you need only one primary partition, use the default value, and then click Next.
- c. Assign Drive Letter or Path: Select a drive letter (in this example, F). Click Next.
- d. Format Partition: select format settings (including the file system, allocation unit size, and volume label), and confirm whether to enable Quick Formatting and File and Folder Compression. Use the default values, then click Next.
- e. Create a new simple volume. When the wizard shows the information below, a new simple volume is created. Click Finish to close the New Simple Volume Wizard.

After the partition formatting is completed, the status of Disk 2 in Disk Management is as shown in the following figure.



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

3.6.2 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 multiple 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.



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](#). We recommend that you use the built-in system tool for partitioning.

Prerequisites

- You have [attached the data disk to an instance](#) in the ECS console. For more information, see [Create a Pay-As-You-Go cloud disk](#).



Note:

For a data disk purchased along with an instance, you do not need to attach it to an instance.

- You have obtained the device name of the data disk.

You can obtain the device name of the data disk by choosing ECS Console > Block Storage > Disks > (Disk ID specific) More > Modify Attributes. By default, device names are assigned by the system, starting from `/ dev / vdb` to `/ dev / vdz`.

Procedure

In this example, a new 20 GiB data disk (device name: `/ 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 your data disk name is `dev / xvd ?` (where, `?` is a letter from a to z), you are using a non-I/O-optimized 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.

```
# 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.

```

# fdisk -lu / dev / vdb

Disk / dev / vdb : 21 . 5 GB , 2147483648 0 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 OSs, 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 @ i ##### ~]# 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 @ i ##### ~]# 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 @ i ##### ~]# 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
```

3.6.3 Partition and format data disk more than 2 TiB

This article describes how to partition and format a large data disk in different operating systems using the GPT format. Note that a large data disk indicates a disk greater than 2 TiB, while a disk size smaller than 2 TiB indicates a small data disk.



Note:

If you want to partition and format a data disk less than 2 TiB, please see [Format a data disk for Linux instances](#) and [Format a data disk for Windows instances](#).

Note

Before partition and formatting a large data disk, note the following:

- Large data disks support the partition tools and file systems shown in the following table.

Operating system	Partition tool	File system
Linux	parted	ext4 or xfs
Windows	Disk management	NTFS

- We recommend that you do not create a large data disk by using a snapshot of a small data disk.

If you create a large data disk by using the snapshot of a small data disk, the following risks may occur:

- The system expands the block-level of the device's disk, but does not automatically convert between the partition format and the file system.
- If the MBR format is used in the snapshot of the small data disk, none of the supported partition tools (`parted` on Linux and Disk Management on

Windows) can convert the MBR to GPT and retain the data. Therefore, even if you create a large data disk by using a snapshot of a small data disk, while partitioning and initializing, you must delete the original data and partition with the GPT format. If you have created large data disk by using a snapshot of a small data disk, see [use windows to partition and format a large data disk created by a snapshot of a small data disk](#) .

**Note:**

This is not the case if the snapshot of the small data disk is in GPT format, or if you have another powerful partitioning tool. You can select based on your own situation.

Instead, create an empty large data disk, or create large data disk by using snapshots of large data disks, because of the following reasons

- Effect of data disk snapshots

No matter if you are using a large data disk or a small data disk, the process for creating a snapshot from a data disk is proportional to the total data volume of the disk. If the disk contains a large amount of compromised or damaged data, or residual data fragments, then the snapshot creation time is extended

Use Windows to partition and format an empty large data disk

The following examples uses a Windows Server 2008 R2 64-bit operating system to describe how to partition and format a large data disk in Windows instance. Assume the data disk to be processed is a 4 TiB empty disk.

Prerequisites

The data disk has been attached to an instance. For more information, see [attach a cloud disk](#).

Procedure

To partition and format a large data disk, follow these steps:

1. [Connect to a Windows instance](#).

2. Click the  icon in the task bar.



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

4. Find the disk that is to be partitioned and formatted (in this example, Disk 4). The disk status should be shown as Offline.
5. Right click the blank area around Disk 4, and then click Online.

After going online, Disk 4 enters the Not Initialized status.
6. Right click the blank area around Disk 4, and then select Initialize Disk in the context menu.
7. In the Initialize Disk dialog box, select Disk 4 and select GPT as the disk partitioning method.
8. In the Disk Management window, right click the Unallocated area of Disk 4, and then select New Simple Volume to create a 4 TiB volume in the NTFS format.
9. In the New Simple Volume Wizard, follow these steps:
 - a. Click Next.
 - b. Choose a volume size: designate size of simple volume. If you want to create a master area only, use the default value. Click Next. You can also partition Disk 4 into several partitions.

**Note:**

The maximum NTFS volume is, theoretically, 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, for a 64 KiB cluster, the maximum NTFS volume is approximately 256 TiB. If you select a 4 KiB cluster, the maximum NTFS volume is 16 TiB. NTFS selects the size of a cluster automatically based on the disk capacity.

- c. Assign drive letter and path: select a drive letter, then select G in this instance. Click Next.
- d. Format Partition: Select the formatting settings, including file system, distributed unit size, and volume label, and then confirm whether to Perform a quick format and Enable file and folder compression. Select Perform a quick format here only. Click Next.
- e. Start creating a new simple volume by following the prompts displayed by the wizard. Then, click Finish.

After the formatted partition is completed, in Disk Management, the status of Disk 4 is shown in the following screenshot.

Use Windows to partition and format a large data disk created by a snapshot of a small data disk

If you create a large data disk by using snapshots of a small data disk, you first need to convert the partition format of data disk from MBR to GPT, and then format the data disk. However, because data of the original snapshots is saved, we recommend you do not create large data disk by using a snapshot of a small data disk.

If you have already created large data disks in this method, perform the following actions to partition and format this data disk. The following example operating system uses a Windows Server 2012 R2 64-bit operating system. Assume the data disk to be processed is a 3 Tib disk.

Prerequisites

The data disk has been [attached](#) to an instance.

Procedure

To partition and format a large data disk, follow these steps:

1. [Connect to a Windows instance](#).
2. On the Windows Server desktop, right click the Start icon, and select Disk Management.

The data disk (Disk 2, in this example) that has not been formatted or partitioned is in the Offline status.
3. Right click the blank area around Disk 2, and then select Offline in the context menu.
4. Right click a simple volume, and then select Delete Volume in the context menu.
5. Right click the blank area around Disk 2, and then select Convert to GPT Disk in the context menu.
6. In the Disk Management window, right click Unallocated area of Disk 2, and then select New Simple Volume to create a 3 TiB volume in the NTFS format.
7. In the New Simple Volume Wizard, follow these steps:
 - a. Click Next.
 - b. Specify the size of the simple volume. If you need only one primary partition, use the default value, and then click Next. You can also partition Disk 2 into several partitions.



Note:

The maximum NTFS volume is, theoretically, 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, for a 64 KiB cluster, the maximum NTFS volume is approximately 256 TiB. If you select a 4 KiB cluster, the maximum NTFS volume is 16 TiB. NTFS selects the size of a cluster automatically based on the disk capacity.

- c. Assign Drive Letter or Path: Select a drive letter, and then click Next.
- d. Format Partition: Select the formatting settings, including file system, distributed unit size and volume label, and then confirm whether to Perform a quick format and Enable file and folder compression. Select Perform a quick format here only. Click Next.
- e. Start creating a new simple volume by following the prompts displayed by the wizard. Then, click Finish.

After the formatted partition is completed, in Disk Management, the status of Disk 4 is shown in the following screenshot.

Use Linux to partition and format a large data disk

To partition and format a large data disk that is attached to a Linux instance, use the GPT format. In Linux system, large data disks normally uses xfs or ext4 file system.

The following example uses a CentOS 7.4 64-bit operating system. This section describes how to use parted and e2fsprogs tools to partition and format a large data disk on a Linux instance. Assume the data disk to be processed is an empty 3 TiB new disk, and the device name is `/dev/vdd`.

Prerequisites

Your Linux instance has installed parted. If not, run `yum install -y parted`.

Your Linux instance has installed e2fsprogs. If not, run `yum install -y e2fsprogs`.

The data disk has been attached to the instance. For more information, see [attach a cloud disk](#).

Procedure

To partition and format a large data disk and mount the file system, follow these steps:

1. Run `fdisk -l` to check whether the data disk exists. If the data disk is successfully mounted, the following result is returned

```
Disk / dev / vdd : 3221 . 2 GB , 3221225472 000 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 `parted / dev / vdd` to start partitioning:
 - a. Run `mklabel gpt` , to convert partitioning format from MBR to GPT.
 - b. Run `mkpart primary ext4 < StartSector > < EndSector >` to partition a primary partition by using the ext4 file system, and specify a start sector and end sector for the partition. If a data disk is partitioned into one partition only, run `mkpart primary ext4 0 - 1` .



Note:

You can also use xfs file system.

- c. Run `print` to check partition table.

```
( parted ) mkpart primary ext4 0 - 1
Warning : The resulting partition is not properly
aligned for best performance .
Ignore / Cancel ? ignore
( parted ) print
Model : Virtio Block Device ( virtblk )
Disk / dev / vdd : 3221 GB
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
```

- d. Run `quit` to exit parted .
3. Run `partprobe` to make system re-read the partition table.
 4. Run the following commands to create an ext4 file system, and make `/dev/vdd1` partition use ext4.

```
mke2fs -O 64bit , has_journal , extents , huge_file , flex_bg
, uninit_bg , dir_nlink , extra_isize / dev / vdd1
```



Note:

- If you want to disable the lazy init function of ext4 file system to avoid its effect on data disk I/O performance, see [disable lazy init function](#)..

- If capacity of the data disk is 16 TiB, you must format it by using `e2fsprogs` in the designated version. See [update e2fsprogs](#).
- If you want to create an xfs file system, run `mkfs -t xfs /dev/vdd1`.

5. Run `mkdir /test` to create a mount point with the name `/test`.

6. Run `mount /dev/vdd1 /test` to mount `/dev/vdd1` to `/test`.

7. Run `df -h` to check current disk space and usage.

If the returned result shows the newly created file system information, the mount operation was successful, and you can use the new file system directly (that is, you do not need to restart the instance).

```
[ root @ izXXXXz ~]# 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 mount partition while the instance is started.

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

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

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

If the new partition information is in the returned result, the write operation was successful.

Appendix 1: Update e2fsprogs

If the disk capacity is 16 TiB, you must use `e2fsprogs` of version 1.42 or later to format its partitions to `ext4` file system. If `e2fsprogs` version is too low (for example, `e2fsprogs 1.41.11`), 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 `e2fsprogs` of later version, such as 1.42.8 in this example, follow these steps:

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

```
$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 current version is earlier than 1.42, update the software by following these steps.

2. Run the following command to download e2fsprogs in version 1.42.8. Go to [e2fsprogs](#) to find 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 in turn to compile tools in later versions.

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

4. Run `rpm -qa | grep e2fsprogs` to check whether the software of the later version has been installed successfully.

Appendix 2: Disable lazy init function

The lazy init function of ext4 file system is enabled by default. While the function is enabled, it will run in the system background and initiate a thread to initialize metadata of ext4 file system continuously to delay metadata initialization. Therefore, immediately after formatting a data disk, IOPS can be affected.

If you need to test performance of data disk immediately after formatting, run the following commands to disable lazy init function while formatting 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
```

If the lazy init is disabled, it may take longer time to format a partition. For example, it may take 10–30 minutes to format a 32 TiB data disk.

3.7 Resize cloud disks

3.7.1 Disk resizing overview

In Alibaba Cloud, you can resize the disk volume of a system disk or a data disk at any time.

Scenarios

You can resize the volume of a disk to meet the needs of different scenarios, including the need to:

- [Increase the system disk size](#). You can extend the existing partitions or newly added partitions.
- [Create a cloud disk](#) and [attach the cloud disk](#) to the instance as a data disk. After, you need to [partition and format the cloud disk](#).
- [Replace the system disk](#) and specify a higher system disk capacity.

Procedures for disk resizing

The following table describes the process of how to resize a cloud disk based on its current status.

Cloud disk status	Cloud disk that is not attached or partitioned	Cloud disk that is attached but is not partitioned	Cloud disk that is attached and partitioned
-------------------	--	--	---

Resizing procedure	<ol style="list-style-type: none"> 1. Use the ECS console or call the API action <code>ResizeDisk</code> to resize the cloud disk. 2. Use the ECS console or call the API action <code>AttachDisk</code> to attach the cloud disk. 3. Partition and format the cloud disk: <ul style="list-style-type: none"> • Partition and format a Windows data disk less than 2 TiB. • Partition and format a Linux data disk less than 2 TiB. • Partition and format a cloud disk greater than 2 TiB. 	<ol style="list-style-type: none"> 1. Use the ECS console or call the API action <code>ResizeDisk</code> to resize the cloud disk. 2. Use the ECS console or call the API action <code>RebootInstance</code> to restart the instance. 3. Partition and format the cloud disk: <ul style="list-style-type: none"> • Partition and format a Windows data disk less than 2 TiB. • Partition and format a Linux data disk less than 2 TiB. • Partition and format a cloud disk greater than 2 TiB. 	<ol style="list-style-type: none"> 1. Use the ECS console or call the API action <code>ResizeDisk</code> to resize the cloud disk. 2. Use the ECS console or call the API action <code>RebootInstance</code> to restart the instance. 3. Extend the partitions of the attached system disk or data disk: <ul style="list-style-type: none"> • Extend a Windows file system. • Extend a Linux file system. • Extend the file system of a Linux data disk.
--------------------	--	---	---

Limits of resizing a system disk

When you resize a system disk, the new size of the system disk must be greater than its current size, but must be less than or equal to 500 GiB. The following table describes the limits of system disk resizing for different images.

Image	Limit of system disk resizing (GiB)
Linux (excluding CoreOS) + FreeBSD	[max{20, current size of the system disk}, 500]
CoreOS	[max{30, current size of the system disk}, 500]
Windows	[max{40, current size of the system disk}, 500]

For example, if the system disk size of a CentOS instance is 35 GiB, it must be greater than 35 GiB and must be less than or equal to 500 GiB after being resized.

Upper limit of data disk resizing

After being resized, the new size of the data disk must be greater than its current size. The following table describes the upper limit of data disk resizing for different cloud disk.

Cloud disk type	Current size (GiB)	Upper limit of data disk resizing (GiB)
Basic Cloud Disk	< 2,000	2,000
SSD Cloud Disk, Ultra Disk, or ESSD Cloud Disk	< 2,048	2,048
SSD Cloud Disk, Ultra Disk, or ESSD Cloud Disk	$\geq 2,048$	N/A

3.7.2 Resize a cloud disk

This topic describes how to resize a cloud disk. If you resize a cloud disk, you need to restart the instance through the ECS console, or by calling the API `RebootInstance` for the operation to take effect.

Limits

System limits

- For detailed information about the limits of system disk resizing and data disk resizing, see [Disk resizing overview](#).
- Resizing a cloud disk does not extend the file system; it only extends the storage capacity. Therefore, you need to allocate storage space after you resize a cloud disk. For more information, see [What to do next](#).

Supported cloud disk types

- Cloud disks that are in use and are attached to Running or Stopped instances
- Basic cloud disks, ultra disks, SSD cloud disks, and ESSD cloud disks
- NTFS file system (for Windows instances)

Unsupported cloud disk types

- Cloud disks whose snapshot is being created
- Cloud disks of Subscription instances that are included in the remaining billing cycle after such instances are [renewed for configuration downgrade](#)
- System disks of Windows Server 2003 instances

Preparations

[Create a snapshot](#) to back up your data.

Procedure

To resize a cloud disk in the ECS console, follow these steps:

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Block Storage > Disks.
3. Select the target region.
4. Locate the target cloud disk, and then choose More > Resize Disk in the Actions column.
5. Set the Capacity after resizing.



Note:

The capacity after resizing cannot be less than the current capacity.

6. Confirm the price, read and confirm you agree to the ECS Service Terms by selecting the check box, and then click Confirm to resize.
7. Complete the payment.
8. Use the ECS console or call the API ([RebootInstance](#)) to restart the instance for the settings to take effect.

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

What to do next

After you resize a cloud disk, you can perform the following operations as needed.

Cloud disk type	Cloud disk that is not attached or partitioned	Cloud disk that is attached but not partitioned	Cloud disk that is attached and partitioned
-----------------	--	---	---

Procedure	<p>If your cloud disk is in the Available state, the resizing settings take effect immediately after you complete the payment. You can:</p> <ol style="list-style-type: none"> 1. Use the ECS console or call the API (AttachDisk) to attach the cloud disk. 2. Partition and format the cloud disk: <ul style="list-style-type: none"> • Partition and format a Linux data disk less than 2 TiB. • Partition and format a Windows data disk less than 2 TiB. • Partition and format a cloud disk greater than 2 TiB. 	<p>You can perform the following operations as needed:</p> <ol style="list-style-type: none"> 1. Partition and format the cloud disk: <ul style="list-style-type: none"> • Partition and format a Linux data disk less than 2 TiB. • Partition and format a Windows data disk less than 2 TiB. • Partition and format a cloud disk greater than 2 TiB. 2. Extend the file system: <ul style="list-style-type: none"> • Extend the file system of a Linux system disk or Extend the file system of a Linux data disk. • Extend a Windows file system. 	<p>Partition and format the cloud disk and extend the file system:</p> <ul style="list-style-type: none"> • Extend the file system of a Linux system disk or Extend the file system of a Linux data disk. • Extend a Windows file system.
-----------	---	---	---

3.7.3 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](#).

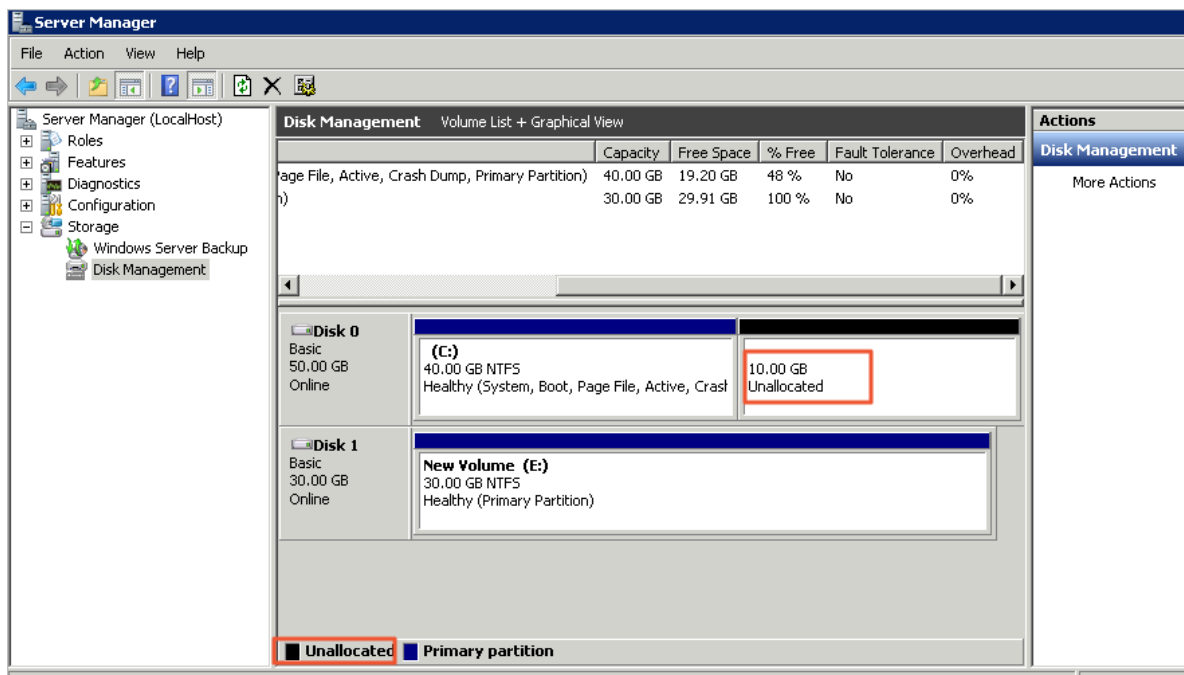
Extend the system disk partition

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.

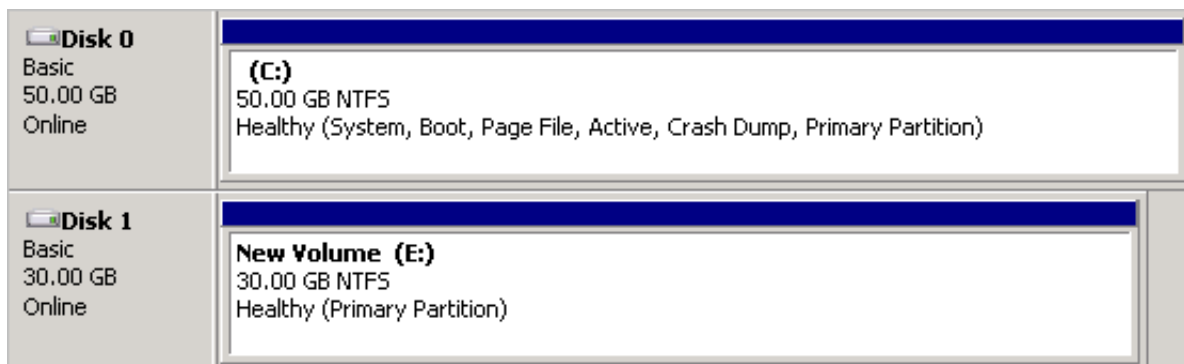


5. In the Disk Management area, view the unallocated capacity. In this example, Disk 0 is the resized system disk.



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.



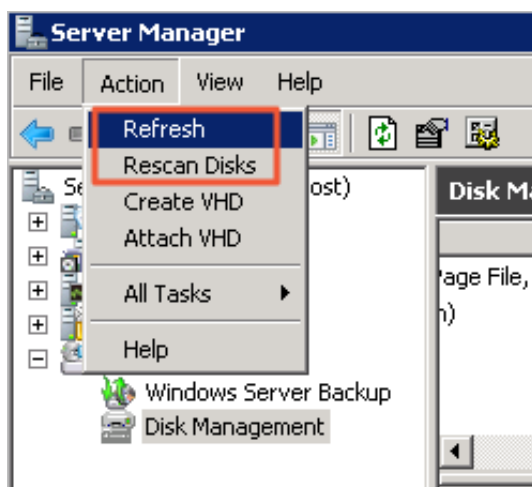
Extend a data disk partition

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.

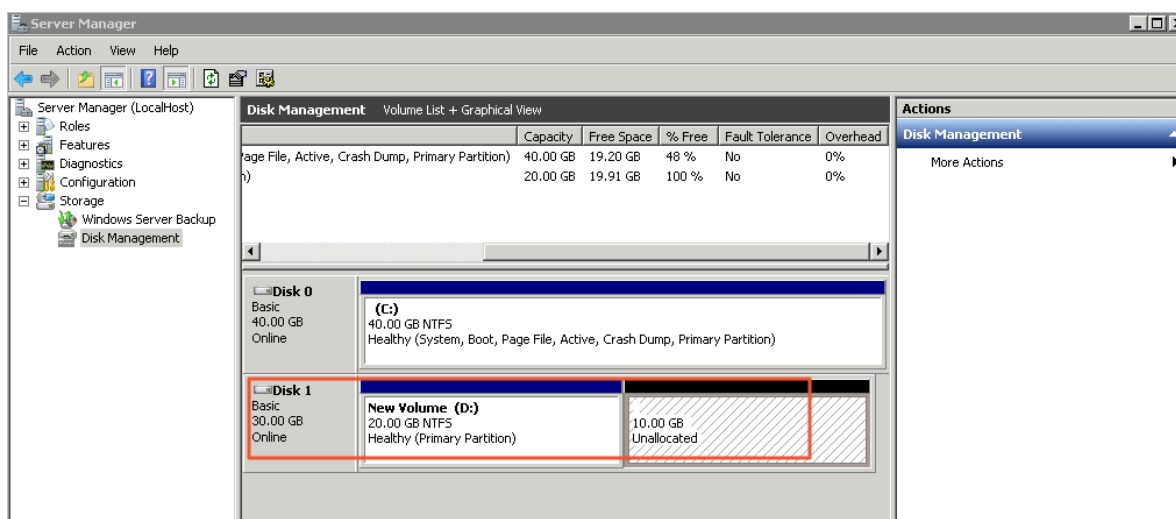
1. [Connect to the Windows instance.](#)
2. Open Server Manager.

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

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

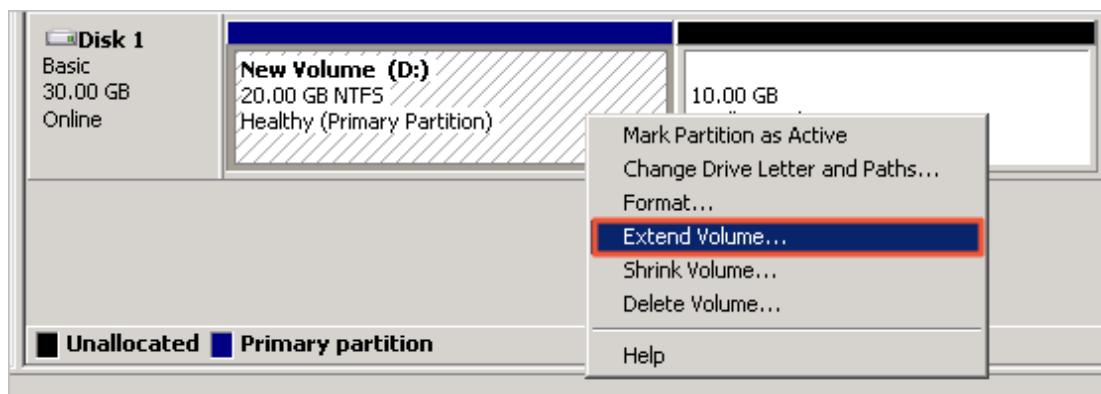


5. In the Disk Management area, view the unallocated capacity. In this example, Disk 1 is the resized data disk.



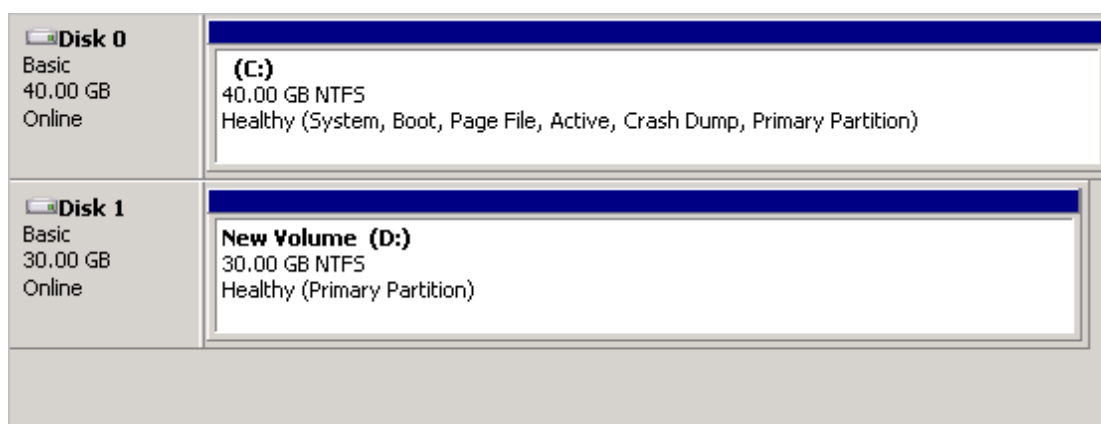
6. Extend Disk 1.

- 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.

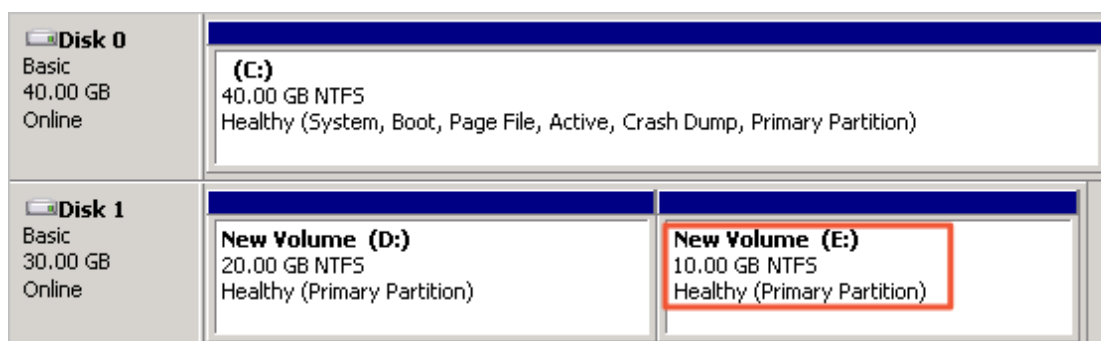


- To use the new disk space to add a new partition, follow these steps:
 - a. Right-click the bank space in the Disk 1 area, and then select New Simple Volume.



- 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.



3.7.4 Extend the file system of the Linux system disk

This topic describes how to use the `growpart` and `resize2fs` tools to resize the system disk and extend the file system of a Linux instance.

Overview

This topic uses the `/dev/vda1` partition of the `/dev/vda` system disk as an example.

- You can resize the disk only if a combination of the following conditions are met. For more information, see [Resize the system disk online](#).
 - Partition format: `mbr` and `gpt`
 - File system: `ext`, `XFS`, `Btrfs`, and `UFS`
 - Operating system: The kernel version is V3.6.0 or later (run the `uname -a` command to check the version).
- If the operating system is earlier than V3.6.0 (for example CentOS 6, Debian 7, and SUSE Linux Enterprise Server 11 SP4), you must restart the operating system before you can resize the system disk. For more information, see [Resize the system disk offline](#).

Preparations

1. Resize the disk by using the ECS console or calling the API.
2. [Create a snapshot](#) to back up data.
3. Mount the disk to an ECS instance that is in the Running state. For information on how to connect to an ECS instance, see [Overview](#).

4. Install the required growpart tool according to your operating system.

Table 3-1: Install the growpart tool in different operating systems

Operating system	Growpart tool
CentOS 7	yum install cloud - utils - growpart
Aliyun Linux	
Ubuntu 14	apt install cloud - guest - utils
Ubuntu 16	
Ubuntu 18	
Debian 9	
Debian 8	Use the upstream growpart tool.
OpenSUSE 42.3	
OpenSUSE 13.1	
SUSE Linux Enterprise Server 12 SP2	

Resize the system disk online

This procedure uses a CentOS 7 instance as an example to describe how to resize a partition online.

1. Install the required growpart tool according to your operating system.

```
[ root @ localhost ~]# yum install - y cloud - utils - growpart
```

2. Run the `fdisk - l` command to view the current disk size. In this example, the disk size (/dev/vda) is 100 GiB.

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

Disk / dev / vda : 107 . 4 GB , 1073741824 00 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
```

3. Run the `df -h` command to view the disk partition size. In this example, the disk partition size (/dev/vda1) is 20 GiB.

```
[ root @ localhost ~]# 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       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
```

4. Run the `growpart <DeviceName><PartionNumber>` command to call the growpart tool to resize the target system disk and partition. In this example, the first partition of the system disk is resized.

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

5. Run the `resize2fs <PartitionName>` command to call the resize2fs tool to extend the file system. In this example, the file system of the /dev/vda1 partition in the system disk is extended.

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

6. Run the `df -h` command to view the size of the disk partition. In this example, the returned partition size (/dev/vda1) is 100 GiB, which means that the partition is extended.

```
[ root @ localhost ~]# 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
```

Resize the system disk offline

This procedure uses a CentOS 6 instance as an example to describe how to resize a partition offline.

1. Install the required `growpart` and `dracut-modules-growroot` tools according to your operating system.

```
[ root @ AliYunOS ~]# yum install -y cloud-utils-growpart dracut-modules-growroot
[ root @ AliYunOS ~]# dracut -f
```

2. Run the `fdisk -l` command to view the current disk size. In this example, the disk size (`/dev/vda`) is 100 GiB.

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

Disk /dev/vda : 107.4 GB, 1073741824 00 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
```

3. Run the `df -h` command to view the disk partition size. In this example, the disk partition size (`/dev/vda1`) is 20 GiB.

```
[ root @ AliYunOS ~]# 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
```

4. Run the `growpart <DeviceName><PartionNumber>` command to call the `growpart` tool to resize the target system disk and partition. In this example, the first partition of the system disk is resized.

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

5. Restart the instance.

```
[ root @ AliYunOS ~]# reboot
```

6. Connect to the instance again.

7. Run the `resize2fs <PartitionName>` command to call the `resize2fs` tool to extend the file system. In this example, the file system of the `/ dev / vda1` partition in the system disk is extended.

```
[ root @ AliYunOS ~]# resize2fs / dev / vda1
resize2fs 1 . 41 . 12 ( 17 - May - 2010 )
Filesystem at / dev / vda1 is mounted on /; on - line
resizing required
old desc_block s = 2 , new_desc_b locks = 7
Performing an on - line resize of / dev / vda1 to
26213807 ( 4k ) blocks .
The filesystem on / dev / vda1 is now 26213807 blocks
long .
```

8. Run the `df - h` command to view the size of the disk partition. In this example, the returned partition size (`/dev/vda1`) is 100 GiB, which means that the partition is extended.

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

Disk / dev / vda : 107 . 4 GB , 1073741824 00 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 13054 104855231 83 Linux
[ root @ AliYunOS ~]# 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
```

3.7.5 Extend the file system of a Linux data disk

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

Limits

The information contained in this topic applies only to disks that are attached to ECS instances (that is, the cloud disk is in a state other than Available). 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](#).

Scenarios

Table 3-2: Scenarios for resizing a data disk of a Linux instance

Scenario		Operation
The data disk is attached to the instance.	The data disk is partitioned and formatted.	You can follow the relevant procedure described in this topic to resize the data disk.
	The data disk is empty and is not partitioned or formatted.	After you resize the data disk in the ECS console, you can follow the relevant procedure describe in this topic to partition and format the data disk .
	The data disk is formatted but not partitioned.	After you resize the data disk in the ECS console, you can follow the procedure to extend the file system .
The data disk is not attached to the instance.		You can resize the data disk offline, or you can resize the data disk after you attach it to the instance.

In this example, the data disk is an ultra disk and the operating system of the ECS instance is CentOS 7.5 64-bit. The mount point of the data disk is `/dev/vdb` and the existing partition is `/dev/vdb1`. The data disk capacity before resizing is 20 GiB whereas the data disk capacity after resizing is 40 GiB. The data disk is attached to the instance and is partitioned and formatted.

- If you want to extend the existing partition of the data disk, see [Procedure 1: Extend the existing partition](#).
- If you want to create a partition for the newly added disk space, see [Procedure 2: Create and format a partition](#).

Preparations

1. Use the ECS console or call the API [ResizeDisk](#) 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](#).

Check the partition and file system

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

In this example, the data disk has the partition `/ dev / vdb1` .

```
[ root @ localhost ~]# fdisk -lu / dev / vdb
Disk / dev / vdb : 42 . 9 GB , 4294967296 0 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 <PartionName>` command to check the file system type.

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

```
[ root @ localhost ~]# blkid / dev / vdb1
/ dev / vdb1 : UUID =" e97bf1e2 - fc84 - 4c11 - 9652 - 73554546c3
24 " TYPE =" ext4 "
```



Note:

If the data disk is not partitioned or formatted, or if the data disk is partitioned but not formatted, no response message is returned.

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

- For ext file systems: `e2fsck -n <dst_dev_part_path>`
- For XFS file systems: `xfs_repair -n <dst_dev_part_path>`

In this example, the file system is in the `clean` state.

```
[ root @ localhost ~]# 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
```



Note:

If the data disk is not partitioned or formatted, you need to partition and format it.
For more information, see [Format a data disk of a Linux instance](#).

Procedure 1: Extend the existing partition

To avoid data loss, we recommend that you do not resize a mounted (`mount`) partition or file system. Instead, we recommend that you unmount (`umount`) the partition first, and then re-mount (`mount`) the partition after you extend it. If you do need to extend a mounted partition, follow these instructions according to the kernel version of the target instance:

- If the kernel version of the instance is earlier than v3.6, unmount the partition, and then modify the partition table by following *Step 1: Modify the partition table* as detailed in the next section.
- If the kernel version of the instance is later than v3.6, you need to update the kernel information after you modify the partition table.

To use the new disk space to extend the existing partition, follow these steps:

Step 1: Modify the partition table

1. Run the `fdisk -lu /dev/vdb` command and note down the start sector and end sector of the old partition.

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

```
[ root @ localhost ~]# fdisk -lu /dev/vdb
Disk /dev/vdb : 42 . 9 GB , 4294967296 0 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 `fdisk` command to delete the old partition.
 - a. Run the `fdisk -u /dev/vdb` command to partition the data disk.
 - b. Enter `p` to print the partition table.
 - c. Enter `d` to delete the partition.
 - d. Enter `p` to confirm that the partition is deleted.
 - e. Enter `w` to save your modifications and exit.

```
[root@localhost ~]# 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.

```

3. 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 print the partition table.
- c. Enter `n` to create a partition.
- d. Enter `p` to set the partition type to primary partition.
- e. Enter `<partition number>` to select a partition number. In this example, select 1.



Warning:

- The start sector of the new partition must be the same as that of the old partition.
- The end sector of the new partition must be greater than that of the old partition. Otherwise, the disk resizing operation fails.

- f. Enter `w` to save your modifications and exit.

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

```
logical[root@localhost ~]# 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.

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 `e2fsck -n / dev / vdb1` command to check the file system to confirm the file system is in the `clean` state.

Step 2: Update the kernel information

If the kernel version of the instance is later than v3.6, after you modify the partition table, run the `partprobe < dst_dev_path >` or `partx -u <dst_dev_path>` command to update the kernel information.

Step 3: Extend the file system

- To extend an ext file system (for example, ext3 or ext4), run the `resize2fs / dev / vdb1` command.

```

[ root @ localhost ~]# 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 .
```

- To extend an XFS file system, run the `mount / dev / vdb1 / mnt /` command, and then run the `xfs_growfs / dev / vdb1` command.

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

meta - data =/ dev / vdb1          isize = 512      agcount =
4 , agsize = 1310720 blks
      =
projid32bi t = 1                  sectsz = 512      attr = 2 ,
      =
spinodes = 0                      crc = 1         finobt = 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
```

Procedure 2: Create and format a partition

To use the new disk space to create a partition, follow these steps:

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

In this example, the new 20 GiB disk space is used to create a partition named `/ dev / vdb2` .

```
[root@localhost ~]# 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 @ localhost ~]# 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 . ext4 / dev / vdb2` command.

```

[ root @ localhost ~]# 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 superblock s and filesystem accounting
informatio n: done

[ root @ localhost ~]# blkid / dev / vdb2

```

```
/ dev / vdb2 : UUID =" e3f336dc - d534 - 4fdd - af71 -
b6ff1a55bd bb " TYPE =" ext4 "
```

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

```
[ root @ localhost ~]# 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 superblock s and filesystem accounting
informatio n : done

[ root @ localhost ~]# blkid / dev / vdb2
/ dev / vdb2 : UUID =" dd5be97d - a630 - 4593 - 9b0f -
5056def914 ea " SEC_TYPE =" ext2 " TYPE =" ext3 "
```

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

```
[ root @ localhost ~]# 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                bsize = 4096      blocks =
data =                    5242880 , imaxpct = 25
                =                sunit = 0        swidth = 0
blks
naming = version 2                bsize = 4096      ascii -
ci = 0      ftype = 1                bsize = 4096      blocks =
log = internal      log                sectsz = 512      sunit = 0
2560 , version = 2
blks , lazy - count = 1                extsz = 4096      blocks = 0
realtime = none
, rtextents = 0

[ root @ localhost ~]# blkid / dev / vdb2
```

```
/ dev / vdb2 : UUID =" 66251477 - 3ae4 - 4b44 - 8b21 -
5604420dbe cb " TYPE =" xfs "
```

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

```
[ root @ localhost ~]# mkfs . btrfs / dev / vdb2
btrfs - progs v4 . 9 . 1
See http :// btrfs . wiki . kernel . org for more
informatio n .

Label : ( null )
UUID : 6fb5779b - 57d7 - 4aaf - bf09 - 82b46f54a4
29
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 @ localhost ~]# blkid / dev / vdb2
/ dev / vdb2 : UUID =" 6fb5779b - 57d7 - 4aaf - bf09 -
82b46f54a4 29 " UUID_SUB =" 9bdd889a - ab69 - 4653 - a583 -
d1b6b87233 78 " TYPE =" btrfs "
```

4. Run the `mount / dev / vdb2 / mnt` command to mount the file system.
5. Run the `df - h` command to view the current disk space and usage.

The information of the new file system is displayed, indicating that the file system is mounted and you can use it.

```
[ root @ localhost ~]# 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
```

3.8 Change the operating system

3.8.1 Replace the system disk (non-public image)

You can replace the system disk if you select an incorrect OS when creating an ECS instance or you need to replace the current OS. The new system disk will be allocated a new ID, and the previous system disk ID will be released.

You can replace the image of the system disk with a public image, shared image, customized image, or any other image from the marketplace.



Note:

Microsoft has ended extended technical support for Windows Server 2003. For the purpose of data security, we recommend that you do not continue running Windows Server 2003 on your ECS instance. Its image is no longer provided. For more information, see [Offline announcement of Windows Server 2003 system image](#).

After you replace the system disk, note that:

- A new system disk with a new disk ID is allocated to your instance, and the original ID is released.
- The cloud type of the cloud disk cannot be replaced.
- The IP address and the MAC address remain unchanged.
- We recommend that you [delete snapshots or automatic snapshot policies](#) to ensure sufficient snapshot quota for executing automatic snapshot policies of the new system disk.

This topic describes how to replace an existing image with a non-public image. If you need to use a public image, see [Replace the system disk \(public image\)](#).

Precautions

Replacing the system disk exposes the system to multiple risks. Read the following sections carefully before you begin:

Risks

Risks of replacing the system disk are as follows:

- Replacing the system disk will stop your instances, which means your business services will be disrupted.
- After replacing the system disk, you must redeploy the service running environment on the new system disk, which may result in a long interruption to services.
- After you replace the system disk, a new system disk with a new disk ID will be assigned to your instance. This means that you cannot use snapshots of the original system disk 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 customized 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 no longer apply and all automatic snapshots of the original system disk will be automatically deleted.

Precautions for cross-OS disk replacement

Cross-OS disk replacement refers to replacing the system disk between Linux and Windows.

**Note:**

Regions outside mainland China do not support disk replacement between Linux and Windows. Disk replacement between Linux editions or Windows editions are supported.

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

- If 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.

- If important data exists in your data disk, perform the following actions as required:
 - From Windows to Linux, you must install a software application, for example, NTFS-3G, because NTFS cannot be identified by Linux.
 - From Linux to Windows, you must install a software application, for example , Ext2Read or Ext2Fsd, because ext3, ext4, and XFS cannot be identified by Windows.

If you replace Windows with Linux, use a password or an SSH key pair for authentication.

Prerequisites

Before replacing the existing image with a non-public image, note the following:

- If the target image is a custom image:
 - If you want to use an image of a specified ECS instance, you must [create a snapshot for the system disk of the specified instance](#) and [create a custom image using a snapshot](#). If the specified instance and the one whose system disk you want to change are located in different regions, you need to [copy the images](#).
 - To use a local physical image file, [import it on the ECS console](#) or [use Packer to create and import the local image](#). The region where the image is located must be the same as that of your instance.
 - To use an image in a region other than that of your instance, you must [copy the image](#) first.



Note:

Imported or duplicated images will be displayed in the Custom Image drop-down list.

- To use an image owned by another Alibaba Cloud account, the account must first [share the image](#).
- If you want to replace the OS to Linux and use an SSH key pair for authentication, you must first [create an SSH key pair](#).
- Replacing the system disk may cause data loss or service interruption. To minimize impact to your business services, we recommend that you [create snapshots](#) for the original system disk before replacement.

- If you want to replace the OS to Linux, make sure that there is sufficient system disk space. We recommend that you reserve 1 GiB in case the OS cannot properly start after system disk replacement.

**Note:**

We recommend that you create snapshots during off-peak hours and plan for sufficient time. For example, to create a snapshot of 40 GiB for the first time, the process takes about 40 minutes. Additionally, creating a snapshot may reduce I/O performance of a block storage device by up to 10%.

Procedure

1. Log on to the [ECS console](#).
2. In the left navigation pane, click Instances.
3. Select the target region.
4. In the Actions column of the target instance, choose More > Instance Status > Stop and follow the instructions in the prompt to stop the instance.

The action is successful when the instance status is Stopped.

5. In the Actions column, choose More > Disk and Image > Replace System Disk.
6. In the displayed dialog box, read the precautionary statement about system disk replacement, and then click OK.
7. On the Replace System Disk page, complete the following settings:
 - a. Image Type: Select Custom Image, Shared Image, or Marketplace 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 space of the system disk you can configured is determined by the current disk space and image type.

Image	Allowed range (GiB)
Linux (excluding CoreOS) + FreeBSD	20-500
CoreOS	30-500
Windows	40-500

**Note:**

If your instance has been configured with [renewal for configuration downgrade](#), you cannot change the system disk space until the next billing cycle.

c. Security enhancement:

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

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

e. Check the configuration and click Confirm to change.

Log on to the ECS console to monitor the system status. It may take about 10 minutes to replace the OS. After the OS is replaced, the instance automatically starts.

What to do next

After replacing the system disk, you can perform the following operations:

- (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 automatically fail. You need to configure automatic snapshot policies for the new system disk.
- If the OS is Linux before and after disk replacement, 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. In this case, you need to write the new partition information to the `/etc/fstab` file of the new system disk and mount the partition, but do not need to partition or format the data disk again. The steps are described as follows. For more information, see [Format a data disk of a Linux instance](#).
 1. (Recommended) Back up the `/etc/fstab` file.
 2. Write information about the new partition into the `/etc/fstab` file.
 3. Check the information in the `/etc/fstab` file.
 4. Run `mount` to mount the partition.
 5. Run `df -h -h` 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](#)

3.8.2 Replace the system disk (public image)

You can replace the system disk if you select an incorrect OS when creating an ECS instance or you need to replace the current OS. The new system disk will be allocated a new ID, and the previous system disk ID will be released.

You can replace the image of the system disk with a public image, shared image, customized image, or any other image from the Alibaba Cloud Marketplace.

**Note:**

Microsoft has ended extended technical support for Windows Server 2003. For the purpose of data security, we recommend that you do not continue running Windows Server 2003 on your ECS instance. Its image is no longer provided. For more information, see [offline announcement of Windows Server 2003 system image](#).

After you replace the system disk, note that:

- A new system disk with a new disk ID is allocated to your instance, and the original ID is released.
- The cloud type of the cloud disk cannot be replaced.
- The IP address and the MAC address remain unchanged.
- We recommend you [delete snapshots or automatic snapshot policies](#) to ensure sufficient snapshot quota for executing automatic snapshot policies of the new system disk.

This article describes how to replace an existing image with a public image. If you need to use a non-public image, see [replace the system disk \(non-public image\)](#).

Precautions

Replacing the system disk exposes the system to multiple risks. Read the following precautions carefully before you begin:

Risks

The risks of replacing the system disk are as follows:

- Replacing the system disk will stop your instances, which means your business services will be disrupted.

- After replacing the system disk, you must redeploy the service running environment on the new system disk, which may result in a long interruption to services.
- After you replace the system disk, a new system disk with a new disk ID will be assigned to your instance. This means that you cannot use snapshots of the original system disk 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 customized 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 no longer apply and all automatic snapshots of the original system disk will be automatically deleted.

Precautions for cross-OS disk replacement

Cross-OS disk replacement refers to replacing the system disk between Linux and Windows.

**Note:**

Regions outside mainland China do not support disk replacement between Linux and Windows. Disk replacement between Linux editions or Windows editions are supported.

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

- If 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.
- If important data exists in your data disk, perform the following actions as required:
 - From Windows to Linux, you must install a software application, for example, NTFS-3G, because NTFS cannot be identified by Linux by default.
 - From Linux to Windows, you must install a software application, for example, Ext2Read or Ext2Fsd, because ext3, ext4, and XFS cannot be recognized by Windows by default.

If you replace Windows with Linux, use a password or an SSH key pair for authentication.

Preparations

- Make sure that there is sufficient system disk space. We recommend that you reserve 1 GiB in case the OS cannot properly start after system disk replacement.
- If you want to replace the OS to Linux and use an SSH key pair for authentication, [create an SSH key pair](#) first.
- Replacing the system disk may cause data loss or service interruption. To minimize impact to your business services, we recommend that you [create snapshots](#) for the original system disk before replacement.



Note:

We recommend that you create snapshots during off-peak hours and plan for sufficient time. For example, to create a snapshot of 40 GiB for the first time, the process takes about 40 minutes. Additionally, creating a snapshot may reduce I/O performance of a block storage device by up to 10%.

Procedure

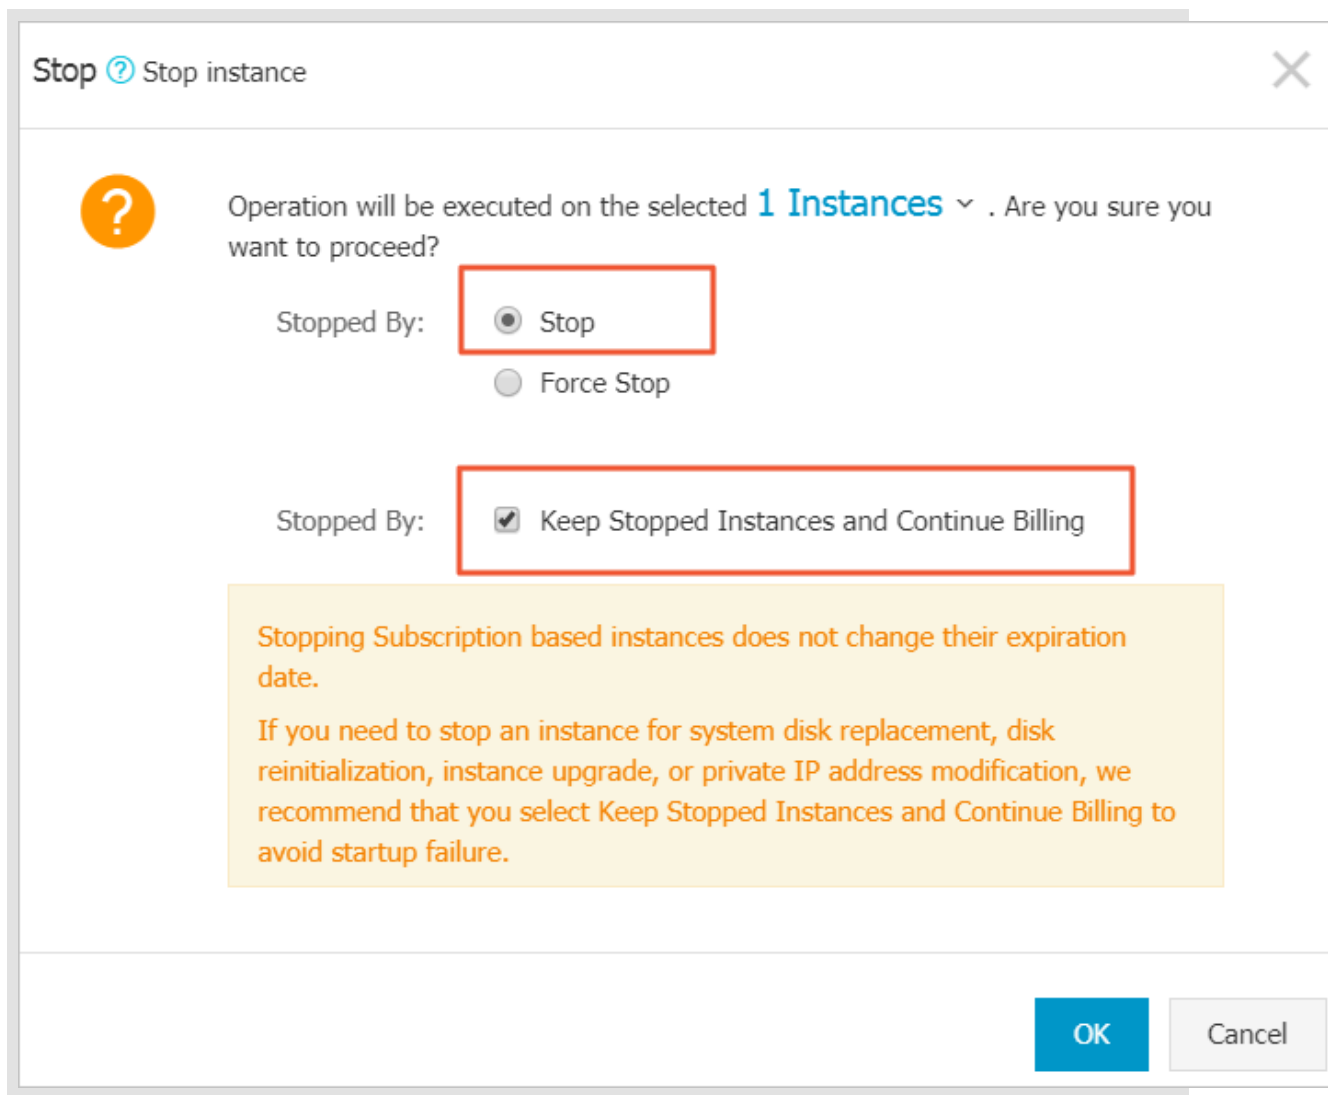
To replace the system disk, follow these steps:

1. Log on to the [ECS console](#).
2. In the left navigation pane, click Instances.
3. Select the target region.
4. In the Actions column of the target instance, select More > Instance Status > Stop and follow the instructions in the prompt to stop the instance.



Note:

If the instance is a Pay-As-You-Go instance using a VPC with the No Fees for Stopped Instances function enabled, in the displayed Notes dialog box, click OK. In the displayed Stop dialog box, select Keep Stopped Instances and Continue Billing. If you select No Fees for Stopped Instances (VPC-Connected), the instance may not be properly started after system disk replacement.



5. After the instance is stopped, in the Actions column, select More > Disk and Image > Replace System Disk.
6. In the displayed dialog box, read the precautionary statement about system disk replacement and then click OK.
7. On the Replace System Disk page, configure the following:
 - a. Image Type: Select Public Image and then select the image version.

**Note:**

If you need to use a non-public image, see [replace the system disk \(non-public image\)](#).

- 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 space of the system disk you can configured is determined by the current disk space and image type.

Image	Allowed range (GiB)
Linux (excluding CoreOS) + FreeBSD	20-500
CoreOS	30-500
Windows	40-500



Note:

If your instance has been configured with renewal for configuration downgrade, you cannot change the system disk space until the next billing cycle.

c. Security enhancement:

- If the new OS is Windows, you can only use a password for 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 either a password or an SSH key pair for authentication. In this case, we recommend you set a login 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 about system disk price, see the [pricing](#) page of ECS.

e. Check the configuration and click Confirm to change.

Log on to the ECS console to monitor the system status. It may take about 10 minutes to replace the OS and update the system status. After the OS is replaced, the instance automatically starts.

Additional operations

After replacing the system disk, you can perform the following operations:

- (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 automatically fail. You need to configure automatic snapshot policies for the new system disk.
- 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. In this case, you need to write the new partition information into the `/etc/fstab` file of the new system disk and mount the partition, but do not need to partition or format the data disk for another time. The steps are described as follows. For more information about operation commands, see [format and mount data disks for Linux instances](#).

1. (Recommended) Back up the `/etc/fstab` file.
2. Write information about the new partition into the `/etc/fstab` file.
3. Check the information in the `/etc/fstab` file.
4. Run `mount` to mount the partition.
5. Run `df -h -h` 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](#)

3.9 Roll back a cloud disk

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

Note

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](#) status.

**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 ? Stop instance

Operation will be executed on the selected **1 Instances**. Are you sure you want to proceed?

Stopped By: ☒ Stop ☐ Force Stop

Stopped By: ☒ Keep Stopped Instances and Continue Billing

Stopping Subscription based instances does not change their expiration date.

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, click Instances.
3. Select the target region.
4. Find the target instance and click its ID to go to its 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 Disk Rollback.
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](#)

Additional operations

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](#)

3.10 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.
- 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 [import ed an SSH key pair](#).

Procedure

To reinitialize a system disk, follow these steps:

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

**Note:**

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

Stop ? Stop instance

?

Operation will be executed on the selected **1 Instances** . Are you sure you want to proceed?

Stopped By:

☒ Stop

☐ Force Stop

Stopped By:

☒ Keep Stopped Instances and Continue Billing

Stopping Subscription based instances does not change their expiration date.

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, in the left-side navigation pane, click Disks.
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:

Security Enhancement: ☒ Activate i

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

- b. (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.
 - c. (Optional) Instance Startup: Select Start Instance Resetting Disk. .
 - d. 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, click Instances.
4. Select the target region.
5. Find the target ECS instance and click its ID to go to its Instance Details page.
6. Click Stop.



Note:

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



7. After the instance is stopped, in the left-side navigation pane, click Disks.
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.

API

[#unique_40](#)

3.11 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.	

3.12 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. Log on to the [ECS console](#).
2. In the left-side navigation pane, select Block Storage > Disks.
3. Select the target region.

4. Find a cloud disk and click its ID to go to the Details page.
5. In the left-side navigation pane, click Disk Monitoring.
6. On the Monitoring Information page, click the  icon and set Start Time and

End Time for monitoring information. You can check the monitoring information of a cloud disk for up to 15 days.




The dialog box for setting monitoring information. At the top, a date range is shown: 2018-06-05 22:38 - 2018-06-05 23:38, with a calendar icon to its right. Below this, there are two rows for Start Time and End Time. Each row has a date selector (showing 2018-06-05), a calendar icon, and a time selector (showing 22:38 for Start and 23:38 for End). At the bottom, there are four tabs: 1Hour(s), 6Hour(s), 1Day(s), and 7Day(s). To the right of these tabs is the text 'Maximum interval of 15 days'. An OK button is located at the bottom right.

7. 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.



3.13 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 status 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 as 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 have been [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, click Instances.
4. Select the target region.
5. Find the target instance and click its ID to go to its 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.
 - Unmountable 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 Unmounted, the disk is detached.

Detach a cloud disks on the Diskspage

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

Prerequisites

The cloud disk has been [attached to the instance](#) and are in the In Use status.

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, select Block Storage > Disks.
4. Select the target 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.
- Unmountable must be Yes.
- Type must be Data Disk.

6. In the dialog box, click Confirm.

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

Related API

[DetachDisk](#)

Additional operation

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

3.14 Release a cloud disk

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

Note

When releasing a cloud disk, note that

- Only cloud disks that are in the Available status 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 status, 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:**

Each cloud disk can have up to 64 snapshots. 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.

- 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, select Block Storage > Disks.
3. Select the target region.
4. Select the disk that you want to release and check it is in the Unmounted status.
Then, in the Actions column, select More > Release.
5. In the Release dialog box, read the note and click Confirm Release.

Related API

[DeleteDisk](#)

4 Shared Block Storage

Shared Block Storage is a block-level storage device that supports concurrent read and write access from multiple ECS instances.



Note:

Currently, Shared Block Storage is available only in public beta release. You can [open a ticket](#) to submit an application to join the testing group.

Shared Block Storage replicates data three times by default, and stores 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), Shared Block Storage can automatically copy your data within the target zone to help ensure data availability. You can partition and format Shared Block Storage attached to an ECS instance, and create a file system and store data on Shared Block Storage. For more information, see [Triplicate technology](#).



Note:

- A single Shared Block Storage device can be attached to a maximum of eight ECS instances in the same zone and region at a time.
- When attached to an ECS instance, Shared Block Storage shares the data disk quota with cloud disks, that is, up to 16 data disks can be attached to each ECS instance.
- Shared Block Storage can only be created separately and used only as data disks.

Background information

In a traditional cluster architecture, multiple computing nodes require access to the same copy of data so that the entire cluster can continue providing business services even when one or more computing nodes fail. If data files are stored in Shared Block Storage devices and these devices are managed through the cluster file system, data consistency can be guaranteed between multiple front-end computing nodes during concurrent read/write operations.

Shared Block Storage is designed for high-availability architectures required by enterprise-level applications.

Limits

Before you use Shared Block Storage, note the following:

Shared Block Storage does not provide a cluster file system. However, you can install a cluster file system to manage Shared Block Storage. Specifically, you can use such cluster file systems as GFS and GPFS to manage Shared Block Storage. If your architecture is a typical Oracle RAC architecture, we recommend that you use ASM to manage storage volumes and data files.

If you use a traditional file system to manage Shared Block Storage that is attached to multiple ECS instances, 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 completed, the space allocation record of Instance A is changed, but the records of 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. As a result, data inconsistency occurs.

Types of Shared Block Storage

The following table describes the available types of Shared Block Storage.

Type	Description
SSD	SSD Shared Block Storage, which uses SSDs as storage media to provide stable and high-performance storage with enhanced random I/O and data reliability.
Ultra	Ultra Shared Block Storage, which uses a combination of SSDs and HDDs as storage media.

Performance

For information about the performance indexes of Shared Block Storage, see [Block storage performance](#).

Billing

Currently, Shared Block Storage is available free of charge. We recommend that you check the official Alibaba Cloud website for the latest billing information.

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 GBps	4 GBps
Read throughput *	min{1.4 * capacity, 2,000} MBps	2 * min{1.4 * capacity, 2,000} MBps
Maximum write throughput	1.2 GBps	2.4 GBps
Write throughput *	min{0.85 * capacity, 1,200} MBps	2 * min{0.85 * capacity, 1,200} MBps
Access latency	Microsecond (μ s)	

* 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 MBps 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 MBps	5,320 MBps
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.



Life cycle

The life cycle 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 life cycle](#).

Impact of instance operations on the data in local disks

The following table describes the impact of instance operations on the data stored in local disks.

Instance operation	Data status of the local disk	Data on the 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
The instance is restored automatically.	Erased	Erased
Release the instance.	Erased	Erased
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	Erased

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

What is an 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 one-way latency.

What is the data reliability of ESSD cloud disks?

ESSD cloud disks deliver 99.9999999% data reliability. This is the same reliability guarantee as with SSD cloud disks and Ultra cloud disks provided by Alibaba Cloud.

What scenarios are ESSD cloud disks suitable for?

ESSD cloud disks are suitable for the following latency-sensitive and I/O-intensive applications:

- OLTP databases such as MySQL, PostgreSQL, Oracle, SQL Server, and other relational databases
- NoSQL databases such as MongoDB, HBase, Cassandra, and other non-relational databases
- Distributed log management systems such as the ELK (Elasticsearch, Logstash, and Kibana)

How do ESSD, SSD, and Ultra cloud disks compare in terms of performance and reliability?

All of these cloud disks types offer high reliability and scalability. They also support creating snapshots and encrypting data. In terms of performance, ESSD cloud disks provide better performance. For more information, see [Block storage performance](#).

How can I measure the performance of an ESSD cloud disk?

The performance of an ESSD cloud disk positively correlates with the relative capacity of the disk. That is, larger capacity corresponds to higher performance. For information about how to measure the performance, see [Block storage performance](#).

How can I test the performance of an ESSD cloud disk and obtain a 1 million IOPS result?

We recommend that you use the FIO tool to test the performance of ESSD cloud disks.

For more information, see [How can I test the performance of an ESSD disk](#).

How is the performance of an ESSD cloud disk and storage performance of the corresponding instance related?

The storage performance of some instances positively correlates to the instance type level. That is, higher instance type levels directly correspond to higher IOPS and throughput numbers.

For example, after you create a g5se instance (a storage-optimized ECS instance type family) and mount ESSD cloud disks to it:

- If the total storage performance of ESSD cloud disks does not exceed that of the instance, the storage performance of ESSD cloud disks takes precedence.
- If the total storage performance of ESSD cloud disks exceeds that of the instance, the storage performance of the instance takes precedence.

Assume that you created a 16 GiB ecs.g5se.xlarge instance that has a maximum IOPS of 60,000. If you mount an ESSD disk that has a capacity of 2 TiB (for which the IOPS is 101,800) to the instance, the maximum IOPS of the instance is 60,000, instead of 101,800.

The performance specifications of other instance types from g5se are described in the following table.

Instance type	vCPU (Core)	Memory (GiB)	IOPS (10 thousand)	Throughput (Mbit/s)
ecs.g5se.large	2	8	3	118
ecs.g5se.xlarge	4	16	6	235
ecs.g5se.2xlarge	8	32	12	470
ecs.g5se.4xlarge	16	64	23	900
ecs.g5se.6xlarge	24	96	34	1350
ecs.g5se.8xlarge	32	128	45	1800

Instance type	vCPU (Core)	Memory (GiB)	IOPS (10 thousand)	Throughput (Mbit/s)
ecs.g5se.16xlarge	64	256	90	3600
ecs.g5se.18xlarge	70	336	100	4000

Can I specify the disk performance attribute of an ESSD cloud disk?

Yes, when you purchase an ESSD cloud disk, you can select a specific disk performance output for that disk. If you require a different disk performance attribute after you purchase the ESSD cloud disk, you can modify the disk performance attribute on the disk details page of the ECS console or by calling `ModifyDiskAttribute`.

How are ESSD cloud disks billed?

You can select the Subscription or Pay-As-You-Go billing method. For more information, see [Prices](#).

In which regions and zones are ESSD cloud disks available?

Currently, ESSD cloud disks are available in the following regions and zones:

- China North 2 Zone G
- China North 3 Zone A
- China East 2 Zone E and Zone F
- China East 1 Zone H and Zone I
- China South 1 Zone E

On what instance types can I mount ESSD cloud disks?

Currently, ESSD cloud disks can be mounted only to instance type families that are based on the 25GE network (namely, c5, ic5, g5, r5, g5se). This is because ESSD cloud disks are available as a beta release only.

What are the common operations that can be performed on a disk?

The following topics describe common operations you can perform on a disk:

- [Format and partition a data disk for Windows Server 2008 instances](#)
- [Format, partition, and mount a data disk for Linux instance](#)

- [Detach a cloud disk](#)
- [Reinitialize a cloud disk](#)
- [Replace the system disk](#)

What is the I/O performance of an SSD cloud disk?

For detailed information, see [Block storage performance](#).

What scenarios are SSD cloud disks suitable for?

SSD cloud disks apply to the following scenarios:

- I/O-intensive applications, such as MySQL, SQL Server, Oracle, PostgreSQL, and other small and medium-sized relational databases
- Small and medium-sized development and testing environments

What is I/O optimization? Can I make a standard ECS instance an I/O-optimized instance?

I/O optimization is a means of providing better storage performance for instances. For example, the storage performance of an SSD cloud disk can be optimized by attaching it to an I/O-optimized instance.

ECS instances that are already purchased do not support I/O optimization, and non-I/O-optimized instances cannot be upgraded to I/O-optimized instances.

Can I use an SSD cloud disk to replace an existing basic cloud disk?

No, as SSD cloud disks use SSD storage media, they cannot be used to replace existing basic cloud disks.

How can I purchase an SSD cloud disk? What are the prices of I/O-optimized instances and SSD cloud disks?

For detailed information, see [Prices](#).

What can I do if one or more of my custom images do not support I/O-optimized instances?

If you want to create an I/O-optimized instance from custom images, open a ticket and include details about the image name. If Alibaba Cloud authorizes the requested image type, you can use it to create I/O-optimized instances.

Can I upgrade the SSD cloud disk after I purchase it?

Yes, an SSD cloud disk can be upgraded and resized.

Why did my Linux system report an error when I attempted to mount an SSD cloud disk to an I/O-optimized instance?

The mount point for SSD cloud disks in the Linux system is `/dev/vdx`, whereas the mount point for basic cloud disks is `/dev/xvdx`. If you use the original tool script for basic cloud disks to mount SSD cloud disks, an error will occur. To resolve this issue, we recommend that you enter the mount point `/dev/vdx` when you manually mount an SSD cloud disk.

Why does my instance shut down when I use FIO to test I/O performance?

The FIO tool supports I/O performance testing on the bare disk partition or on the file system. However, if you directly test bare disk partitions, the metadata in the bare disk partition may be damaged, meaning that you may be unable to access related files. This will cause the instance to shut down. This problem does not occur when you use the FIO file system to test I/O performance.

What test tools and methods can I use to evaluate the performance of my block storage?

For detailed information, see [Block storage performance](#).

What is Shared Block Storage?

ECS Shared Block Storage is a data block-level storage device that supports concurrent reads/writes by multiple ECS instances. It features a high level of concurrency, performance, and reliability. A single Shared Block Storage device can be attached to a maximum of eight ECS instances at the same time.

What are the benefits of Shared Block Storage?

In a traditional cluster architecture, multiple computing nodes require access to the same copy of data so that the entire cluster can continue providing services even if one or more computing nodes fail. If data files are stored in Shared Block Storage devices and these devices are managed through the cluster file system, data consistency can be guaranteed between multiple front-end computing nodes during concurrent read/write operations.

How is Shared Block Storage billed?

Currently, Shared Block Storage is available across all regions free of charge. For the latest billing information, see the official [Alibaba Cloud website](#).

How can I apply for Shared Block Storage?

To use Shared Block Storage, you must first [apply for trial use](#).

What scenarios are Shared Block Storage suitable for?

Shared Block Storage is designed for the high availability architecture required by 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 the Oracle database with Oracle RAC architecture. The Oracle RAC architecture is common among government departments, enterprises, and financial customers.

How can I use Shared Block Storage?

To use block devices, you must first install a cluster file system separately.

You can use cluster file systems such as GFS and GPFS to manage Shared Block Storage. If your architecture is a typical Oracle RAC architecture, we recommend that you use ASM to manage storage volumes and data files.

If you attach Shared Block Storage devices to multiple ECS instances but use a conventional file system to manage them, disk space allocation conflicts and data file inconsistencies 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 completed, the space allocation record of Instance A is changed, but the records of 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. As a result, data inconsistency occurs.

Can I attach Shared Block Storage to instances across regions and zones?

No, a Shared Block Storage device can only be attached to instances in the same zone of the same region.

How many Shared Block Storage devices can be attached to one ECS instance?

Up to 16 data disks can be attached to one ECS instance.

What are the specifications and performance of Shared Block Storage?

Currently, SSD and Ultra Shared Block Storage devices are supported. For more information about their specifications and performance, see [Block storage performance](#).

How can I test the performance of Shared Block Storage?

According to the number of instances involved and the required test item, run the corresponding command as needed.



Note:

When you use the FIO for the performance test, the total iodepth value of multiple clients cannot exceed 384. If the test is conducted simultaneously on four instances, the iodepth value of each client cannot exceed 96.

Number of instances involved	Test item	Command
2	Random write IOPS	<pre>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</pre>
	Random read IOPS	<pre>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</pre>
	Random write IOPS	<pre>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</pre>
	Random read throughput	<pre>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</pre>
4	Random write IOPS	<pre>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</pre>
	Random read IOPS	<pre>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</pre>
116	Write throughput	<pre>FIO - direct = 1 - iodepth = 64 - rw = write - ioengine = libaio - bs = 64k - size = 1G - numjobs = 1</pre>

- If you require lower network latency, we recommend that you deploy ECS instances in the same zone.

Note that an independent Pay-As-You-Go cloud disk must be mounted to an ECS instance in the same zone.

What is an independent cloud disk?

An independent cloud disk is a cloud disk that you can purchase separately, and can be mounted to or unmounted from any ECS instance in the same zone. Before you can use an independent cloud disk, you must first mount it to an instance and set the relevant configurations. For more information, see [Create a cloud disk](#).

Can I mount one cloud disk to multiple ECS instances?

No, a cloud disk can be mounted to only one ECS instance in the same zone.

Do I need to partition and format a Pay-As-You-Go cloud disk after I purchase it and mount it to an ECS instance?

After purchasing an independent Pay-As-You-Go cloud disk, you must mount the disk to an ECS instance and format it. For more information, see [Format a data disk of a Linux instance](#) and [Format a data disk for Windows instances](#).

How is an independent Pay-As-You-Go data disk billed?

The fee for a Pay-As-You-Go data disk is calculated based on the amount of data space you have used. Note that if your account balance is insufficient, the data disk service will be suspended.

Can I attach an independent Pay-As-You-Go data disk to a Subscription instance?

Yes.

Can I detach a data disk from a Subscription instance?

You cannot detach a data disk from a Subscription instance. The data disk is released with instance when the instance expires. If you want to release the data disk, first convert the Subscription data disk to a Pay-As-You-Go data disk, and then detach and release the data disk. For information about how to convert the billing method of cloud disks, see [Convert the billing methods of cloud disks](#).

I changed the configuration of an instance during its renewal. Can I convert a Subscription cloud disk to a Pay-As-You-Go cloud disk within the instance renewal period?

No, you can convert a Subscription cloud disk to a Pay-As-You-Go cloud disk only after the instance renewal period elapses.

Why was my cloud disk that was created separately from its corresponding instance released at the same time as the instance?

This is related to the release setting that is selected when you mount a separately created data disk. That is, when you mount a separately created data disk to an instance, you can choose whether this data disk is released along with the instance. You can change the release setting at any time through the ECS console or by calling the related API.

Why am I unable to find my ECS instance when I attempt to mount a cloud disk?

Check whether the target ECS instance has been released. If the instance is displayed (indicating it has not been released), make sure it is in the same zone of the same region as the cloud disk.

When I delete a cloud disk, will the corresponding snapshots also be deleted?

If you have selected Delete Automatic Snapshots While Releasing Disk, the corresponding automatic snapshots will be deleted when you delete the associated cloud disk. However, corresponding manual snapshots are retained. You can modify this setting at any time by setting the required policy. For more information, see [Apply automatic snapshot policies to disks](#).

Why are some automatic snapshots on my cloud disk missing?

Some automatic snapshots may be missing because the system deletes automatic snapshots (starting with the oldest snapshots) when the snapshot quota is reached. Manual snapshots are not affected.

How many cloud disks can be mounted to one ECS instance?

Up to 16 data disks can be mounted to one ECS instance.

Will the data on my disk be deleted when I unmount a cloud disk (data disk)?

There is a possibility that the data on your disk will be deleted when you unmount a cloud disk (data disk). Therefore, we strongly recommend that you plan accordingly.

- In Windows, we recommend that you stop the read and write operations on all file systems of the cloud 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. After the command is executed successfully, log on to the ECS console to unmount the disk.

Can I unmount the system disk?

You cannot unmount the system disk, but you can replace it. For information about how to replace the system disk, see [Replace the system disk \(non-public image\)](#).

What is a mount point?

A mount point is the location of the Alibaba Cloud ECS disk on the disk controller bus. The mount point matches the disk device number in Linux, and matches the disk order in the Disk Management in Windows.

Are my snapshots retained if I reinitialize a cloud disk?

Yes, both manual snapshots and automatic snapshots are retained when you initialize a cloud disk.

Are my snapshots retained if I replace the system disk?

This depends on the creation method of the snapshot. Manual snapshots will be retained, but automatic snapshots will be deleted.



Note:

The disk ID is changed when you replace the system disk. This means that the snapshots of the original system disk cannot be used for the new system disk.

However, you can use the manual snapshots preserved on the original system disk to create a custom image.

Can I use a snapshot to create an independent cloud disk?

Yes, you can use an existing snapshot to create an independent Pay-As-You-Go cloud disk. The newly created data disk has the same size and use the same data as the snapshot. However, the size of the cloud disk cannot be modified.

What can I do if I cannot access a Linux data disk because of possible data disk mounting errors?

If data of a Linux data disk cannot be accessed, follow these steps to fix the error:

1. Locate the disk where the data is stored and check whether the data disk is mounted to the corresponding ECS instance by using either of the following methods:
 - Perform the check in the ECS console. For more information, see [Monitor a cloud disk](#).
 - Log on to the instance, run the `fdisk -l` command to check whether the data disk partition information is correct, and 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 mount two cloud disks to the same directory.
 - If two cloud disks are mounted to the same directory, the previously mounted cloud disk will be replaced by the new one, resulting in data inaccessibility. We recommend that you mount one of the cloud 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 there are links between the two directories. If there is a link, run the `mkdir` command to create a new directory, mount one of the cloud disks to this directory, and check whether data can be accessed.

How can I resize a system disk?

You can resize a system disk through the ECS console or by calling [ResizeDisk](#). You can also resize a system disk by replacing it with a disk of larger specification.

What considerations should I be aware of before I replace a system disk?

Before replacing a system disk, you must:

- Create snapshots and custom images of the current system disk.
- Make sure that your disk has enough space for storing the snapshots and images. The recommended disk space is 1GiB or larger. If the disk space is insufficient, the system may not start properly after you replace the system disk.

What block storage devices support system disk resizing, and is there any regional restriction for this action?

Ultra cloud disks, SSD cloud disks, and ESSD cloud disks support system disk resizing, and all regions of the Alibaba Cloud ecosystem support system disk resizing.

Can I resize the system disk of Subscription and Pay-As-You-Go instances?

Yes, the system disk of both Subscription and Pay-As-You-Go instances can be resized.

What is the storage capacity range of a system disk?

The capacity of a system disk varies depending on the operating system. For more information, see [Disk resizing overview](#).

I downgraded the configuration of a Subscription instance during its renewal. Can I specify a new system disk capacity after the renewal is completed?

After you [downgrade the configuration](#) of a Subscription instance through instance renewal, you can resize the system disk of the instance only at the start of the next billing cycle.

Can I scale in a system disk after I scale it out?

No, you cannot scale in a system disk after you scale it out.

Can I partition a data disk for data storage?

Yes, you can divide a data disk into multiple partitions according to your specific 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, rather than for a specific partition.

What considerations should I be aware of before I repartition a disk?

To ensure data security, we recommend that you [create a snapshot](#) of the target disk before you can repartition it. This way, you can roll back the disk if an exception occurs.

I rolled back the snapshot of a data disk after repartitioning it. How is the number of partitions changed?

Snapshot rollback returns a snapshot to its previous state. This means that if the disk has not been repartitioned at the time point of the previous state, there is only one partition.

How can I ensure no data is lost during a data disk scale-out if I need to create a cloud disk from a snapshot?

If the original data disk cannot be scaled out due to a disk error, you can purchase a temporary Pay-As-You-Go cloud disk to store data and then format the original data disk. To do so, follow these steps:

1. Create a snapshot of the current data disk.
2. Go to the [disk purchase](#) page to purchase a Pay-As-You-Go cloud disk, select the same region and zone as the ECS instance, and click Create from snapshot.
3. Log on to the ECS console and then mount the data disk you purchased to an ECS instance.
4. Log on to the ECS instance, run the `mount` command to mount the purchased disk to the ECS instance, and then check whether files in this disk are the same as those in the original data disk.
5. Run the `fdisk` command to delete the original partition table, and then run commands such as `fdisk` and `mkfs . ext3` to format and partition the original data disk again, so that the available space of the original data disk is the same as that after the scale-out action.
6. Run the `cp -R` command to copy all the data in the purchased disk back to the original data disk. You can add the `--preserve = all` parameter to preserve the file's properties when copying files.
7. Run the `umount` command to unmount the purchased data disk.
8. Log on to the ECS console, unmount the purchased data disk from the ECS instance, and release it.

In a Linux instance, what considerations should I be aware of before I add mount information to a basic cloud disk or an SSD cloud disk?

When you add a data disk to the Linux system and add the partition information according to [Format a data disk of a Linux instance](#), note that the mount point of a basic cloud disk is `/dev/xvdb1` and that of an SSD cloud disk is `/dev/vdb1`. If you add invalid information, disk mounting will fail. To avoid this issue, follow these steps:

1. Run the `fdisk -l` command to view data disk information.

2. Check whether the information added to `/ etc / fstab` is valid. Do not add mount information repeatedly 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, add valid mount information, and then run the `mount - a` command to check whether the data disk is successfully mounted.

How can I re-mount data disks after initializing a Linux instance?

Initializing the system disk of a Linux instance does not change data on the data disk. However, all the data disks will be unmounted. Therefore, you need to mount the data disks again. To do so, follow these steps:



Note:

In this example, assume that the mount point is `/ dev / vdb1` and the mount point name is `/ InitTest` before the system disk is initialized.

1. Run the `mount` command to view the data disk mount status.

The command output does not contain information related to `/ dev / vdb1` .
2. Run the `fdisk - l` command to view partitioned data disks.
3. Run the `mkdir / InitTest` command to re-create the mount point of the data disk partition. The mount point name must be the same as that of the mount point `/ dev / vdb1` before the system disk initialization. You can view the original mount point name by running the `cat / etc / fstab` command.
4. Run the `mount / dev / vdb1 / InitTest` command to remount the data disk partition.
5. Run the `df - h` command to view the mount results.
6. 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 successfully unmounted. If this is the case, the returned result does not contain information related to `/ dev / vdb1` .
 - c. Run the `mount - a` command to automatically mount `/ dev / vdb1` .
 - d. Run the `mount` command to check whether the auto-mount succeeded. If this is the case, the returned result contains information related to `/ dev / vdb1` .

What can I do if data disks of a Linux instance cannot be found after I restart the instance or initialize the system?

Symptom: After restarting a Linux instance or initializing the system, I log on to the instance and run the `df -h` command to check the disk mount status. However, no data disk can be found.

Cause:

- **Restarting an instance:** Mount information was not written to the `/etc/fstab` file when the data disk was mounted. As a result, data disks are not automatically mounted after the instance restarts.
- **Initializing the system:** The `/etc/fstab` file is reset after the system disk is initialized. As a result, data disks are not automatically mounted upon system startup.

Solution:

1. Run the `mount /dev/xvdb1` command to mount the data disk.
2. Run the `mount` command to check the format of files in the `/dev/xvdb1` partition.

In this example, the file format is ext3.

3. Run the following command to write the data disk mount information to `/etc/fstab` :

```
echo '/dev/xvdb1 /data ext3 defaults 0 0' >> /etc/fstab
```

What can I do if data is lost after I restart a Linux instance?

Symptom: All data in the `/alidata` directory is lost after I restart a Linux instance.

Cause: Output of the `df -h` command shows that no data disk is mounted to the instance.

Solution:



Note:

This solution uses an I/O-optimized instance as an example. If you use a non-I/O-optimized instance, the disk partition is `/dev/xvdb1` .

1. Run the `fdisk -l` command.

The output shows that one disk is not mounted.

2. Run the `mount /dev/vdb1 /alidata` command to mount this disk to the instance.

We recommend that you also enable auto-mount at startup for `/etc/fstab` to avoid encountering this issue again.

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 the cloud disk of a Linux instance?

Symptom: The error message "Bad magic number in super-block while trying to open /dev/xvdb1" is returned when I run the `e2fsck -f /dev/xvdb` command to resize a cloud disk.

Cause: There is no disk partition

Solution: Run the `e2fsck -f /dev/xvdb` and `resize2fs /dev/xvdb` commands to resize the disk. Then mount the disk to the instance.

What can I do if an error message similar to "Only disks on stopped instances and disks with no ongoing snapshot tasks and no replaced operating system can be rolled back." is returned?

Symptom: An error message similar to "Only disks on stopped instances and disks with no ongoing snapshot tasks and no replaced operating system can be rolled back ." is returned when I attempt to roll back a cloud disk.

Cause: The problem is usually caused by incorrect disk attributes or disk status.

Solution:

- Check whether the operating system of instances associated with related snapshots has been replaced. For information about how to replace an operating system, see [Replace the system disk](#).

If the operating system has been replaced, the system disk of the instance will be automatically re-created from the new image and the system disk ID will change. Therefore, original snapshots cannot be used for the rollback. However, you can [create a custom image](#) from related snapshots and then specify the custom image by [replacing the system disk](#) to switch the instance to the corresponding snapshot state.

- Check whether the instance has stopped.

You can only roll back a stopped instance. You can view the instance status on the Instances page of the ECS console.

- Check whether the disk has an ongoing snapshot task.

To ensure data consistency, do not roll back a disk that has an ongoing snapshot task. You can check the snapshot status on the Snapshots page. A snapshot task is ongoing if its Progress is not 100% and Status is Progressing.

If you need to forcibly terminate related snapshot tasks to accelerate the rollback, select the corresponding snapshots and click Delete.

How can I copy data across instances?

You can choose data copy methods based on your operating system:

- Data copy between Linux instances

- Using the lrzsz tool

Log on to the Linux instance, install the lrzsz tool, run the `rz` command to upload files, and run the `sz` command to download these files.

You can also first run the `sz` command to download files to your local PC and run the `rz` command to upload these files to another instance.

- Using 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.

- Using 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 another instance to download the file or folder.

- Data copy between Linux and Windows

We recommend that you use the SFTP tool to download files from a Linux instance to your local PC and use the FTP tool to upload data to a Windows instance.

- Data copy between Windows

- FTP tool
- Alibaba TradeManager

Can I mount a cloud disk to an ECS instance that is in a different zone from the cloud disk?

No, you can only mount a Pay-As-You-Go cloud disk to an ECS instance that is in the same zone as the cloud disk.

Why am I unable to find the data disk I purchased for a Linux instance?

If you purchase a Pay-As-You-Go data disk separately, you need to [format the data disk](#) and [attach it](#) to a Linux instance before you can use it.

Where can I purchase a last-generation disk (a local SSD disk) and how can I maintain an existing local SSD disk?

Local SSD disks are no longer available for purchase. If you are still using a local SSD disk, see [Local disks](#) and [FAQs about SSD cloud disks](#).

How can I test the performance of an ESSD cloud disk?

This FAQ describes how to configure the appropriate environments for a performance test and how to obtain a 1 million IOPS result. Note that the cloud disk type and actual environment used will impact the test result.



Warning:

You can obtain accurate test results by testing a bare disk partition, but you may destroy the file system structure by testing it directly. Before such a test, you must [create snapshots](#) 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 any data loss.

We recommend that you use the latest versions of Linux images in Alibaba Cloud official images, such as CentOS 7.4/7.3/7.2 (64-bit) and AliyunLinux 17.1 (64-bit) because Linux images of earlier versions and Windows images may use defective drivers.

We also recommend that you use the [FIO](#) for the performance test.

Assume that the instance type is ecs.g5se.18xlarge and the device name of the ESSD disk is `/dev/vdb`. This example describes how to test the random write (`randwrite`) performance of an ESSD cloud disk.

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 change the directory.
4. Run the `vim test100w . sh` command to create a 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 tested threads ,
such as 8 in this example .
    iodepth=$2 # The maximum concurrent I / O requests ,
such as 64 in this example .
    bs=$3 # The size of the data block for one
I / O , such as 4k in this example .
    rw=$4 # The read and write policy , such as
randwrite in this example .
    filename=$5 # The name of the tested file , such
as / dev / vdb in this example .
    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 8 64 4k randwrite / dev / vdb
```



Note:

- You must modify the following commands according to your test environment.
 - `vdb` in the command line `list = `cat / sys / block / vdb / mq / */ cpu_list | awk '{ if (i <= NF) print $ i ;}' i ="$ i " | tr - d ',' | tr '\ n ' ','``
 - `8 , 64 , 4k , randwrite , and / dev / vdb` in the command line `RunFio 8 64 4k randwrite / dev / vdb`
- The file system structure may be damaged if you test the bare disk partition directly. If you choose to proceed with this action, set `filename` as the device name, such as `/ dev / vdb` . If you do not want to risk data loss, set `filename` as a file path, such as `/ mnt / test . image` .

5. Run `sh test100w . sh` to start testing the performance of the ESSD disk.

```
[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 parameter

The command `echo 2 > / sys / block / vdb / queue / rq_affinit y` in the script is used to set the value of the `rq_affinit y` parameter to `2` .

- If the value of the `rq_affinit y` parameter is `1` , the block device migrates the I/O Completions to the vCPU group that originally submitted the request. In this way, I/O Completions may run on the same vCPU in the case of concurrent processing I/O requests by multiple threads. This may cause a performance bottleneck.
- If the value of the `rq_affinit y` parameter is `2` , the I/O Completions are forced to run on the requesting vCPU. As a result, performance of each vCPU is fully used in the case of concurrent processing I/O requests by multiple threads.

- Binding threads to corresponding vCPU cores
 - Generally, a device has only one Request-Queue. This unique Request-Queue is a bottleneck in the case of concurrent processing I/O requests by multiple threads.
 - 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 vCPU core corresponding to each Request-Queue. This allows you to fully use the Multi-Queue mode to improve the storage performance.

To fully output the performance of the device, 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 vCPU cores. `/ dev / vd *` is the device name of your ESSD cloud disk, 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_repo rting - cpus_allow ed = $ spincpu - cpus_allow
ed_policy = split
```

FIO provides the `cpusallowe d` and `cpus_allow ed_policy` parameters to bind vCPUs. The preceding command runs multiple `jobs`. They are bound to different CPU cores and correspond to different `Queue_Id`.

To view the `cpu_core_i d` corresponding to a `Queue_Id`, follow the instructions:

1. Run the `ls / sys / block / vd / mq /` command to check the `Queue_Id` of the ESSD disk with the device name `/ dev / vd *`, for example, `/ dev / vdb`.
2. Run the `cat / sys / block / vd / mq // cpu_list` command to check the `cpu_core_i d` corresponding to the `Queue_Id`.