Alibaba Cloud Elastic Compute Service

Block storage

Issue: 20190409

MORE THAN JUST CLOUD |

Legal disclaimer

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

- 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 cor	nventions
-----------	-----------	-----------

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 informatio n, 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 cd / d C :/ windows command to enter the Windows system folder.
Italics	It is used for parameters and variables.	bae log list instanceid Instance_ID
[] or [a b]	It indicates that it is a optional value, and only one item can be selected.	ipconfig [-all -t]

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

Contents

Legal disclaimer	I
Generic conventions	I
1 What is block storage?	1
2 Storage parameters and performance test	3
3 Block storage	12
2 1 Cloud diales and Charad Plack Starson	12
2.2 Triviliante technology	12
3.2 Implicate technology	16
3.3 ECS disk encryption	1/
	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	28
3.6 Format a data disk	32
3.6.1 Partition and format data disk more than 2 TiB	32
3.7 Monitor a cloud disk	40
3.8 Resize cloud disks	
3.8.1 Overview	42
3.8.2 Increase system disk size	43
3.8.3 Windows - Resize a data disk	49
3.8.4 Linux - Resize a data disk	53
3.9 Roll back a cloud disk	59
3.10 Reinitialize a cloud disk	61
3.11 Change the operating system	69
3.11.1 Replace the system disk (non-public image)	69
3.11.2 Replace the system disk (public image)	74
3.12 Convert billing methods of cloud disks	81
3.13 Detach a cloud disk	82
3.14 Release a cloud disk	85
4 Local disks	87
T LUCAI UISKS	07

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 distribute d storage architecture and local disks located on the physical servers where ECS instances are hosted. Specifically, the storage products are as follows:

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

For more information about the performance of block-level storage products, see *#unique_6*.

Block storage, OSS and NAS

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

the following three types of data storage products:

- Block storage: A high-performance and low-latency block-level storage device for ECS. It supports random reads and writes. You can format block storage and create a file system on it as you would with a hard disk. , thereby enabling block storage to meet the data needs of numerous business scenarios.
- OSS: A huge storage space designed for storing massive amounts of unstructur ed 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.

 \cdot 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 op	erations 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).



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 **	IOPS = min{ 1800 + 50 x capacity, 1,000 ,000} Throughput =	IOPS = min{ 1800 + 30 x capacity, 25, 000} Throughput =	IOPS = min {1800 + 8 x capacity, 5,000 } Throughput =	N/A N/A
	min {120 + 0.5 x capacity, 4, 000} MBps	min{120 + 0.5 x capacity, 300} MBps	min{100 + 0.15 x capacity, 140 } MBps	
Data reliability	99.9999999%	99.9999999%	99.9999999%	99.9999999%
API name	cloud_essd	cloud_ssd	cloud_effi ciency	cloud

Parameter	ESSD Cloud Disk	SSD Cloud Disk	Ultra Cloud Disk	Basic Cloud Disk
Scenarios	 OLTP databases : relational databases such as MySQL, PostgreSQL, Oracle, and SQL Server NoSQL databases : non- relational databases such as MongoDB, HBase, and Cassandra ElasticSea rch distribute d logs: Elasticsea rch, Logstash and Kibana (ELK) log analysis 	 Large and medium -sized relational databases , such as MySQL, SQL Server, PostgreSQL, and Oracle Large or medium -sized developmen t or testing applications that require high data reliability 	 Small or medium sized relational databases such as MySQL, SQL Server, and PostgreSQL Large or medium sized developmen t or testing applications that require high data reliabilit y and medium performanc e 	 Applicatio ns with infrequent access or low I/O load. If higher I/O performanc e is needed , we recommend that you use SSD disks. Applicatio ns 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 Bock Storage	Ultra Shared Block Storage
Capacity	Singe disk: 32,768 GiBSingle instance: 128 TiB	Singe disk: 32,768 GiBSingle 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	IOPS = min{1600 + 40 x capacity, 30,000}	IOPS = min{1000 + 6 x capacity, 5,000}
disk**	Throughput = min{100 + 0 .5 x capacity, 512} MBps	Throughput = min{50 + 0. 15 x capacity, 160} MBps

Parameter	SSD Shared Bock Storage	Ultra Shared Block Storage
Scenarios	 Oracle RAC SQL Server Failover cluster High-availability architecture of servers 	 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 in 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.



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	Indicates that the read and write policy is random write. Other options include: • randread (random read) • read (sequential read) • write (sequential write) • randrw (random read and write)
-ioengine=libaio	 Use libaio as the testing method (Linux AIO, Asynchronous I/O). Usually there are two ways for an application to use I/ O: Synchronous 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. Asynchronous The asynchronous method uses libaio to submit a batch of I/O requests each time, thus reducing interaction times and making interactions more effective.
-bs=4k	Indicates the size of each block for one I/O is 4 KiB. If not specified, the default value 4 KiB is used. When IOPS is tested, we recommend that you set bs to a small value, for example, such as 4k in this example command. When throughput is tested, we recommend that you set bs to a large value, such as 1024k in this example command.

Parameter	Meaning
-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 - size block by block, with the data block size specified by - bs.
-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 performanc e, and high data reliability
- Basic cloud disks: with high data reliability and general random I/O performanc e
- 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. A data disk created with an ECS instance has the same life cycle as the instance, and is created and 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 unauthoriz ed access. Your data and the associated keys are encrypted based on the industry standard AES-256 algorithm.

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

The key management infrastructure of Alibaba Cloud conforms to the recommenda tions 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 *RelnitDisk*.

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 unencrypte d, 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.



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 Disks list 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 CreateDisk.

To create a cloud disk when creating an instance, see RunInstances or CreateInstance.

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

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, click Instances.



3. 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-

4. Click Configuration Change > Add Subscription Cloud Disk.

	Buy Same Type				
Subscription Manage 1 March 2019, 00.00	Instance Status				
Expire	Instance Settings				
Pay- Switch to Pay-As-You-Go	Password/Key Pair				
11.5 Add Subscription Cloud Disk	Configuration Change				
Pay-As-You-Go	Disk and Image				
8 November 2018, 13.57 Create	Network and Security Group				
	Operations and Troubleshooting				

- 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.

😫 Disk	SSD Cloud Disk 🗸	20	GiB 2400 IOPS Create	from snapshot 🔄 Encrypted				
To compare SSD Disks, Ultra Disks, and Basic Disks, see Details>								
	Quantity:	1 Attachable: 15, A	Unit ttached: 1					
	Disk Name:	Enter a disk nar	ne	The name can be 2 to 128 characters in length and can contain letters, Chinese characters, numbers, hyphens (-), underscores (_), and periods (.). It must start with a letter or Chinese character.				
	Description :	Enter a descript	ion	The description can contain 2 to 256 characters. It cannot start with http:// or https://.				

- 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.

Instance Details	o launch-advisor-20190	13								c
Disks Instance Snapshots	Disks									Create Disk Mount
Security Groups	ID/Name	Tag	Disk Type(All) 👻	Status (All) 👻	Billing Method(All) -	Unmountable(All) -	Zone	Type(All) -	Encrypted/Unencrypted	Actions
Security Protection	d-möshavarsi21000amilo 0	۲	SSD 20GB	In Use	Subscription	No	Qingdao Zone C	Data Disk	Unencrypted	Create Snapshot Reinitialize Disk Create Automatic Snapshot Policy More+
	 d-maneskacija invikancija O - 	۲	Ultra Disk 20GB	In Use	Subscription	No	Qingdao Zone C	Data Disk	Unencrypted	Create Snapshot Reinitialize Disk Create Automatic Snapshot Policy More -
	indexplotedourhowske (i 0	۲	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 article describes how to create a cloud disk by using a snapshot in the ECS console. You can take a snapshot of an existing system disk, or data disk, and create a cloud disk by using the snapshot. The new disk can be attached to any instance in the same zone of the same region.

Scenario

If you need to access data from a snapshot, but do not want to *roll back a cloud disk*, you can create a cloud disk by using the snapshot to access data that you need. For example, if your instance encounters a system disk failure, you can use an existing snapshot to create a cloud disk, and attach the disk to a healthy instance. By doing so, you can restore the data of the affected instance.

Disk Performance

If a cloud disk is created by using a snapshot, the initial disk performance decreases because data needs to be accessed from OSS before being written into the disk. We recommend that you write and read every data block at least once before production use. For more information about OSS, see *what is OSS*.

Considerations

• Only *Pay-As-You-Go* cloud disks can be created in this way, and they can only be used as data disks.

Note:

You can set cloud disks as data disks when creating an ECS instance. The disks then have the same billing method as that of the instance.

You can create a new empty cloud disk. For more information, see create a cloud disk.

- The quota of Pay-As-You-Go cloud disks that are used as data disks of each account in all regions is five times the quota of 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 of each other, and you cannot merge their space by
 formatting. We recommend that you confirm the amount and size required before
 you create cloud disks.
- You can create a snapshot for a single cloud disk, so we do not recommend that you create LVM (Logical Volume Manager) volumes, which may cause data loss when you use the snapshot to rollback the cloud disk.
- After a Pay-As-You-Go cloud disk is created, you can convert its billing method to Subscription:
 - If it is attached to a Subscription instance, use the *upgrade* configurations of Subscription instances feature.
 - If it is attached to a Subscription instance, use the *switch from Pay-As-You-Go to Subscription* feature.
- If a cloud disk is created in this way, and its billing method is not converted, you can detach a cloud disk and release a cloud disk at any time.

Prerequisites

- You have created a snapshot for your instance, and you make sure the region and zone. For specific actions, see *create snapshots*.
- *Attach a cloud disk*. The instance and the cloud disk must be 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 Disks list page, click Create Disk.
- 4. Select a 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. Configure the cloud disk:
 - a. Select a cloud disk category. The category of the source disk of the snapshot does not modify the configuration.
 - b. Click Create from snapshot and select a snapshot.
 - c. Specify the size of the cloud disk. The size range is 20 GiB to 32768 GiB. If the selected snapshot is smaller than 20 GiB, you can adjust the size manually. For a snapshot larger than 20 GiB, the size is adjusted automatically according to the snapshot size. However, if you replace the snapshot, you must manually set the size.
 - d. For Purchase Plan, set the quantity.
- 6. Check the 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 operation

Attach a cloud disk.

Related API

Create a cloud disk: CreateDisk

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.
- $\cdot\,\,$ 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.
A cloud disk created from a snapshot	Linux	Connect to the Linux instance and run the mount < partition > < mount point > 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 mount 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 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.



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

Operating system	Partition tool	File system
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:

- he system expands the block-level of the device's disk, but does not automatica lly 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_2 -1 clusters. However, in Windows XP Pro, the maximum volume of NTFS is 2 -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 MNR 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 2₂ -1 clusters. Howver, in Windows XP Pro, the maximum volume of NTFS is 2 -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
                                    3221225472
                              GB ,
                                                000
                                                      bytes ,
6291456000
            sectors
                   of
                          * 512 = 512
Units = sectors
                        1
                                           bytes
        size ( logical / physical ):
Sector
                                     512
                                            bytes
                                                   /
                                                      512
bytes
I / 0
       size ( minimum / optimal ): 512
                                                    512
                                          bytes /
                                                          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 < StartSecto r > < 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
                                                      properly
                                           is
                                                not
                        performanc e.
aligned
          for
                 best
Ignore / Cancel ? ignore
( parted ) print
Model : Virtio Block
                           Device (virtblk)
                   Block
Disk / dev / vdd : 3221
                            GB
Sector size (logical / physical): 512B / 512B
Partition Table : gpt
Disk Flags :
                  End
                        Size File
Number
          Start
                                       system
                                                Name
                                                       Flags
    17 . 4kB
                                  primary
1
                3221GB
                         3221GB
```

- 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 - 0 64bit , has_journa l , extents , huge_file , flex_bg
, uninit_bg , dir_nlink , extra_isiz e / 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 @	j izXXX	(Xz ~]	# df	– h			
Filesys	stem	Size	Used	Avail	Use %	Mounted	on
/ dev /	vda1	40G	6.4G	31G	18 % /		
devtmpf	fs 48	37M 0	487M	0%/	dev		
tmpfs	497M	0	497M	0 % / dev	v / shm		
tmpfs	497M	364K	496M	1%/	run		
tmpfs	497M	0	497M	0 % / sys	s / fs ,	/ cgroup	
tmpfs	100M	0	100M	0 % / ru	n / usei	r / 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.

of / dev / vdd mkfs . ext4 : Size device too big to be in expressed 32 bits using а 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
e2fsproas-libs-	1.39-33.1.alios5
e2fsprogs-devel	-1.39-33.1.alios5

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 version1.42.8. 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 xvzf 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 - 0 64bit , has_journa l , extents , huge_file , flex_bg ,
uninit_bg , dir_nlink , extra_isiz e - E lazy_itabl e_init = 0 ,
lazy_journ al_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 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.

	2018-06-05 2	2:38 - 2	2018-	06-0	05 23:	38	
Start Time :	2018-06-05	i		22	^ ~	38	^ ~
End Time :	2018-06-05	i		23	<u>^</u> :	38	^ ~
1Hour(s) 6 15 days	Hour(s) 1Day	y(s) 7Day	(s)	Maxi	imum ir	nterva	l of OK

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





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

3.8 Resize cloud disks

3.8.1 Overview

Depending on the disk type, you can resize a disk as follows:

- For a system disk: Change System Disk
- For a data disk: Resize Disk

Limitations

Limitations of resizing disks vary between system disks and data disks.

System disks

The Change system disk feature allows you to increase the disk size only. The size limit for disk resizing is determined by the image and the current size of the system disk, as displayed in the following table.

Image	Size limit (GiB)
Linux (excluding CoreOS) and FreeBSD	20-500
CoreOS	30-500
Windows	40-500

Data disk

The Resize Disk feature allows you to increase the disk size only. The following table lists the capacity limits of different data disk typics after resizing, which is determined by the cloud disk types.

Cloud disk type	Current capacity	Capacity after resizing
Basic Cloud Disk	Any	2,000 GiB
SSD Cloud Disk or Ultra Cloud Disk	equal or less than 2,048 GiB	2,048 GiB
SSD Cloud Disk or Ultra Cloud Disk	> 2,048 GiB	Cannot be resized
ESSD Cloud Disk	Any	32,768 GiB

Additional operations

- To increase the size of the system disk of an ECS instance, see *increase system disk* size.
- To resize a data disk attached to a Windows instance, see *Windows Resize a data disk*.
- To resize a data disk attached to a Linux instance, see Linux Resize a data disk.

3.8.2 Increase system disk size

You can increase the size of the system disk of your ECS instance by using Change System Disk feature. This article describes how to increase the size of a system disk while keeping the operating system and environment intact.



You can change the operating system while increasing the size of a system disk. For more information, see *change the operating system*.

Notes

Before you begin, consider the following.

Risks

Risks that may occur when you replace a system disk are as follows:

- If you attempt to replace the system disk while the instance is running, your business services may be disrupted. We recommend you stop your instance before replacing the system disk.
- After disk replacement, you must redeploy the business runtime environment on the new system disk. This may result in a long period of disruption to your business services.
- After the system disk is changed, a new cloud disk with a new disk ID is assigned, and the old disk ID is released. Therefore, you cannot roll back the system disk by using any snapshot of the released cloud disk.

Note:

After the system disk is changed, you can still use manually created snapshots of the released disk to create custom images. If you have applied an automatic snapshot policy to the old system disk and set the automatic snapshots to release when the disk is released, you must apply the policy to the new disk. Furthermore, all the automatic snapshots of the old disk are released.

Limits and recommendations

When changing the system disk, you must consider the following:

- After the system disk is changed, your instance is assigned a new cloud disk as the system disk, with a new disk ID, and the old one is released.
- You cannot replace the Cloud Type of the system disk.
- You cannot reduce the capacity of a system disk. You can only increase it. The maximum capacity of a system disk is 500 GiB.
- You cannot increase the size of the system disk that runs Windows 2003.
- If your Subscription instance has been *renewed for configuration downgrade*, you cannot modify the system disk capacity until you enter the next billing cycle.

- The IP address and the MAC address remain unchanged after the system disk is changed.
- We recommend that you create a snapshot for the system disk before you change the disk. When creating a snapshot, consider the following:
 - We recommend that you create snapshots during off-peak business hours as it may take an extended amount of time to complete. For example, it takes about 40 minutes to create a snapshot of 40 GiB. Creation of a snapshot may also reduce the I/O performance of a block storage device.
 - Make sure the system disk has enough available storage space when creating a snapshot (at least 1 GiB). Otherwise, the system may fail to start after the system disk is changed.
- To make sure you have enough quota for automatic snapshots of the new system disk, delete any unnecessary snapshots of the old system disk. For more information, see *delete snapshots or automatic snapshot policies*.

Procedure

If you want to increase the size of the system disk while keeping the operating system and environment intact, 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 instance for which to change the system disk and click its instance ID to go to the Instance Details page.
- 5. Follow these steps to create a snapshot of the system disk:
 - a. In the left-side navigation pane, click Disks.
 - b. Locate the required system disk and then, in the Actions column, click Create Snapshot.

Note:

For more information about the limits or note for creating a snapshot, see *create snapshots*.

- 6. Follow these steps to create a custom image by using the snapshot:
 - a. In the left-side navigation pane, click Instance Snapshots to check the creation status and progress. When the progress is 100% and the status is Success, go to the Actions column and click Create Custom Image.



- For more information about the limitations of creating a custom image, see create a custom mirror using a snapshot.
- The custom image is displayed in the dropdown ist of the Custom Image l on the Replace System Disk page.
- b. Go back to the Instances page and then, in the left-side navigation pane, select Snapshots and Images > Image to check the creation status and progress of the custom image.
- 7. When the progress is 100% and the status is Available, in the left-side navigation pane, click Instances.
- 8. In the Instances list, find the instance, and in the Actions column, select More > Instance Status > Stop.



For a Pay-As-You-Go VPC-Connected ECS instance, if the *No fees for stopped instances* (*VPC-Connected*) feature is enabled, in the Notes dialog box, click OK. Then, in the Stop dialog box, select Keep Stopped Instances and Continue Billing, and click OK. If you use the No Fees for Stopped Instances (VPC-Connected), you may not be able to start the instance successfully after changing the system disk.

Stop ⑦ Stop	instance	×
?	Operation will be executed on the selected 1 Instances ~ . Are you sure you want to proceed? Stopped By:	
	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.	
	ОК С	Cancel

9. When the instance is in the Stopped status, go to the Actions column and select More > Disk and Image > Replace System Disk.

10.In the pop-up dialog box, read and confirm you agree to the notice by clicking OK.

11.On the Replace System Disk page, complete the configurations as follows:

- a. Image Type: Click the Custom Image tab and select the created custom image in the drop-down list.
- b. Security enhancement:
 - System Disk: Specify a new size for the system disk according to your business needs. The maximum size is 500 GiB. The size limit for changing is determined by the image and the current size of the system disk, as displayed in the following table.

Image	Limit for capacity expansion (GiB)
Linux (excluding CoreOS) and FreeBSD	20-500
CoreOS	30-500

Image	Limit for capacity expansion (GiB)
Windows	40-500

Note:

You cannot modify the Cloud Type of the system tray.

- If a Windows image is used, set a logon password.
- If a Linux image is used and the instance is I/O optimized, you can choose to set a password or bind an SSH key pair for logon.
- c. Confirm the Instance Cost, which includes the price of the mirror and the price of the system disk. For more information, see *cloud product price*.
- d. Read and confirm you agree to the ECS Service Terms and Product Terms of Service, check the box, and then click Confirm to change.

Go back to the ECS console to check the status of the process. It may take a few minutes to process the change. After the system disk is changed, the instance starts automatically.

Follow-up operations

After the system disk is changed, you may have to perform the following:

- If your instance is running a Linux image, and a data disk was attached to the instance and set to automatically mount the file systems at the beginning, the mount information is now lost while changing the system disk. Therefore, you must write the new partition and mounting information to the / etc / fstab file on the new system disk and mount the file systems. You must not partition or format the data disk again. For more information about the commands, see Linux _ Format and mount a data disk. Follow these steps to mount the file systems:
 - 1. (Optional) Create a backup of / etc / fstab file .
 - 2. Write the new partition and mounting information to the / etc / fstab file.
 - 3. Check the new partition information in the / etc / fstab file.
 - 4. Mount the file systems.
 - 5. To view disk space and usage: run the command df h.

After mounting, you do not need to restart the instance to use the new file system directly.

• *Apply automatic snapshot policies to disks*. Note that the link between an automatic snapshot policy ID and a disk ID is broken after the system disk is changed. You must set up an automatic snapshot policy for the new system disk.

3.8.3 Windows - Resize a data disk

As your business grows, the current capacity of your data disks may not be able to meet your data storage needs. You can use the Resize Disk function to resize your data disks as necessary.

Note:

- We recommend that you manually create a snapshot to back up your data before resizing a data disk.
- You can resize a data disk when the data disk is either in the Available status or in the In Use status.
- · If a snapshot is being created for a data disk, you cannot resize the data disk.
- If you have renewed a Subscription ECS instance for configration downgrade (*renew for configuration downgrade*) during its current billing cycle, you cannot resize the attached Subscription cloud disks, including its data or system disks.
- You can resize data disks, but not file system.
- You can resize data disks, but not system disks or local disks.
- Resize the data disks that are attached to the instance only when the instance is in the Running (Running) or Stopped (Stopped) status. The changes are applied when you restart the instance in the ECS console. This action stops your instance from working and interrupts your business. Hence, proceed with caution.

This example uses a data disk of the ultra cloud disk type and an ECS instance running 64-bit Windows Server 2008 R2 Enterprise Edition to show how to resize a data disk and extend the available capacity. In this example, the current disk capacity is 20 GiB, and we resize it to 24 GiB.

To resize a data disk, follow these steps:

Step 1. Resize a data disk in the ECS console

Step 2. Log on to the instance to enable the extended storage space

Step 1. Resize a data disk in the ECS console

To resize a data 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.

Note:

If the data disk you want to resize is attached to an instance, click Instances in the left-side navigation pane, find the instance, go to the Instance Details page, and then click Disks.

- 3. Select a region.
- 4. Find the disk to be resized, and in the Actions column, select More > Resize Disk.
- 5. On the Resize Disk page, set Capacity after resizing. In this example, 24 GiB. The capacity after resizing must be larger than the current capacity.
- 6. When the cost is calculated, click Confirm to resize.

Note:

If your data disk is attached to an instance, *restart the instance* in the ECS console to make the disk resize take effect.

Once the data disk resizing completes, you can do the following:

- If the data disk is attached to an instance, log on to the instance to enable the extended storage space.
- If the data disk is not attached to an instance, attach the disk to an instance in the console first, and then proceed depending on the data disk:
 - If it is not formatted or partitioned, format and mount the data disk. For more information, see *format a data disk for Windows instances*.
 - If it is formatted and partitioned, log on to the instance to enable the extended storage space.

Step 2. Log on to the instance to enable the extended storage space

To resize a data disk within the instance, follow these steps:

- 1. Connect to a Windows instance.
- 2. On the Windows Server desktop, click the Server Manager icon



3. In the left-side navigation pane of Server Manager, select Storage > Disk Management. In the disk management area,

you can see the relationship between the new and the original

52

data disk spaces. In this example, Disk 1 is the resized data

disk.	E Server Manager		
	File Action View Help		
		3	
	 Server Manager (iZd7p5zc94o8gaZ Roles Features Diagnostics Configuration Storage Windows Server Backup Disk Management 	Disk Manageme Volume Layout (C:) Simple N Simple	nt Volume List + Graphical View Type File System Status Basic NTFS Healthy (Sys Basic NTFS Healthy (Prim
		✓ Disk 0 Basic 40.00 GB Online	(C:) 40.00 GB NTFS Healthy (System, Boot, Page Fil
		Disk 1 Basic 20.00 GB Online	New Volume (D:) 10.23 GB NTFS Healthy (Primary Partition)
		Unallocated	Primary partition

4. Right click an empty area of the New Volume of Disk 1, and select Extend Volume.



5. Follow the Extend Volume Wizard to extend the volume.

When the wizard is complete, the new data disk space is automatically merged into the original volume and the Disk 1 information showed in the Disk Manager as follows.





On Windows Server 2003, the extended storage space is added to the data disk but it is displayed as a separate volume in Disk Manager. On Windows Server 2003, one separate volume is created for each expansion and is not merged into the original volume, which does not affect the availability of the extended storage space.

You have resized a data disk successfully and the extended storage space is ready for use.

3.8.4 Linux - Resize a data disk

As your business grows, the current capacity of your data disks may not be able to meet your data storage needs. You can use the Resize Disk feature to resize your data disks as necessary.



 Resize the data disks that are attached to an instance only when the instance is in the Running or Stopped status. You must restart the instance in the ECS console to apply the changes. This action causes your instance to stop working and may cause your business to be interrupted, so please proceed with caution.

- We recommend that you manually create a snapshot to back up your data before resizing your data disk.
- You can resize a data disk when the data disk is either in the Available status or in the In Use status.
- If you have renewed a Subscription ECS instance for configration downgrade (*renew for configuration downgrade*), during its current billing cycle, you cannot resize the attached Subscription cloud disks, including its data or system disks.
- · If a snapshot is being created for a data disk, you cannot resize the data disk.
- · You can resize data disks, but not system disks or local disks.

This example uses a data disk of the ultra cloud disk type and an ECS instance running 64-bit CentOS 7.3 to describe how to resize data disk and extend the available capacity.

To resize a data disk, follow these steps:

Step 1. Increase the size of a data disk in the ECS console

Step 2. Log on to the instance to resize the file system

Step 1. Increase the size of a data disk in the ECS console

To increase the size of a data 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.



If the data disk you want to resize has been attached to an instance, in the leftside navigation pane, click Instances, find the corresponding instance, go to the instance details page, and click Disks.

- 3. Select a region.
- 4. Find the disk to be resized, and in the Actions column, select More > Resize Disk..
- 5. On the Resize Disk page, set Capacity after resizing (in this example, 30 GiB). The capacity after resizing must be larger than the current capacity.
- 6. When the cost is calculated, click Confirm to resize.



After the resizing, you can view the new disk size in the console. However, if the data disk is attached to an ECS instance, you must *restart the instance* in the ECS console to view the new disk size when you log on to the instance.

After the disk size is increased,

- · If the data disk is attached to an instance, log on to the instance to resize the file system.
- If the data disk is not attached to an instance, attach the disk to an instance in the console (*attach a cloud disk*) first, and then proceed depending on the data disk:
 - If it is a new data disk, which has not been formatted, format it. For more information, see *format a data disk for Linux instances*.
 - If it has been formatted and partitioned, log on to the instance to resize the file system.

Step 2. Log on to the instance to resize the file system

After the disk size is increased, you must log on to the instance to resize the file system.

In this example, the data disk is attached to a Linux instance running the 64-bit CentOS 7.3. The data disk before resizing has only one primary partition (/dev/vdb1, ext4 file system), the mount point of the file system is / resizetest , and after resizing is completed, the data disk still has only one primary partition.

- 1. Connect to a Linux instance by using a password.
- 2. Run the umount [file system name] command to unmount the primary partition.

umount / dev / vdb1

Note:

Run the df – h command to check whether the unmounting is successful. If you do not see the /dev/vdb1 information, unmounting is successful. The following is the sample output.

```
[ root @ iXXXXXX ~]# df
                                 h
Filesystem Size Used
dev / vda1 40G 1.5G
                                                  Mounted
                                Avail
                                         Use %
                                                             on
                                  36G
                                         4 % /
         s 487M 0 487M 0 % / dev
497M 0 497M 0 % / dev / shm
devtmpfs 487M
 tmpfs
                         496M
                                  1 % / run
          497M
                  312K
 tmpfs
                              0 % / sys / fs / cgroup
 tmpfs
          497M
                  0
                       497M
```

tmpfs 100M 0 100M 0 % / run / user / 0

3. Run the fdisk command to delete the original partition and create a new partition:

Note:

If you use the parted tool to manipulate partitions, you cannot use it in conjunction with fdisk . Otherwise, this results in an inconsistent first sector of the partition. Instructions on how to use the parted tool can be found *here*.

- a. Run the fdisk l command to list the partition details and record the final size of the partition and its first sector before resizing.
- b. Run the fdisk [device name of data disk] command to go tofdisk . In this example, the device name is / dev / vdb .
- c. Type d and press the Enter key to delete the original partition.

Note:

Deleting a partition does not cause loss of data in the data disk.

- d. Type d and press the Enter key to start creating a new partition.
- e. Type **p** and press the Enter key to create a primary partition. In this example, you are creating a single-partition data disk, so it is sufficient to create one primary partition.

Note:

If you want to create more than four partitions, create at least one extended partition, that is, type e.

- f. Type the partition number and press the Enter key. In this example, only one partition is to be created, so type 1.
- g. Type a number for the First sector: For data consistency, the number for the First sector must be identical with that of the original partition. In this example, press the Enter key to use the default value of 1.



If you find that the First sector is not identical with the recorded one, you may have used the parted tool for partitioning. In that case, stop the current fdisk operation and use *parted* to start over again.

- h. Type a number for the last sector: Because only one partition is to be created in this example, press the Enter key to use the default value.
- i. Type wq and press the Enter key to start partitioning.

```
[ root @ iXXXXXX ~]#
                     fdisk
                            / dev / vdb
         to fdisk
Welcome
                     ( util - linux
                                      2.23.2).
Changes
          will remain
                        in
                              memory
                                       only, until
                                                      you
decide
        to
            write them .
             before
Be
    careful
                      using
                               the
                                     write
                                            command .
Command ( m
              for help):
                             d
         partition
Selected
                      1
Partition
           1 is
                    deleted
Command ( m for
                    help ):
                            n
Partition type :
р
    primary
             (0)
                  primary, 0
                                 extended, 4
                                                free )
е
    extended
Select ( default
                   p):
Partition number (1 - 4,
First soctor
Using
       default response
                            р
                            default
                                        1):
       sector (2048 - 62914559, default
default value 2048
                                             2048):
First
       default value 2048
sector , + sectors or
Using
                              + size { K , M , G } ( 2048 -
Last
62914559 , default 62914559 ):
Using default value 62914559
Partition 1 of
                                                         GiB
                                         of
                                                    30
                    type
                           Linux
                                   and
                                             size
    set
is
        (m for help): wq
Command
The partition table has
                              been
                                      altered !
          ioctl () to
Calling
                       re – read
                                    partition
                                               table .
Syncing
          disks .
```

Note:

If you are using the parted tool, type p in the parted window to list the current partition details. If any partition is displayed, use rm + serial number to delete the original partition table, then run the unit s command to specify the start unit, calculated by the number of sectors, and finally run the mkpart command to create it, as shown in the following figure.

```
[root@:
                  ~]# parted /dev/xvdb
NU Parted 3.1
Jsing /dev/xvdb
Velcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
odel: Xen Virtual Block Device (xvd)
)isk /dev/xvdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
Number Start End Size File system Name Flags
(parted) unit s
(parted) mkpart primary ext3 56 5369MB
Marning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? i
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 10485760s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
lumber Start End
                          Size
                                     File system
                                                  Name
                                                            Flags
1
        56s
               10485726s 10485671s
                                     ext3
                                                   primary
```

- 4. For some operating systems, the file system may be automatically mounted to the mount point after partitioning. We recommend that you run the df h command to check the file system space and usage. Run the umount [file system name] to unmount the file system again.
- 5. Check the file system and resize the file system.

```
e2fsck - f / dev / vdb1 # check the file system
resize2fs / dev / vdb1 # resize the file system
```

Note:

- Running the e2fsck command is time-consuming because the system needs to check and revise the file system metadata during that process, so be patient.
- Properly running the e2fsck command and the resize2fs command does not cause data loss.

The following is the sample output.

```
[ root @ iXXXXXX ~]#
                    e2fsck - f / dev / vdb1
e2fsck
         1 . 42 . 9 ( 28 - Dec - 2013 )
           Checking
                      inodes ,
Pass
       1:
                               blocks ,
                                         and
                                              sizes
Pass
       2 : Checking
                      directory
                                 structure
Pass
       3:
           Checking
                      directory
                                 connectivi
                                           ty
Pass
      4 :
           Checking
                      reference
                                 counts
Pass 5 : Checking
                                       informatio n
                      group summary
```

```
/ dev / vdb1 : 11 / 1835008
                              files ( 0 . 0 % non - contiguous
    159218 / 7339776
),
                       blocks
[ root @ iXXXXXX ~]#
                      resize2fs
                                 / dev / vdb1
            1.42.9 (28 - Dec - 2013)
 resize2fs
                              on / dev / vdb1
                                                               (4k
           the filesystem
                                                      7864064
Resizing
                                                 to
   blocks .
                   on / dev / vdb1
      filesystem
                                      is
                                                 7864064
                                                           blocks
The
                                           now
long .
```

6. Mount the resized file system to the original mount point (in this example, /

```
resizetest ).
```

mount / dev / vdb1 / resizetest

7. Run the df - h command to check file system space and usage. If the correct information about the resized file system is displayed, the mounting is successful and the resized file system is ready for use.



After the mounting is completed, you can use the resized file system without restarting the instance.

The following is the sample output.

```
[ root @ iXXXXXX ~]#
                      df
                           – h
                    Used
                                    Use %
Filesystem
             Size
                            Avail
                                           Mounted
                                                     on
/ dev / vda1
              40G
                    1.5G
                                    4 % /
                              36G
devtmpfs
                  0 487M
          487M
                              0 % / dev
tmpfs
                         0 % / dev / shm
        497M
              0
                   497M
                             1 % / run
tmpfs
        497M
               312K 496M
                   497M
tmpfs
        497M
               0
                           0 % / sys / fs / cgroup
                           0 % / run / user / 0
tmpfs
        100M
               0
                   100M
                    44M
                           28G
/ dev / vdb1
              30G
                                 1 % / resizetest
```

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.



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 in	nstance		\times
?	Operation will be ex want to proceed? Stopped By:	 ecuted on the selected 1 Instances ~ . Are you sure you Stop Force Stop 	
	Stopped By:	Keep Stopped Instances and Continue Billing	
	Stopping Subscri date. If you need to st reinitialization, ir recommend that avoid startup fail	ption based instances does not change their expiration op an instance for system disk replacement, disk istance upgrade, or private IP address modification, we you select Keep Stopped Instances and Continue Billing to ure.	
		OK Car	ncel

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.



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.



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 ins	peration will be ex vant to proceed? Stopped By:	 ecuted on the selected 1 Instances ~ . Are you sure you 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.		
		ОК Са	ncel

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.

System Disk.		
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 (/).	
onfirm Password:		
Security Enhancement:	🗹 Activate 🕜	
Instance Startup Policy:	✓ Start Instance after Resetting Disk	
Note: 1. After reir 2. You must the disk.	nitialization, the selected system disk will be restored to its image. t reset the password for logging on to the instance when reinitializing	

• 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.
Security:	💽 Set SSH Key 🔘 Set Password
SSH Key Pair:	Search by SSH key pair name 🔻
Security Enhancement:	Activate
stance Startup Policy:	✓ Start Instance after Resetting Disk
te: 1. After rein 2. You must the disk.	itialization, the selected system disk will be restored to its image. reset the password for logging on to the instance when reinitializing

- 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.

Stop ⑦ Stop in	nstance		\times
?	Operation will be ex want to proceed? Stopped By:	 ecuted on the selected 1 Instances ~ . Are you sure you Stop Force Stop 	
	Stopped By:	 Keep Stopped Instances and Continue Billing 	
	Stopping Subscri date. If you need to st reinitialization, in recommend that avoid startup fail	ption based instances does not change their expiration op an instance for system disk replacement, disk istance upgrade, or private IP address modification, we you select Keep Stopped Instances and Continue Billing to ure.	
		OK Ca	ncel

- 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

RelnitDisk

3.11 Change the operating system

3.11.1 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 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 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 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 authentica tion.

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 dropdown 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, select 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, 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 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



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.

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.
- Rf the OS before and after disk replacement 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 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

3.11.2 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 authentica tion.

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.

Stop ⑦ Stop ins	stance		\times
?	Operation will be ex want to proceed? Stopped By:	 ecuted on the selected 1 Instances ~ . Are you sure you Stop Force Stop 	
	Stopped By:	 Keep Stopped Instances and Continue Billing 	
	Stopping Subscri date. If you need to sto reinitialization, in recommend that	ption based instances does not change their expiration op an instance for system disk replacement, disk stance upgrade, or private IP address modification, we you select Keep Stopped Instances and Continue Billing to	
	avoid startup fail	ure.	
		OK Ca	incel

- 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.



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	nt ⑦		
System Disk:			
Ultra Cloud Disk 🗸	40 GB 212	0 IOPS The default system	disk device name : /dev/xvda
To learn how to select SSD c	loud 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			J
commin 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	ent 🕐		
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, a	and basic cloud disks, Learn	n More >
Security:			
Key Pair	Password Set Now]	
A key pair includes a public	key and a private key. Current	J ly only I/O-optimized insta	nces support the use of key p
using a key pair, you canno	t log on with a user name and	password.	
Key Pair:			
Select the Key Pair	~ C		
Also, you can go to the cor	sole 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

3.12 Convert billing methods of cloud disks

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

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

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

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

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

3.13 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 configurat ions 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 umount [partition].For example, umount / dev / vdb1 .
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 umount [partition].For example, umount / dev / vdb1 .
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.



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 Local disks

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. Local disks provide local storage and access for instances , and feature low latency, high random IOPS, high throughput, and cost-effective performance.

Because a local disk is attached to a single physical server, the data reliability depends on the reliability of the physical server, which may create single points of failure in your architecture. We recommend that you implement data redundancy at the application layer to guarantee data availability.

\rm Marning:

Using a local disk for data storage carries the risk of data loss (for example, if the host machine is down). Therefore, we recommend you never use a local disk to store any data that requires long-term persistence. If no data reliability architecture is available for your application, we strongly recommend that you build your ECS with *cloud disks or Shared Block Storage devices*.

This document details information about local disks and instances that support local disks. If you are using a previous generation local SSD disk, see *local SSD disks in the previous generation disks*.

Disk types

Currently, Alibaba Cloud provides two types of local disks:

- Local NVMe SSD: This disk is used together with instances of the following type families: i2, i1, and gn5. The instance type families i1 and i2 apply to the following scenarios:
 - Online games, e-businesses, livestreaming, and other industries that provide online businesses and have low latency and high I/O performance requirements on block level storage for I/O-intensive applications.
 - Business scenarios that have high requirements on the storage I/O performance and availability of the application layer, such as NoSQL non-relational databases , MPP data warehouses, and distributed file systems.

 Local SATA HDD: This disk is used together with instances of the d1ne and d1 type families. It is applicable to businesses that require big data computing and storage analysis for massive data storage and offline computing business scenarios. It fully meets the needs of distributed computing business models (such as those built on the Hadoop framework) across instance storage performance, capacity, and intranet bandwidth.

Performance of local NVMe SSD

The following table lists the performance of local NVMe SSD of an i1 ECS instance.

Parameters	Local NVMe SSD
Maximum capacity	Single disk: 1,456 GiB Total: 2,912 GiB
Maximum IOPS	Single disk: 240,000 Total: 480,000
Maximum throughput	Read throughput per disk: 2 GBps Total read throughput: 4 GBps Write throughput per disk: 1.2 GBps Total write throughput: 2.4 GBps
Single-disk performance *	 Write performance: Single-disk IOPS: IOPS = min{165 * capacity, 240,000} Single disk throughput: Throughput = min{0.85 * capacity, 1,200} MBps Read performance: Single disk IOPS: IOPS = min{165 * capacity, 240,000} Single disk throughput: Throughput = min{1.4 * capacity, 2,000} MBps
Access latency	Microsecond-level

* Single disk performance calculations are as follows:

- Write IOPS for a single local NVMe SSD: 165 IOPS for each GiB, up to 240,000 IOPS.
- Write throughput for a single local NVMe SSD: 0.85 MBps for each GiB, up to 1,200 Mbit/s.

Performance of local SATA HDD

The following table lists the performance of local SATA HDD of a d1ne or d1 ECS instance.

Parameters	Local SATA HDD
Maximum capacity	Single disk: 5,500 GiB Total capacity per instance: 154,000 GiB
Maximum throughput	Single disk: 190 MBps Total throughput per instance: 5,320 MBps
Access latency	Millisecond-level

Billing

Local disks are charged according to the instances to which they are attached. For

more information about instance billing methods, see Subscription and Pay-As-You-Go.

Lifecycle

A local disk has the same lifecycle as the instance that it is attached to. This means that:

- You can create a local disk only when creating an instance that has local storage
 The capacity of a local disk is determined by the ECS instance type. You cannot increase or decrease it.
- $\cdot\,$ When the instance is released, the local disk is released with it.

Instance operations

The following table details how operations on an instance that has local storage affect the state of the data on the local disk.

Operation	State of the data on the local disk	Result
Restart within the operating system/restart or force restart in the ECS console	Retained	Both the storage volumes and data on the local disk are retained.

Operation	State of the data on the local disk	Result
Shutdown within the operating system/stop or force stop in the ECS console	Retained	Both the storage volumes and data on the local disk are retained.
Release in the ECS console	Erased	The storage volumes on the local disk are erased and the data on it is not retained.
Downtime migration	Erased	The storage volumes on the local disk are erased and the data on it is not retained.
Out-of-service (before the computing resources of an instance is released)	Retained	Both the storage volumes and data on the local disk are retained.
Out-of-service (after the computing resources of an instance is released)	Erased	The storage volumes on the local disk are erased and the data on it is not retained.

Related operations

If your ECS instance is attached with local disks, you must connect to the instance to *format the disk*. Unlike cloud disks, you cannot perform the following operations on local disks:

- Independently create an empty local disk or create a local disk from a snapshot.
- Attach a local disk in the ECS console.
- Detach and release a local disk.
- Increase the size of a local disk.
- Re-initialize a local disk.
- Create a snapshot for a local disk and use the snapshot to roll back the local disk.