Alibaba Cloud
Elastic Compute Service

Block Storage

Legal disclaimer

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

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed due to product version upgrades , adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults " and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity , applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use

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

5. By law, all the content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of the Alibaba Cloud website, product programs, or content shall be used, modified , reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates . The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates).

6. Please contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
⛔	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	⛔ Danger: Resetting will result in the loss of user configuration data.
⚠️	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	⚠️ Warning: Restarting will cause business interruption. About 10 minutes are required to restore business.
📋	This indicates warning 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.**	`swich` *{stand | slave}*

Contents

1 What is block storage?

Overview

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

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

· Cloud Disk, which is a block-level data storage product provided by Alibaba Cloud for ECS, uses a multiple distributed system, and features low latency, high performance, persistence, high reliability, and more. Cloud disks can be created, resized, and released at any time.

· Shared block storage is a block-level data storage device that supports simultaneous read and write access to multiple ECS instances. Similar to the cloud disk, shared block storage uses a multiple distributed system. It supports simultaneous access to multiple instances, and features low latency, high performance, and high reliability. Shared Block Storage applies to shared access scenarios for block storage devices under a shared everything architecture.

· Local disks are the disks attached to the physical servers (host machines) on which ECS instances are hosted. They are designed for business scenarios requiring high storage I/O performance and massive storage cost performance. Local disks provide local storage and access for instances, and features low latency, high random IOPS, high throughput, and cost-effective performance.

For more information about the performance of block-level storage products, see Block storage performance.

Block storage, OSS and NAS

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

the following three types of data storage products:

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

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

· OSS: A huge storage space designed for storing massive amounts of 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 Block storage performance

This topic describes the performance indexes, performance, and performance test methods of block storage devices such as cloud disks, Shared Block Storage, and local disks.

Performance indexes

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

· IOPS

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

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

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

· Throughput

Throughput measures the size of data 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 ESSD disks, 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

Block storage capacity is measured in binary units, such as Kibibyte (KiB), Mebibyte (MiB), Gibibyte (GiB), or Tebibyte (TiB). This section describes the performance of different block storage devices.

> **Note:**
> 1 KiB equals 1,024 bytes. 1MiB equals 1,024 KiB. 1GiB equals 1,024 MiB. 1TiB equals 1,024 GiB.

· Cloud disks

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

Measurement	ESSD cloud disk			SSD cloud disk	Ultra cloud disk	Basic cloud disk***
Performance level (PL)	PL3	PL2	PL1	N/A	N/A	N/A
Capacity of a single disk	32,768 GiB			32,768 GiB	32,768 GiB	2,000 GiB
Max. IOPS	1,000,000	100,000	50,000	25,000*	5,000	Several hundreds
Max. throughput	4,000 Mbit/s	750 Mbit/s	350 Mbit/s	300 Mbit/s*	140 Mbit/s	30−40 Mbit/s

Measurement	ESSD cloud disk			SSD cloud disk	Ultra cloud disk	Basic cloud disk[***]
Formulas to calculate performance of a single disk[**]	IOPS = min{1,800 + 50 × capacity, 1,000,000}	IOPS = min{1,800 + 50 × capacity, 100,000}	IOPS = min{1,800 + 50 × capacity, 50,000}	IOPS = min{1,800 + 30 × capacity, 25,000}	IOPS = min{1,800 + 8 × capacity, 5,000}	N/A
	Throughput = min {120 + 0.5 × capacity, 4,000} Mbit/s	Throughput = min {120 + 0.5 × capacity, 750} Mbit/s	Throughput = min {120 + 0.5 × capacity, 3500} Mbit/s	Throughput = min{120 + 0.5 × capacity, 300} Mbit/s	Throughput = min{100 + 0.15 × capacity, 140} Mbit/s	N/A
Data reliability	99.9999999%			99.9999999%	99.9999999%	99.9999999%
API name	cloud_essd			cloud_ssd	cloud_efficiency	cloud

Measurement	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 - ElasticSearch distributed logs : Elasticsearch, Logstash, and Kibana (ELK) log analysis	- Large and medium -sized relational databases , such as MySQL , SQL Server, PostgreSQL , and Oracle - Large and medium -sized developmen t or testing applicatio ns that require high data reliability	- Developmen t and testing applicatio ns - System disk	- 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 SSD cloud disks varies according to the size of the data blocks . Smaller data blocks result in lower throughput and higher IOPS, as described in the following table. An SSD cloud disk can achieve a suitable IOPS performance only when it is attached to an I/O-optimized instance.

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

Data block size	Maximum IOPS	Throughput
64 KiB	About 4,800	About 300 Mbit/s

**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 Mbit/s. It increases by 0.5 Mbit/s
  per GiB of storage. The maximum throughput is 300 Mbit/s.

The random write latency varies according to disk categories. For the test
commands, see Performance tests.

- ESSD cloud disks: 0.2 ms
- SSD cloud disks: 0.5−2 ms
- Ultra cloud disks: 1−3 ms
- Basic cloud disks: 5−10 ms

***Basic cloud disks are last-generation cloud disks that are scheduled to be
discontinued. We recommend that you use one of the newer cloud disk products
that have better performance, such as Ultra cloud disks or SSD cloud disks.

· Shared Block Storage

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

Measurement	SSD Shared Block Storage	Ultra Shared Block Storage
Maximum capacity	- Singe disk: 32,768 GiB - Single instance: 128 TiB	- Singe disk: 32,768 GiB - Single instance: 128 TiB
Maximum random read/ write IOPS*	30,000	5,000
Maximum sequential read/write throughput*	512 Mbit/s	160 Mbit/s
Formulas to calculate performance of a single disk**	IOPS = min{1,600 + 40 × capacity, 30,000}	IOPS = min{1,000 + 6 × capacity, 5,000}
	Throughput = min{100 + 0.5 × capacity, 512} Mbit/s	Throughput = min{50 + 0.15 × capacity, 160} Mbit/s

Measurement	SSD Shared Block Storage	Ultra Shared Block Storage
Scenarios	- Oracle RAC - 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 device 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 Mbit/s. It increases by 0.5 Mbit/s per GiB of storage. The maximum throughput is 512 Mbit/s.

The latency varies according to the Shared Block Storage categories as follows:

- 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.

Performance tests

The FIO tool is recommended for testing block storage performance for both Linux and Windows instances.

Note:

You can also use other tools to test block storage performance, but the resulting benchmark performance may be different. For example, tools like dd, sysbench, and iometer may be affected by test parameter configuration and file systems, and may not reflect the actual disk performance. The performance results in this topic are from a Linux instance tested with FIO, and are used as the performance reference of

Shared Block Storage products. Before you test the disk, verify that the disk is 4 KiB aligned.

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

- Test random write IOPS:

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

- Test random read IOPS:

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

- Test sequential write throughput:

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

- Test sequential read throughput:

```
fio  – direct = 1  – iodepth = 64  – rw = read  – ioengine = libaio
 – bs = 1024k  – size = 1G  – numjobs = 1  – runtime = 1000  –
group_repo  rting  – filename = iotest  – name = Read_PPS_T  esting
```

- Test random write latency:

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

- Test random read latency:

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

```
= 4k  – size = 1G  – numjobs = 1  – group_repo  rting  – filename =
iotest  – name = Rand_Write  _Latency_T  esting
```

The command for testing random write IOPS is used as an example to describe the meaning of the parameters of a FIO command, as described in the following table.

Parameter	Description
-direct=1	Ignore I/O buffer when testing. Data is written directly.
-iodepth=128	Indicates that when you use AIO, up to 128 I/O requests can be made at the same time.
-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)

Parameter	Description
-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 I/O In synchronous I/O, a thread sends out only one I/O request at a time and waits until the I/O request has completed. In this case, the iodepth is always less than 1 for a single job, which means low processing efficiency. You can increase the iodepth by using multiple concurrent jobs. Usually 16–32 concurrent jobs can fill up the iodepth. · Asynchronous I/O The asynchronous method uses libaio to submit a batch of I/O requests each time, thus reducing interaction times and making interactions more efficient.
-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, such as 4k in this example command. When throughput is tested, we recommend that you set `bs` to a large value, such as 1,024k in this example command.

Parameter	Description
-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 of the testing results . Group_reporting means the statistics of each job are summed up, instead of showing all statistics of each job.
-filename=iotest	The 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 disk overview

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

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

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

Cloud disk types

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

· By performance

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

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

· **By function**

Type	Description
System disk	A system disk has the same life cycle as the ECS instance to which it is mounted, and is created and released along with the ECS instance. Shared access to the system disk is not allowed. The available size of a system disk varies according to the operating system as follows: - Linux (excluding CoreOS) and FreeBSD: 20–500 GiB - CoreOS: 30–500 GiB - Windows: 40–500 GiB
Data disk	Data disks can be created separately or at the same time as ECS instances. - Data disks created along with ECS instances have the same life cycle as the corresponding instance, and are released along with the instance. - Data disks created separately can be released separately or at the same time as the corresponding ECS instance. Shared access to a data disk is not allowed. The performance of data disks depends on the cloud disk type. For more information, see Block storage performance. 📋 Note: When used as data disks, up to 16 cloud disks can be attached to one ECS instance.

Performance

For the performance of each type of cloud disk, see Block storage performance.

Billing method

· Cloud disks created along with Subscription instances or created separately for Subscription instances are billed according to the selected billing cycle and are charged before you use the service. For more information, see Subscription.

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

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

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

Note:

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

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

3.2 ESSD cloud disk

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

Scenarios

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

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

Specifications

The API value of ESSD cloud disks is `cloud_essd` . ESSD cloud disks are available in three performance levels (PLs). For more information, see How can I test the performance of an ESSD cloud disk?

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

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

Billing

ESSD cloud disks support Pay-As-You-Go and Subscription billing methods. For more information, see Create a Pay-As-You-Go cloud disk and Create a Subscription cloud disk.

For more information, see ECS products.

PL selection for ESSD cloud disks

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

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

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

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

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

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

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

Disk capacity and performance

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

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

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

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

If you require higher performance, resize the ESSD cloud disk or change the performance level of the ESSD cloud disk.

Related operations

- Create a Pay-As-You-Go cloud disk.
- Create a Subscription cloud disk.
- Create a cloud disk by using a snapshot.
- Attach a cloud disk.
- Convert the billing methods of cloud disks.
- Resize a cloud disk.
- Modify the performance level of an ESSD cloud disk.

3.3 Triplicate technology

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

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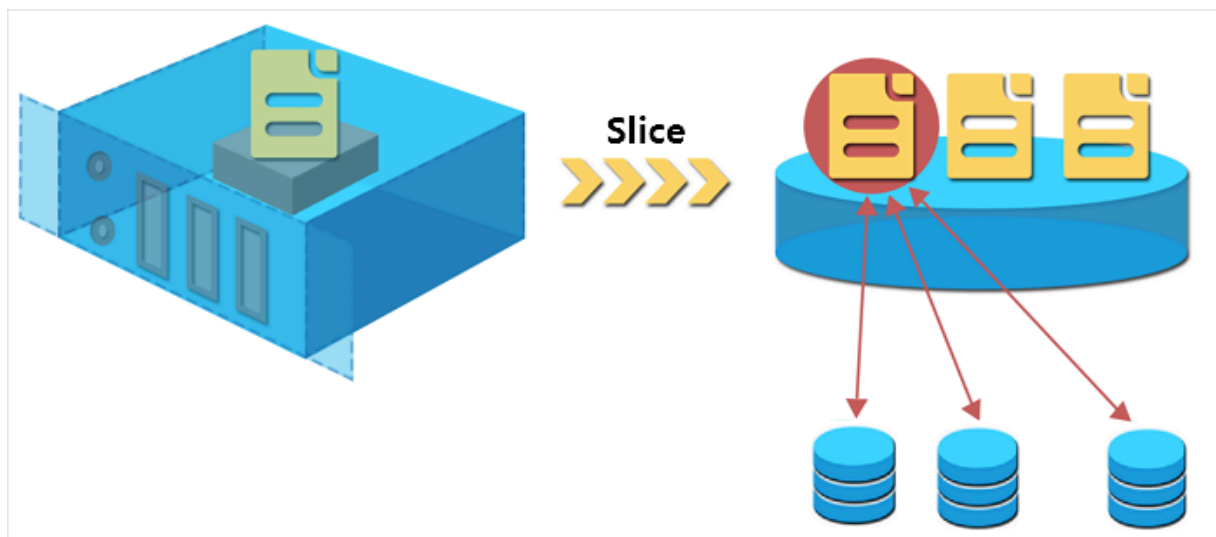
How triplicate technology works

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

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

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

Data protection

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

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

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

3.4 ECS disk encryption

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

What is ECS disk encryption?

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

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

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

· Stored directly on the ECS disk.

· Transmitted between the ECS disk and the instance. However, data in the instance operating system is not encrypted.

· Created from the encrypted ECS disk, such as snapshots. These snapshots are called encrypted snapshots.

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

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

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

ECS disk encryption dependencies

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

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

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

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

Key management for ECS disk encryption

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

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

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

The key management infrastructure of Alibaba Cloud conforms to the 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 ReInitDisk.

Create an encrypted ECS disk

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

- Create a cloud disk as a data disk when creating an ECS instance or :

  - Check Encrypted to create a encrypted blank cloud disk.

  - Select an encrypted screenshot to create a cloud disk.

- When using APIs or the CLI:

  - Set the parameter `DataDisk . n . Encrypted` ([CreateInstance](#)) or `Encrypted` ([CreateDisk](#)) to `true` .

  - Specify the `SnapshotId` parameter of the encrypted snapshot in CreateInstance or CreateDisk.

Convert unencrypted data to encrypted data

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

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

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

Therefore, if you must switch the existing data from status encrypted to 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.

3.5 Create a cloud disk

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

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

Limits

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

· You cannot create a separate system disk in the ECS console.

· Across all regions, the number of Pay-As-You-Go data disks that you can create cannot be more than five times the number of Pay-As-You-Go instances under your account. For more information, see Limits.

· You cannot merge multiple cloud disks by formatting them because they are independent of each other. Therefore, we recommend that you estimate the number and size of cloud disks required before you create them.

· We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes for multiple cloud disks. This is because a snapshot can only back up data of a single cloud disk. If you use LVM, data discrepancies will occur when you roll back these cloud disks.

Billing

The cloud disk created in this topic is billed as Pay-As-You-Go. For more information, see Pay-As-You-Go.

If you need to convert the billing method of a Pay-As-You-Go cloud disk to Subscription, see Convert the billing methods of cloud disks.

> **Note:**
>
> If you need to create a Subscription cloud disk, use either of the following methods:
>
> · Create a Subscription cloud disk and mount it to a Subscription instance.
> · When you create an ECS instance, set its billing method to Subscription. In this case, all the cloud disks attached to the instance use the Subscription billing method.
>
> You cannot unmount or release a Subscription cloud disk on demand. The Subscription cloud disk is released along with the corresponding instance.

Procedure

1. Log on to the ECS console.
2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
3. In the upper-right corner of the Disks page, click Create Disk.
4. Select the region and zone that are the same as the target instance.
5. Specify the disk type, capacity, disk encryption mode, and quantity.

> **Note:**
>
> · To check whether a cloud disk needs to be encrypted, see #unique_21.
> · You can also create a cloud disk by using a snapshot.
> · For information about how to select between ESSD cloud disks of different specifications, see ESSD cloud disk.

6. Select the *ECS Service Level Agreement* check box.
7. Confirm the cloud disk configuration and fee.
8. Click Preview Order.

9. In the displayed dialog box, click Create.

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

The Status of the cloud disk is displayed as Unattached.

What to do next

Attach a cloud disk.

Related APIs

· CreateDisk: Creates a Pay-As-You-Go or Subscription cloud disk.

· RunInstances or CreateInstance: Creates a data disk along with an ECS instance.

References

· Create a cloud disk by using a snapshot

· Create an instance by using the wizard

3.5.2 Create a Subscription cloud disk

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

Limits

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

· You cannot merge multiple cloud disks by formatting them because they are independent of each other. Therefore, we recommend that you estimate the number and size of cloud disks required before you create them.

· We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes for multiple cloud disks. This is because a snapshot can only back up data of a single cloud disk. If you use LVM, data discrepancies will occur when you roll back these cloud disks.

· You cannot detach a Subscription cloud disk on demand. The Subscription cloud disk is released along with the instance. If you need to release a Subscription cloud disk, convert it to a Pay-As-You-Go cloud disk, and then detach and release it.

Billing details

The cloud disk created in this topic is billed as Subscription. For more information, see Subscription.

If you need to convert a Subscription cloud disk to a Pay-As-You-Go cloud disk, see **Convert the billing methods of cloud disks**.

> 📋 **Note:**
>
> If you need to create a Pay-As-You-Go cloud disk, use either of the following methods:
>
> - **Create a Pay-As-You-Go cloud disk** on the Disks page of the ECS console.
> - When you **create an ECS instance**, set its billing method to Pay-As-You-Go. In this case, all the cloud disks attached to the instance use the Pay-As-You-Go billing method.

Procedure

To create a Subscription cloud disk, follow these steps:

1. Log on to the **ECS console**.
2. In the left-side navigation pane, choose Instances & Images > Instances.
3. On the Instances page, find the target Subscription instance. In the Actions column, choose More > Configuration Change > Add Subscription Cloud Disk.
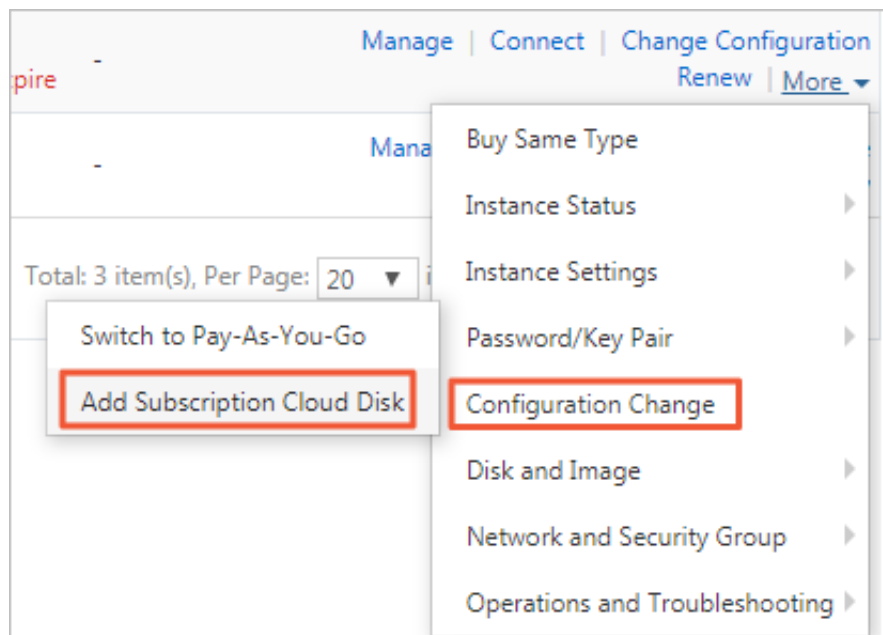
4. **In the Disk area of the displayed page, complete the following configurations:**

   · Cloud disk type: Select a cloud disk type in the drop-down list. For information about how to select between ESSD cloud disks of various specifications, see ESSD cloud disk.

   · Cloud disk capacity: Enter a cloud disk capacity in the text box. The disk capacity ranges from 20 GiB to 32,768 GiB.

   · Cloud disk encryption: If ECS disk encryption is needed, select the Disk Encryption check box.

   · Quantity: Enter the number of cloud disks you want to purchase.

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

   · Disk Name: optional. The name for the cloud disk.

   · Description: optional. The description of the cloud disk.

   · If you need to create a cloud disk from a snapshot, click Create from snapshot.

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

6. Click Preview Order.

7. In the displayed dialog box, click Create.

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

9. Click Console to return to the Instances page and click the instance for which the Subscription cloud disk has just been added.

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

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

What to do next

Format and partition the cloud disk. For more information, see Format a data disk of a Linux instance or Format a data disk for Windows instances.

Related APIs

· CreateDisk: Creates a Pay-As-You-Go or Subscription cloud disk.

· RunInstances or CreateInstance: Creates a data disk along with an ECS instance.

References

· Create a Pay-As-You-Go cloud disk

· Create a cloud disk by using a snapshot

3.5.3 Create a cloud disk by using a snapshot

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

Scenarios

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

· If you need to obtain data from a snapshot, but you do not want to roll back a cloud disk.

· If your instance encounters a system disk failure, you can use an existing system disk snapshot to create a cloud disk. Then, you can attach the cloud disk as a data disk to a healthy ECS instance so that you can continue to read the data of the system disk.

Limits

· By default, a cloud disk created by using a snapshot uses the Pay-As-You-Go billing method, and can only be used as a data disk. However, you can change the billing method of the cloud disk. For more information, see What to do next.

· When you access a cloud disk created by a snapshot the first time, the performance is reduced because it takes some time for ECS to read data from OSS and write data to the cloud disk. Therefore, we recommend that you do not use the cloud disk until it has read from all data blocks. For more information, see What is OSS?

- Across all regions, the number of Pay-As-You-Go data disks that you can create cannot be more than five times the number of Pay-As-You-Go instances under your account. For more information, see Limits.

- You cannot merge multiple cloud disks by formatting them because they are independent of each other. Therefore, we recommend that you estimate the number and size of cloud disks required before you create them.

- We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes for multiple cloud disks. This is because a snapshot can only back up data of a single cloud disk. If you use LVM, data discrepancies will occur when you roll back these cloud disks.

Prerequisites

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

Procedure

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

1. Log on to the ECS console.
2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
3. In the upper-right corner of the Disks page, click Create Disk.
4. Select the target region and zone.

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

5. Complete the following settings:

   a. Select a disk type: Ultra Disk, ESSD Cloud Disk, or SSD Cloud Disk.

   b. Click Create from snapshot, and then select a snapshot.

   c. Set the capacity of the cloud disk. Note that it must be greater than 20 GiB and less than 32768 GiB.

      > 📋 Note:

- If you do not set the capacity of the cloud disk, the system automatically sets a capacity for the new cloud disk according to the cloud disk corresponding to the snapshot.
- If you set a capacity greater than the snapshot size, you must partition the cloud disk before you can use the full capacity.
- If the snapshot size is less than 2048 GiB and the cloud disk to be created is greater than 2048 GiB, you must ensure that the cloud disk corresponding to the snapshot uses GPT partitions (Globally Unique Identifier Partition Table). Otherwise, we recommend that you set the cloud disk size to less than 2048 GiB to avoid data loss that may occur in partitioning. For more information, see Partition and format data disk more than 2 TiB.

   d. Set the Quantity of cloud disks that you want to create.

   e. Read and confirm that you agree with the ECS Service Level Agreement by selecting the check box.

6. Confirm your settings and the estimated cost displayed.

7. Click Preview, and then click Create in the displayed dialog box.

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

What to do next

After you create a cloud disk, you can:

- Attach the cloud disk.
- Change the billing method of the cloud disk from Pay-As-You-Go to Subscription.

   - If the cloud disk is attached to a Subscription instance, see Upgrade configurations of Subscription instances.

   - If the cloud disk is attached to a Pay-As-You-Go instance, see Switch from Pay-As-You-Go to Subscription billing.

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

3.6 Attach a cloud disk

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

Limits

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

- You can attach a cloud disk as a data disk only. You cannot attach a cloud disk as a system disk.
- To attach a cloud disk to an ECS instance, the instance must meet the following requirements:

  - The instance must be in the Running or Stopped state, but not in the Locked state.
  - The instance must not have any overdue payments.

- The disk to be attached must be in the Unmounted state.
- The cloud disk and the ECS instance must be in the same region and the same zone.
- Up to 16 cloud disks can be attached to an ECS instance to work as data disks. One cloud disk cannot be attached to multiple instances simultaneously.
- If a cloud disk is created independently on the Disks page in the ECS console, it can be attached to any ECS instance in the same region and the same zone, regardless of the billing method of the instance.

Prerequisites

You have created an ECS instance and a cloud disk in the same region and zone. For more information, see Create an instance by using the wizard and Create a Pay-As-You-Go cloud disk.

Attach a cloud disk on the Instances page

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

1. Log on to the ECS console.
2. In the left-side navigation pane, choose Instances & Images > Instances.
3. Find the target ECS instance and click its ID to go to the Instance Details page.

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

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

    · Target Disk: Select a cloud disk in the Unmounted state in the same zone.

    · Release Disk with Instance: If you select this option, the disk is released when you release the corresponding instance.

    · Delete Automatic Snapshots While Releasing Disk: If you select this option, all the automatic snapshots of the selected disk are deleted when you release the disk. However, all the manual snapshots are retained. To keep a complete data backup history, we recommend that you do not select this option.

    Click OK, and then Mount.

6. Refresh the Disks page.

    When the cloud disk is in the In Use state, 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 the different follow-up operations.

Disk content	Operating system of the ECS instance	Follow-up operations
A new empty cloud disk	Linux	Format a data disk of a Linux instance. If the cloud disk is larger than 2 TiB, see Partition and format data disk more than 2 TiB.
	Windows	Format a data disk for Windows instances. If the cloud disk is larger than 2 TiB, see Partition and format data disk more than 2 TiB.

Disk content	Operating system of the ECS instance	Follow-up operations
A cloud disk created from a snapshot	Linux	Connect to the Linux instance and run the `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, choose Storage & Snapshots > Disks.

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

4. 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.

5. Refresh the Disks page.

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

6. 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

3.7 Format a data disk

3.7.1 Format a data disk for Windows instances

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

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

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

Prerequisites

For a separately purchased data disk, you must attach the data disk to an instance before partitioning and formatting.

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

Procedure

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

1. Connect to an instance.

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

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

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

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

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

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

   · MBR is still the most common partitioning method. However, this method only supports data disks that no greater than 2 TB and can divide a disk into up to four primary partitions. If you want to divide a disk into more than four partitions, you need to take a primary partition as an extended partition and create logical partitions within it.

   · GPT is a new partitioning method, and cannot be recognized by earlier versions of Windows. The size of GPT-partitioned data disk is determined by the operating system and the file system. In the Windows operating system, GPT supports up to 128 primary partitions.

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

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

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

   a. Click Next.

   b. Specify Volume Size: Specify the size of the simple volume to create. If you need only one primary partition, use the default value, and then click Next.

   c. Assign Drive Letter or Path: Select a drive letter (in this example, F). Click Next.

   d. Format Partition: select format settings (including the file system, allocation unit size, and volume label), and confirm whether to enable Quick Formatting and File and Folder Compression. Use the default values, then click Next.

   e. Create a new simple volume. When the wizard shows the information below, a new simple volume is created. Click Finish to close the New Simple Volume Wizard.

After the partition formatting is completed, the status of Disk 2 in Disk Managementis as shownin the following figure.
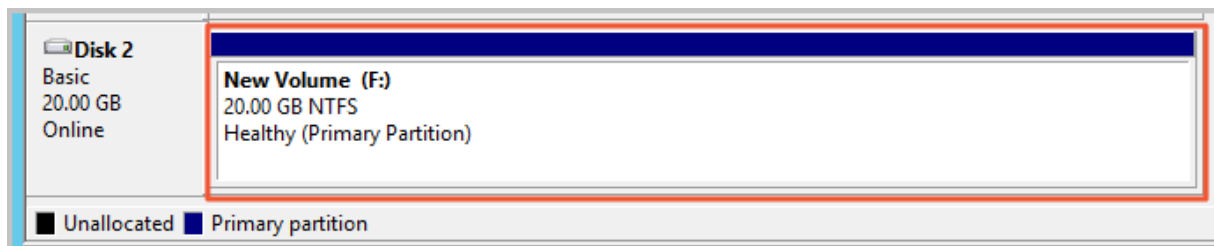
InThis PC, you can view a new drive namedNew Volume (F:). The data disk is now ready to use.

3.7.2 Format a data disk of a Linux instance

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

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

> 📋 **Note:**
>
> The following procedure applies only to data disks less than 2 TiB. If your data disk is greater than 2 TiB, see Partition and format data disk greater than 2 TiB.

Prerequisites

· The ECS instance is attached with a data disk that was created separately. You do not need to perform this operation for data disks created along with ECS instances.

· **The device name of the data disk is obtained.**

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

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

Procedure

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

1. **[Connect to the instance](#) to which the data disk is attached.**

2. **Run the** `fdisk - l` **command to view the data disks of the instance.**

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

3. **Create a single-partition data disk by running the following commands in sequence :**

   a. **Run the** `fdisk - u / dev / vdb` **command to partition the data disk.**

   b. **Enter** `p` **and press Enter to view the partitions of the data disk. In this example, the data disk is not partitioned.**

   c. **Enter** `n` **and press Enter to create a new partition.**

   d. **Enter** `p` **and press Enter to select the primary partition.**

   > **Note:**

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

e. Enter the partition number and press Enter. In this example, enter `1` .

f. Enter a number for the first available sector, or press Enter to use the default value of `2048` .

g. Press Enter to use the default number for the last sector.

h. Enter `p` and press Enter to view the planned partitions of the data disk.

i. Enter `w` and press Enter to start partitioning and exit after partitioning.

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

Command (m for help): p
Disk /dev/vdb: 21.5 GB, 21474836480 bytes, 41943040 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x3e60020e
Device Boot Start End Blocks Id System

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

Command (m for help): p

Disk /dev/vdb: 21.5 GB, 21474836480 bytes, 41943040 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x3e60020e
Device Boot Start End Blocks Id System
/dev/vdb1 2048 41943039 20970496 83 Linux

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

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

```
Syncing disks.
```

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

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

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

    Disk  / dev / vdb :  21 . 5   GB ,  2147483648  0   bytes ,
    41943040   sectors
    Units  =  sectors   of  1  *  512  =  512   bytes
    Sector   size  ( logical / physical ):  512   bytes  /  512
    bytes
    I / O   size  ( minimum / optimal ):  512   bytes  /  512   bytes
    Disk   label   type :  dos
    Disk   identifier :  0x3e60020e

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

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

   Note:

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

   ```
   [ root @ ecshost ~ ]#  mkfs . ext4  / dev / vdb1
    mke2fs  1 . 42 . 9  ( 28 - Dec - 2013 )
    Filesystem   label =
    OS   type :  Linux
    Block   size = 4096  ( log = 2 )
    Fragment   size = 4096  ( log = 2 )
    Stride = 0   blocks ,  Stripe   width = 0   blocks
    1310720   inodes ,  5242624   blocks
    262131   blocks  ( 5 . 00 %)  reserved   for   the   super   user
    First   data   block = 0
    Maximum   filesystem   blocks = 2153775104
    160   block   groups
    32768   blocks   per   group ,  32768   fragments   per   group
    8192   inodes   per   group
    Superblock   backups   stored   on   blocks :
    32768 ,  98304 ,  163840 ,  229376 ,  294912 ,  819200 ,  884736 ,
    1605632 ,  2654208 ,
    4096000

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

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

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

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

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

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

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

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

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

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

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

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

3.7.3 Partition and format data disk more than 2 TiB

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

> **Note:**
>
> If you want to partition and format a data disk less than 2 TiB, please see Format a data disk for Linux instances and Format a data disk for Windows instances.

Note

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

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

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

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

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

- 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 theUnallocatedareaof Disk 4, and then
   select New Simple Volume to create a 4 TiB volume in the NTFS format.

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

   a. Click Next.

   b. Choose a volume size: designate size of simple volume. If you want to create a
      master area only, use the default value. Click Next. You can also partition Disk 4
      into several partitions.

      > 📋 Note:
      >
      > The maximum NTFS volume is, theoretically, the maximum volume of NTFS
      > containing 2_{32} -1 clusters. However, in Windows XP Pro, the maximum volume
      > of NTFS is 2   -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 2_{32} -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   GB ,  3221225472  000   bytes ,
6291456000   sectors
Units  =  sectors   of   1  *  512  =  512   bytes
Sector   size  ( logical / physical ):  512   bytes  /  512
bytes
I / O   size  ( minimum / optimal ):  512   bytes  /  512   bytes
```

2. Run `parted  / dev / vdd` to start partitioning:

   a. Run `mklabel   gpt` , to convert partitioning format from MBR to GPT.

   b. Run `mkpart   primary   ext4  < 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   is   not   properly
    aligned   for   best   performanc  e .
    Ignore / Cancel ?  ignore
   ( parted )  print
    Model :  Virtio   Block   Device ( virtblk )
    Disk  / dev / vdd :  3221   GB
    Sector   size  ( logical / physical ):  512B / 512B
    Partition   Table :  gpt
    Disk   Flags :
    Number   Start   End   Size   File   system   Name   Flags
    1   17 . 4kB   3221GB   3221GB   primary
   ```

   d. Run `quit` to exit parted .

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

4. Run the following commands to create an ext4 file system, and make /dev/vdd1 partition use ext4.

   ```
   mke2fs  - O   64bit , has_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 @ izXXXXz ~]#  df  - h
 Filesystem   Size   Used   Avail   Use %  Mounted   on
/ dev / vda1   40G    6 . 4G   31G    18 % /
 devtmpfs   487M   0   487M   0 % / dev
 tmpfs   497M   0   497M   0 % / dev / shm
 tmpfs   497M   364K   496M   1 % / run
 tmpfs   497M   0   497M   0 % / sys / fs / cgroup
 tmpfs   100M   0   100M   0 % / run / user / 0
/ dev / vdd1   2 . 9T   89M   2 . 8T   1 % / test
```

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

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

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

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

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

Appendix 1: Update e2fsprogs

If the disk capacity is 16 TiB, you must use e2fsprogs of version 1.42 or later to format its partitions to ext4 file system. If e2fsprogs version is too low (for example, e2fsprogs 1.41.11), the following error occurs.

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

To install e2fsprogs of later version, such as 1.42.8 in this example, follow these steps:

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

   e2fsprogs.

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

3.8 Resize cloud disks

3.8.1 Disk resizing overview

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

Scenarios

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

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

Procedures for disk resizing

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

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

Resizing procedure	1. Use the ECS console or call the API action ResizeDisk to resize the cloud disk.
	
	 2. Use the ECS console or call the API action AttachDisk to attach the cloud disk.
	
	 3. Partition and format the cloud disk:
	
	  · Partition and format a Windows data disk less than 2 TiB.
	  · Partition and format a Linux data disk less than 2 TiB.
	  · Partition and format a cloud disk greater than 2 TiB.
		1. Use the ECS console or call the API action ResizeDisk to resize the cloud disk.
		
		 2. Use the ECS console or call the API action RebootInstance to restart the instance.
		
		 3. Partition and format the cloud disk:
		
		  · Partition and format a Windows data disk less than 2 TiB.
		  · Partition and format a Linux data disk less than 2 TiB.
		  · Partition and format a cloud disk greater than 2 TiB.
			1. Use the ECS console or call the API action ResizeDisk to resize the cloud disk.
			
			 2. Use the ECS console or call the API action RebootInstance to restart the instance.
			
			 3. Extend the partitions of the attached system disk or data disk:
			
			  · Extend a Windows file system.
			  · Extend a Linux file system.
			  · Extend the file system of a Linux data disk.

Limits of resizing a system disk

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

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

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

Upper limit of data disk resizing

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

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

3.8.2 Resize cloud disks online

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

What is online resizing?

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

· Online mode: You do not need to restart instances and instances must be in the
  running (Running) state.

· Offline mode: You must restart instances and instances can be in the running (
   Running) or stopped (Stopped) state.

Notes

System limits

· For more information about thresholds of extended system disks and data disks,
  see Overview.

- The operating system must meet the following conditions when you resize system disks. Otherwise, you must select the offline mode.

  - Windows operating system: cannot be Windows Server 2003.

  - Linux operating system: The kernel version must be later than 3.6 when you run the `uname  - a` command.

- When you resize cloud disks, only storage capacity is extended, but not file systems. You can allocate your own storage space after resizing. For more information, see Next operations.

What is not supported

- You cannot resize a cloud disk for which a snapshot is being created.

- After you click Renew for Configuration Downgrade for the subscription instance, you cannot resize subscription cloud disks of the instance during the remaining period of the current billing cycle.

- If a data disk adopts the MBR partition format, you cannot resize the data disk to more than 2 TiB. If you want to resize a data disk to 2 TiB while the MBR partition format is used, we recommend that you create and attach another data disk. Format a GPT partition and copy the data in the MBR partition to the GPT partition.

What is supported

- You can resize cloud disks that are in the in use state. This occurs when instances to which the disks are attached are in the Running(`Running`) state.

- You can resize cloud disks of I/O-optimized instances.

- You can resize ultra disks and SSD disks.

- For Windows instances, you can only resize the NTFS file system.

Preparations

1. Create a snapshot for data backup, to prevent data loss caused by any errors.

2. Check whether you need to update the RedHat VirtIO SICI driver for the instance if your Windows instance was created earlier than March 30, 2019.

Procedure

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

1. Log on to the ECS console.

2. In the left-side navigation pane, choose Storage & Snapshots > Disks.

3. Locate the cloud disk to be extended. Choose More > Resize Disk in the Actions column.

4. Select ✅Resize Online. You do not need to restart the instance to validate new capacity changes (resize and partition operations).

5. Set the extended capacity, which must be greater than the current capacity.

6. Confirm the fee, read ECS Service Terms and Product Terms of Service, and then click Confirm to Resize.

7. Complete the payment.

📋 **Note:**

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

Related API: You can also call ResizeDisk to perform the resize operation.

Next operations

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

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

| Next | If a data disk is in the available (Available) state, the extended capacity takes effect after you complete the payment. Then you can perform the following steps:

1. Log on to the console or use the API (AttachDisk) to attach a cloud disk.
2. Resize or format a partition:

· Partition or format a Linux data disk of less than 2 TiB
· Partition or format a Windows data disk of less than 2 TiB
· Partition or format a data disk of more than 2 TiB | You can perform the following steps:

1. Resize or format a partition:

· Partition or format a Linux data disk of less than 2 TiB
· Partition or format a Windows data disk of less than 2 TiB
· Partition or format a data disk of more than 2 TiB

2. Resize a file system:

· For a Linux instance, see Resize a partition and file system of a Linux data disk or Resize a partition and file system of a Linux system disk
· For a Windows instance, see Resize a file system of a Windows disk | Resize a partition and file system:

· For a Linux instance, see Resize a partition and file system of a Linux data disk or Resize a partition and file system of a Linux system disk
· For a Windows instance, see Resize a file system of a Windows disk |

3.8.3 Resize cloud disks offline

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

What is offline resizing?

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

Notes

System limits

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

What is not supported

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

What is supported

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

Preparations

Create a snapshot for data backup, to prevent data loss caused by any errors.

Procedure

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

1. Log on to the ECS console.
2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
3. Locate the cloud disk to be extended. Choose More > Resize Disk in the Actions column.

4. Set the extended capacity, which must be greater than the current capacity.

5. Confirm the fee, read ECS Service Terms and Product Terms of Service, and then click Confirm to Resize.

6. Complete the payment.

7. Restart the instance for the changes to take effect.

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

Related API: You can also call ResizeDisk to perform the resize operation.

Next operations

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

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

| Next | If a data disk is in the available (Available) state, the extended capacity takes effect after you complete the payment. Then you can perform the following steps:

1. Log on to the console or use the API (AttachDisk) to attach a cloud disk.
2. Resize or format a partition:
 · Partition or format a Linux data disk of less than 2 TiB
 · Partition or format a Windows data disk of less than 2 TiB
 · Partition or format a data disk of more than 2 TiB | You can perform the following steps:

1. Resize or format a partition:
 · Partition or format a Linux data disk of less than 2 TiB
 · Partition or format a Windows data disk of less than 2 TiB
 · Partition or format a data disk of more than 2 TiB
2. Resize a file system:
 · For a Linux instance, see Resize a partition and file system of a Linux data disk or Resize a partition and file system of a Linux system disk
 · For a Windows instance, see Resize a file system of a Windows disk | Resize a partition and file system:

· For a Linux instance, see Resize a partition and file system of a Linux data disk or Resize a partition and file system of a Linux system disk
· For a Windows instance, see Resize a file system of a Windows disk |

3.8.4 Extend a Windows file system

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

Limits

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

disks that are in the Available state, see Attach a cloud disk and Partition and format a
data disk.

Preparations

1. Use the ECS console or call the API to resize the cloud disk.

2. Create a snapshot to back up your data.

3. Attach the cloud disk to the instance and make sure that the instance is in the
   Running state. For information about the connection methods, see Overview.

4. Format and partition the data disk. For more information, see Partition and format
   a data disk.

Extend the system disk partition

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

1. Connect to the Windows instance.

2. Open Server Manager.

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

4. Choose Action > Refresh or Action > Rescan Disks.
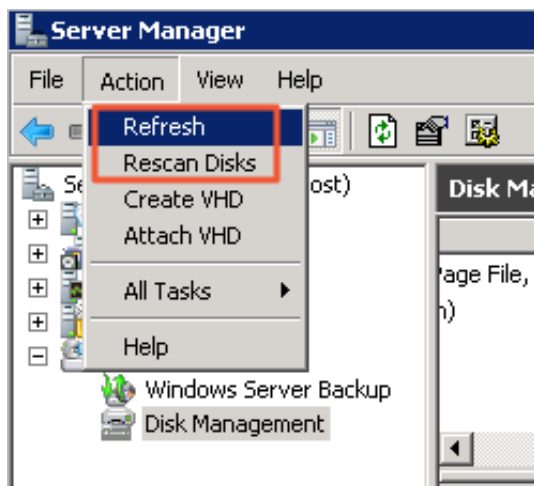
5. In the Disk Management area, view the unallocated capacity. In this example, Disk 0 is the resized system disk.
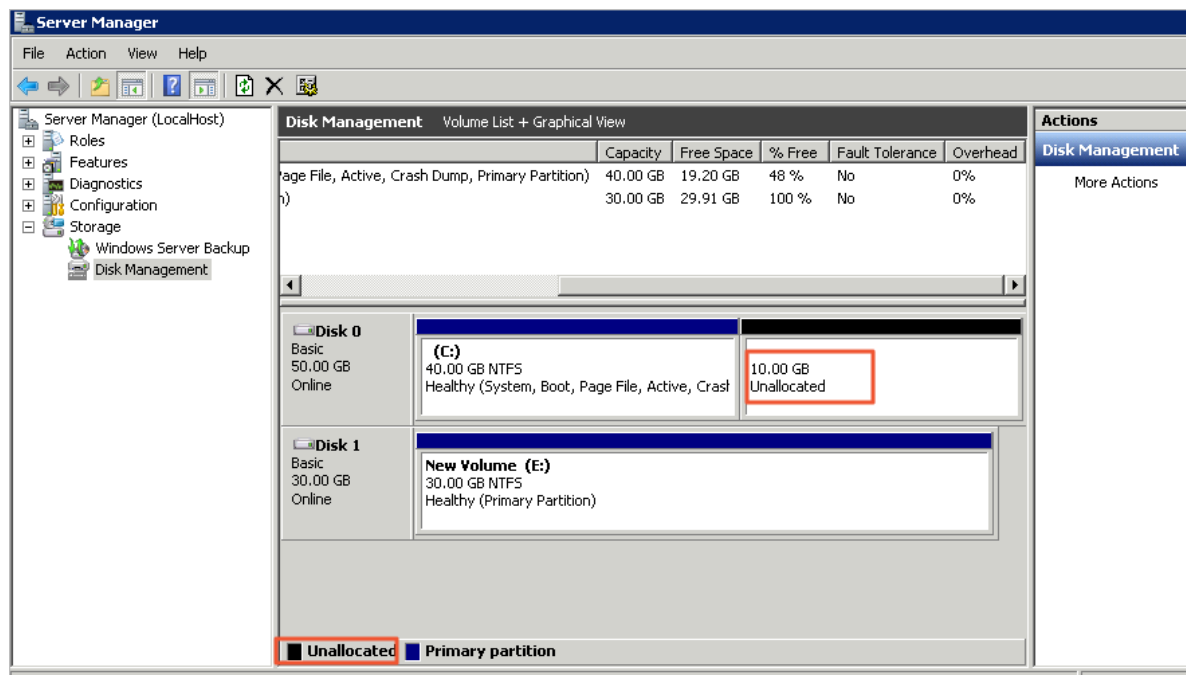
6. Right-click the blank space in the Disk 0 area, and then select Extend Volume.

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

The new disk space is automatically added to the original volume.
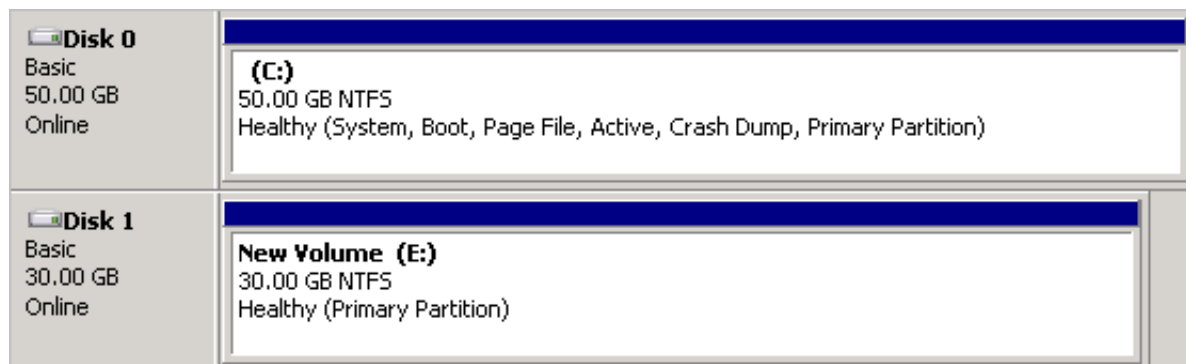
Extend a data disk partition

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

1. Connect to the Windows instance.

2. Open Server Manager.

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

4. Choose Action > Refresh or Action > Rescan Disks.
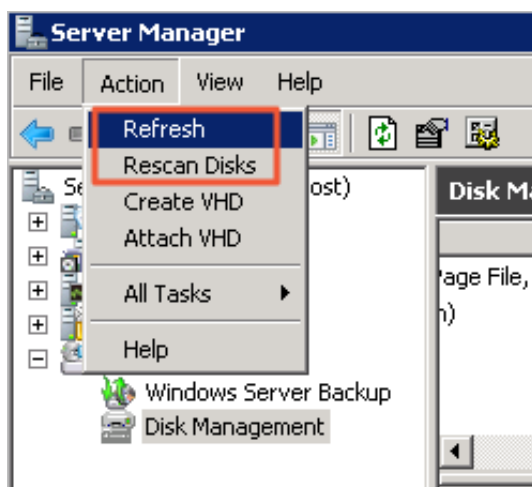
5. In the Disk Management area, view the unallocated capacity. In this example, Disk 1 is the resized data disk.
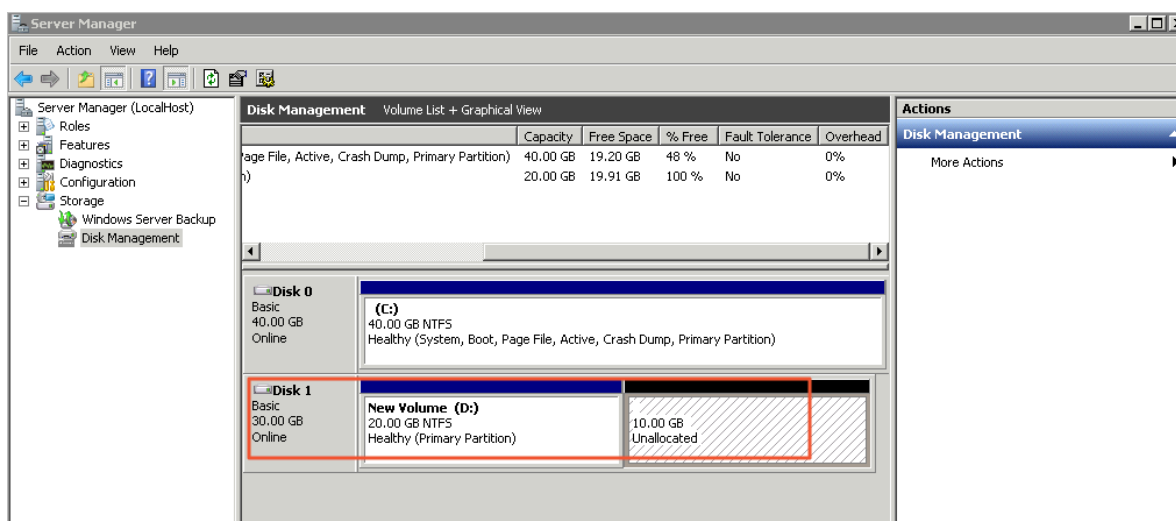
6. **Extend Disk 1.**

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

   a. Right-click the blank space in the Disk 1 area, and then select Extend Volume.
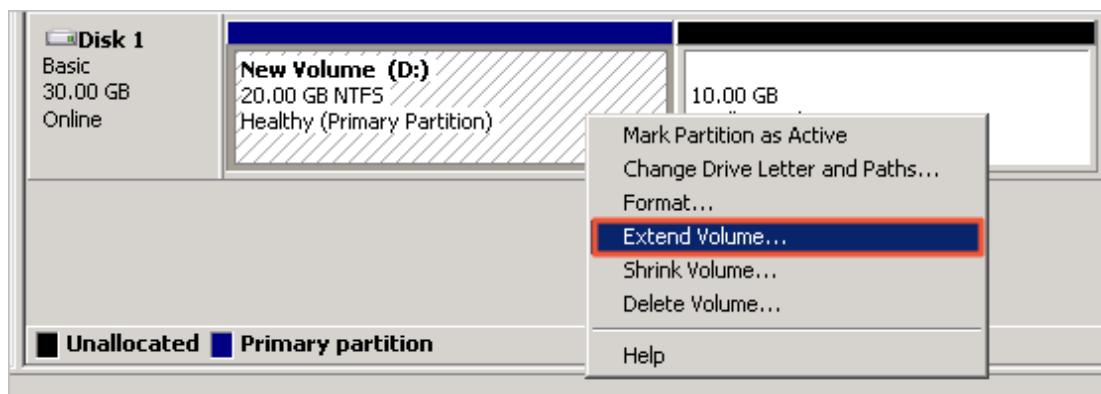
   b. Follow the instructions provided by the Extend Volume Wizard to extend the volume.

   The new disk space is automatically added to the original volume.
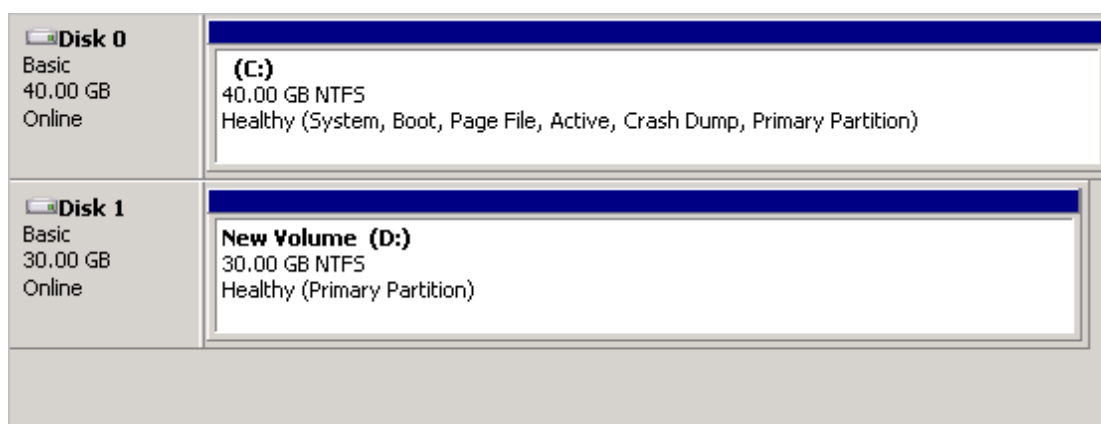
   · To use the new disk space to add a new partition, follow these steps:

   a. Right-click the bank space in the Disk 1 area, and then select New Simple Volume.
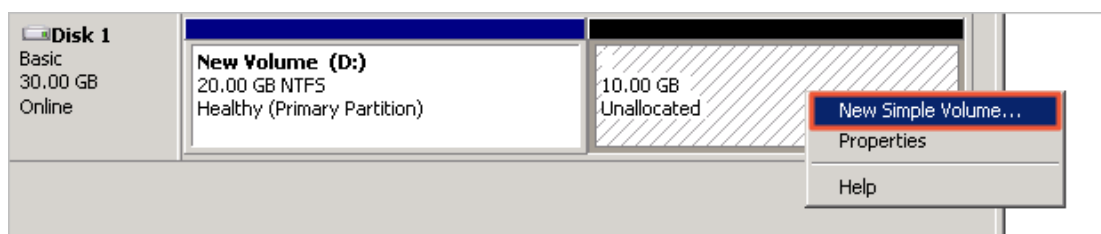
   b. Follow the instructions provided by the New Simple Volume Wizard to extend the volume.

   The new disk space is added to a new partition.

3.8.5 Extend the file system of the Linux system disk

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

Overview

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

· You can resize the disk only if a combination of the following conditions are met. For more information, see Resize the system disk online.

    - Partition format: `mbr` and `gpt`

    - File system: ext, XFS, Btrfs, and UFS

    - Operating system: The kernel version is V3.6.0 or later (run the `uname -a` command to check the version).

· If the operating system is earlier than V3.6.0 (for example CentOS 6, Debian 7, and SUSE Linux Enterprise Server 11 SP4), you must restart the operating system before you can resize the system disk. For more information, see Resize the system disk offline.

Preparations

1. Resize the disk by using the ECS console or calling the API.

2. Create a snapshot to back up data.

3. Mount the disk to an ECS instance that is in the Running state. For information on how to connect to an ECS instance, see Overview.

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

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

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

Resize the system disk online

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

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

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

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

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

 Disk  / dev / vda :  107 . 4  GB ,  1073741824  00  bytes ,
 209715200  sectors
 Units  =  sectors  of  1  *  512  =  512  bytes
 Sector  size  ( logical / physical ):  512  bytes  /  512
 bytes
 I / O  size  ( minimum / optimal ):  512  bytes  /  512  bytes
 Disk  label  type :  dos
 Disk  identifier :  0x0008d73a

 Device  Boot  Start  End  Blocks  Id
 System
```

```
/ dev / vda1     *         2048        41943039          20970496         83
  Linux
```

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

```
[ root @ localhost  ~]#  df  - h
 Filesystem         Size    Used   Avail   Use %  Mounted    on
/ dev / vda1          20G    1 . 5G      18G     8 % /
 devtmpfs         7 . 8G       0    7 . 8G     0 % / dev
 tmpfs            7 . 8G     344K    7 . 8G     1 % / run
 tmpfs            7 . 8G       0    7 . 8G     0 % / sys / fs /
 cgroup
 tmpfs            1 . 6G       0    1 . 6G     0 % / run / user / 0
```

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

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

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

```
[ root @ localhost  ~]#  resize2fs  / dev / vda1
 resize2fs   1 . 42 . 9  ( 28 - Dec - 2013 )
 Filesystem   at / dev / vda1   is   mounted   on  /;  on - line
 resizing   required
 old_desc_b  locks  =  2 ,  new_desc_b  locks  =  7
 The   filesystem   on  / dev / vda1   is   now   26213807   blocks
   long .
```

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

```
[ root @ localhost  ~]#  df  - h
 Filesystem         Size    Used   Avail    Use %  Mounted    on
/ dev / vda1          99G    1 . 6G      93G      2 % /
 devtmpfs         7 . 8G       0    7 . 8G     0 % / dev
 tmpfs            7 . 8G       0    7 . 8G     0 % / dev / shm
 tmpfs            7 . 8G     500K    7 . 8G     1 % / run
 tmpfs            7 . 8G       0    7 . 8G     0 % / sys / fs /
 cgroup
```

```
tmpfs              1 . 6G       0    1 . 6G      0 % / run / user / 0
```

Resize the system disk offline

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

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

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

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

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

 Disk  / dev / vda :  107 . 4   GB ,  1073741824  00   bytes
 255   heads ,  63   sectors / track ,  13054   cylinders
 Units  =  cylinders   of   16065  *  512  =  8225280   bytes
 Sector   size  ( logical / physical ):  512   bytes  /  512
 bytes
 I / O  size  ( minimum / optimal ):  512   bytes  /  512   bytes
 Disk   identifier :   0x0003a7b4

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

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

```
[ root @ AliYunOS  ~]#  df  - h
 Filesystem        Size    Used   Avail   Use %  Mounted    on
/ dev / vda1         20G    1 . 1G      18G     6 % /
 tmpfs            7 . 8G       0    7 . 8G     0 % / dev / shm
```

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

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

5. Restart the instance.

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

6. Connect to the instance again.

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

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

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

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

 Disk  / dev / vda :  107 . 4   GB ,  1073741824  00   bytes
 255   heads ,  63   sectors / track ,  13054   cylinders
 Units  = cylinders   of   16065  *  512  =  8225280   bytes
 Sector   size  ( logical / physical ):  512   bytes  /  512
 bytes
 I / O   size  ( minimum / optimal ):  512   bytes  /  512   bytes
 Disk   identifier :  0x0003a7b4

 Device   Boot       Start         End       Blocks    Id
 System
/ dev / vda1    *          1      13054    104855231    83
 Linux
[ root @ AliYunOS  ~]#  df  - h
 Filesystem       Size    Used   Avail   Use %  Mounted   on
/ dev / vda1       99G    1 . 1G     93G     2 % /
 tmpfs           7 . 8G       0   7 . 8G     0 % / dev / shm
```

3.8.6 Resize partitions and file systems of Linux data disks

When you resize cloud disks, only the storage capacity is extended. File systems are not affected. You must follow the steps in this topic to resize file systems to extend the capacity for ECS instances.

Preparations

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

1. Create a snapshot for data backup, to prevent data loss caused by misoperations.

2. Use the ECS console or API to Resize cloud disks offline.

3. Connect to the ECS instance remotely. For more information about connection methods, see Methods to connect to an ECS instance.

Check the partition format and the file system type

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

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

   In the example, the disk has partition `/ dev / vdb1 ."  System "=" Linux "` indicates that the partition scheme of the disk is MBR. If `" System "=" GPT "` is displayed, it indicates that the partition scheme of the disk is GPT.

   ```
   [root@ecshost ~]# fdisk -lu /dev/vdb
   Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 sectors
   Units = sectors of 1 * 512 = 512 bytes
   Sector size (logical/physical): 512 bytes / 512 bytes
   I/O size (minimum/optimal): 512 bytes / 512 bytes
   Disk label type: dos
   Disk identifier: 0x9277b47b

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

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

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

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

   📋 **Note:**

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

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

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

   · For the xfs file system: `xfs_repair -n <dst_dev_part_path>`

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

   ```
   [root@ecshost ~]# e2fsck -n /dev/vdb1
   e2fsck 1.42.9 (28-Dec-2013)
   Warning! /dev/vdb1 is mounted.
   ```

```
    Warning: skipping journal recovery because doing a read-only
     filesystem check.
    /dev/vdb1: clean, 11/1310720 files, 126322/5242624 blocks
```

Select a method to resize partitions or file systems

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

Scenario	Operation
A data disk has partitions and file systems.	· To resize the existing MBR partitions of the data disk, see Option 1: Resize existing MBR partitions. · If the new disk space is used to add new MBR partitions, see Option 2: Add and format MBR partitions. · To resize the existing GPT partitions of the data disk, see Option 3: Resize existing GPT partitions. · If the new disk space is used to add new GPT partitions, see Option 4: Add and format GPT partitions.
A new data disk does not have partitions or file systems.	Partition and format a data disk or Partition and format data disk more than 2 TiB after you resize the disk space in the console.
A data disk has file systems, but not partitions.	Resize a file system after you resize the disk space in the console.
A data disk is not attached to any instance.	After you attach the data disk to an instance, perform the steps in this topic to resize the data disk.

Note:

If an existing partition of a data disk is created by using MBR, you cannot resize the data disk to greater than or equal to 2 TiB. To prevent data loss, we recommend that you create a cloud disk larger than 2 TiB. Format a GPT partition as specified in Partition and format data disk more than 2 TiB and copy the data in the MBR partition to the GPT partition.

Option 1: Resize existing MBR partitions

Note:

To prevent data loss, we do not recommend that you resize partitions and file systems when they are attached to ECS instances. Detach the attached partitions (unmount) first. After you resize the partitions and they can be normally used, attach the

partitions (`mount`) again. The following operation methods are recommended for different Linux kernel versions:

· If the instance kernel version is earlier than 3.6, detach the partition, modify the partition table, and then resize the file system.

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

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

Step 1: Modify the partition table

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

   In this example, the start sector of `/ dev / vdb1` is `2 , 048` and the end sector is `41 , 943 , 039` .

```
[root@ecshost ~]# fdisk -lu /dev/vdb
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x9277b47b

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

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

   a. Run the `fdisk - u / dev / vdb` command to partition the data disk.

   b. Enter p to display the partition table.

   c. Enter d to delete the partition.

   d. Enter p to confirm that the partition has been deleted.

   e. Enter w to save changes and exit.

```
[root@ecshost ~]# fdisk -u /dev/vdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Command (m for help): p
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x9277b47b
Device Boot Start End Blocks Id System
/dev/vdb1 2048 41943039 20970496 83 Linux
```

```
Command (m for help): d
Selected partition 1
Partition 1 is deleted
Command (m for help): p
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x9277b47b
Device Boot Start End Blocks Id System
Command (m for help): w
The partition table has been altered!
Calling ioctl() to re-read partition table.
WARNING: Re-reading the partition table failed with error 16: Device
 or resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
```

3. Run the `fdisk` command to create a partition.

   a. Run the `fdisk - u / dev / vdb` command to partition the data disk.

   b. Enter p to display the partition table.

   c. Enter n to create a partition.

   d. Enter p to set the partition as the primary partition.

   e. Enter `<partition number>` to select a partition number. 1 is selected in the example.

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

   f. Enter w to save changes and exit.

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

```
[root@ecshost ~]# fdisk -u /dev/vdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Command (m for help): p
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x9277b47b
Device Boot Start End Blocks Id System
Command (m for help): n
Partition type:
p primary (0 primary, 0 extended, 4 free)
```

```
e extended
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-83886079, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-83886079, default
 83886079):
Partition 1 of type Linux and of size 30 GiB is set
Command (m for help): p
Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x9277b47b
Device Boot Start End Blocks Id System
/dev/vdb1 2048 62916607 31457280 83 Linux
Command (m for help): w
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
```

4. Run the `lsblk` */dev/vdb* command to confirm that the partition table has been
   added.

5. Run the `e2fsck -n` */dev/vdb1* command to confirm that the file system after the
   resize operation is in the `clean` state.

Step 2: Notify the kernel of updating the partition table

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

Step 3: Resize the file system

· For the ext* file system (such as ext3 and ext4): Run the ` resize2fs   / dev / `
  ` vdb1 ` command.

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

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

```
[root@ecshost ~]# mount /dev/vdb1 /mnt/
[root@ecshost ~]# xfs_growfs /dev/vdb1
meta-data=/dev/vdb1                isize=512    agcount=4,
 agsize=1310720 blks
         =                         sectsz=512   attr=2, projid32bit=1
         =                         crc=1        finobt=0 spinodes=0
data     =                         bsize=4096   blocks=5242880,
 imaxpct=25
         =                         sunit=0      swidth=0 blks
```

```
naming    =version 2                  bsize=4096    ascii-ci=0 ftype=1
log       =internal                   bsize=4096    blocks=2560, version=2
          =                           sectsz=512    sunit=0 blks, lazy-
count=1
realtime =none                        extsz=4096    blocks=0, rtextents=0
data blocks changed from 5242880 to 7864320
```

Option 2: Add and format MBR partitions

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

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

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

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

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

   Command (m for help): p

   Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 sectors
   Units = sectors of 1 * 512 = 512 bytes
   Sector size (logical/physical): 512 bytes / 512 bytes
   I/O size (minimum/optimal): 512 bytes / 512 bytes
   Disk label type: dos
   Disk identifier: 0x2b31a2a3

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

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

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

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

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

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

```
└─vdb2 253:18   0  20G  0 part
```

3. **Format the new partition.**

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

    ```
    [root@ecshost ~]# mkfs.ext4 /dev/vdb2
    mke2fs 1.42.9 (28-Dec-2013)
    Filesystem label=
    OS type: Linux
    Block size=4096 (log=2)
    Fragment size=4096 (log=2)
    Stride=0 blocks, Stripe width=0 blocks
    1310720 inodes, 5242880 blocks
    262144 blocks (5.00%) reserved for the super user
    First data block=0
    Maximum filesystem blocks=2153775104
    160 block groups
    32768 blocks per group, 32768 fragments per group
    8192 inodes per group
    Superblock backups stored on blocks:
            32768, 98304, 163840, 229376, 294912, 819200, 884736,
     1605632, 2654208,
            4096000

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

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

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

    Allocating group tables: done
    Writing inode tables: done
    Creating journal (32768 blocks): done
    Writing superblocks and filesystem accounting information: done
    [root@ecshost ~]# blkid /dev/vdb2
    ```

```
/dev/vdb2: UUID="dd5be97d-a630-4593-****-5056def914ea"
 SEC_TYPE="ext2" TYPE="ext3"
```

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

```
[root@ecshost ~]# mkfs.xfs -f /dev/vdb2
meta-data=/dev/vdb2                   isize=512     agcount=4,
 agsize=1310720 blks
         =                           sectsz=512    attr=2,
 projid32bit=1
         =                           crc=1         finobt=0, sparse=0
data     =                           bsize=4096    blocks=5242880,
 imaxpct=25
         =                           sunit=0       swidth=0 blks
naming   =version 2                  bsize=4096    ascii-ci=0 ftype=1
log      =internal log               bsize=4096    blocks=2560,
 version=2
         =                           sectsz=512    sunit=0 blks, lazy-
count=1
realtime =none                      extsz=4096    blocks=0,
 rtextents=0
[root@ecshost ~]# blkid /dev/vdb2
/dev/vdb2: UUID="66251477-3ae4-4b44-****-5604420dbecb" TYPE="xfs"
```

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

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

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

4. **Run the** `mount  / dev / vdb2  / mnt` **command to attach the file system to the data disk.**

5. Run the `df  - h` command to check the current capacity and usage of the data disk.

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

```
[root@ecshost ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/vda1 40G 1.6G 36G 5% /
devtmpfs 3.9G 0 3.9G 0% /dev
tmpfs 3.9G 0 3.9G 0% /dev/shm
tmpfs 3.9G 460K 3.9G 1% /run
tmpfs 3.9G 0 3.9G 0% /sys/fs/cgroup
/dev/vdb2 9.8G 37M 9.2G 1% /mnt
tmpfs 783M 0 783M 0% /run/user/0
```

Option 3: Resize existing GPT partitions

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

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

```
[root@ecshost ~]# fdisk -l
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000b1b45

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

Disk /dev/vdb: 35184.4 GB, 35184372088832 bytes, 68719476736 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: gpt
Disk identifier: BCE92401-F427-45CC-8B0D-B30EDF279C2F

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

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

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

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

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

   a. Run the `parted   / dev / vdb` command to start the `parted` tool.

   b. (Optional) Run the `help` command to view the help information.

   c. Run the `print` command to view the number (`Number`) and capacity (`Size`) of the partition to be extended.

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

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

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

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

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

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

Number  Start    End      Size     File system  Name  Flags
 1      1049kB   1100GB   1100GB   ext4         mnt
(parted) resizepart 1 100%
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 35.2TB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

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

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

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

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

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

5. Resize the file system.

   · For the ext* file system (such as ext3 and ext4): run the `resize2fs / dev / vdb1` command.

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

   · For the xfs file system: run the `xfs_growfs / dev / vdb1` command.

6. Attach the partition again.

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

Option 4: Add and format GPT partitions

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

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

   ```
   [root@ecshost ~]# fdisk -l
   Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
   Units = sectors of 1 * 512 = 512 bytes
   Sector size (logical/physical): 512 bytes / 512 bytes
   I/O size (minimum/optimal): 512 bytes / 512 bytes
   Disk label type: dos
   Disk identifier: 0x000b1b45

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

   Disk /dev/vdb: 35184.4 GB, 35184372088832 bytes, 68719476736 sectors
   Units = sectors of 1 * 512 = 512 bytes
   Sector size (logical/physical): 512 bytes / 512 bytes
   I/O size (minimum/optimal): 512 bytes / 512 bytes
   Disk label type: gpt
   Disk identifier: BCE92401-F427-45CC-8B0D-B30EDF279C2F

   #         Start          End    Size  Type            Name
   ```

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

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

   a. Run the `parted   / dev / vdb` command to start the parted tool.

   b. Run the `print    free` command to view the disk space to be allocated.

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

      In the example, the start sector of `/ dev / vdb1` is 1,049 KB, the end sector is

      5,278 GB, and the capacity is 5,278 GiB.

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

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

   c. Run the `mkpart   < partition    name > < start    sector > < capacity`

        `allocation    percentage >` command.

      In the example, the `/ dev / vdb2` partition named test is created. The start

      sector of the new partition is the end sector of the existing partition. The new

      capacity is allocated to the new partition.

   d. Run the `print` command to check whether the capacity (`Size`) of the

      partition changes.

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

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

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

3. **Create a file system for the new partition.**

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

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

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

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

  In the example, an xfs file system is created.

```
[root@ecshost ~]# mkfs –t xfs /dev/vdb2
meta-data=/dev/vdb2              isize=512    agcount=28,
 agsize=268435455 blks
         =                       sectsz=512   attr=2, projid32bit=1
         =                       crc=1        finobt=0, sparse=0
data     =                       bsize=4096   blocks=7301444096,
 imaxpct=5
         =                       sunit=0      swidth=0 blks
naming   =version 2             bsize=4096   ascii-ci=0 ftype=1
log      =internal log          bsize=4096   blocks=521728,
 version=2
         =                       sectsz=512   sunit=0 blks, lazy-
count=1
realtime =none                  extsz=4096   blocks=0, rtextents=0
```

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

```
[root@ecshost ~]# fdisk –l
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000b1b45

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

Disk /dev/vdb: 35184.4 GB, 35184372088832 bytes, 68719476736 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: gpt
Disk identifier: BCE92401-F427-45CC-8B0D-B30EDF279C2F

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

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

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

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

6. Attach the new partition.

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

Related operations

- Extend the file system of the Linux system disk

- Extend a Windows file system

3.8.7 Update a Red Hat VirtIO SICI driver in Windows

This topic describes how to update a Red Hat VirtIO SICI driver running in a Windows instance in Alibaba Cloud. In Alibaba Cloud, you can resize your disk online without the need to restart the instance. However, if your instance is created before March 30, 2019, you need to check whether you need to update the RedHat VirtIO SICI driver according to the information in this topic.

Limits

- The RedHat VirtIO SICI driver is only supported by instances running Windows Server 2008 or later versions.

- If you have multiple data disks in your Windows instance, a driver update may take up to two minutes to be completed.

Download the driver

Download and decompress the instdrv driver package. The following table describes the files and folders that are extracted from the decompressed package.

File/folder name	Description
win7	Windows 7
Wlh	Windows 2008
Win8	Windows 2012
win10	Windows 2016/1709/10
amd64	64-bit
x86	32-bit
instdrv.exe	Checks whether your instance supports online disk resizing programs.

Check the driver version

Method 1: Run the instdrv.exe file

1. Connect to the target Windows instance.

2. Open the Command Prompt (cmd).

3. Enter the directory where `instdrv . exe` is located. In this example, the directory is `cd C:\User|lizr\Desktop\instdrv\`.

4. Optional. Run `instdrv . exe` to view the user guide of the tool. Run the `instdrv . exe - v` command to view the tool version.

5. Run the `instdrv . exe - c` command, and check whether the operating system supports online disk resizing.

> **Note:**
>
> If the response message indicates that support for online disk resizing is available, you can directly resize a disk online. Otherwise, you need to update the RedHat VirtIO SICI driver.

Method 2: Check the driver version manually

1. Connect to the target Windows instance.

2. Enter the system directory `C :\ Windows \ System32 \ drivers .`

3. Right-click the `viostor . sys` file, and then select Properties. Choose Details >
   File Version.

> 📋 **Note:**
>
> If the version of the `viostor . sys` file is `58011` or later, you can directly
> resize a disk online. If the file version is earlier than `58011` , you need to update
> the RedHat VirtIO SICI driver.
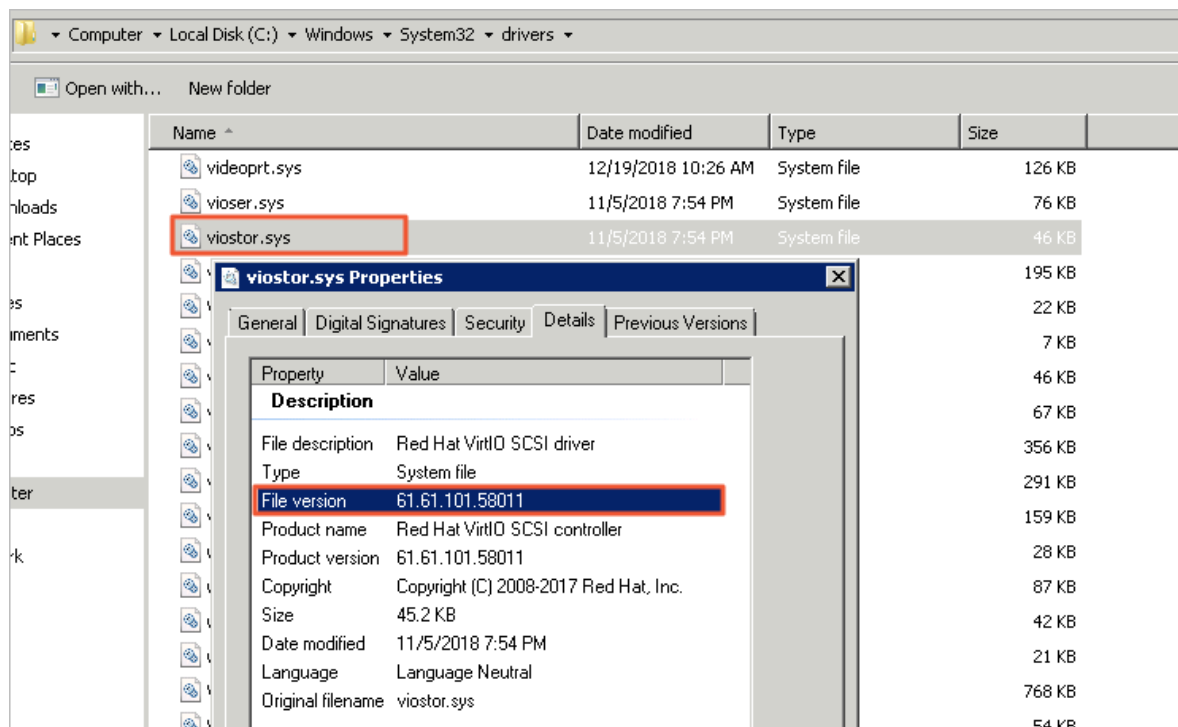
Update the RedHat VirtIO SICI driver

Method 1: Run the `pnputil` command

1. Open the Command Prompt (cmd).
2. Enter the directory where the `viostor . inf` file is located. In this example, the
   directory is `cd C:\Users\Administrator\Desktop\amd`.
3. Decompress the target `. inf` file to the specified `<directory>`.

4. Run the `pnputil -i -a <directory>` command to add the driver package.

5. Log on to the ECS console and restart the instance, or call the API RebootInstance to restart the instance to validate the updated driver.

Method 2: Update the driver manually

1. Open the Device Manager window.

2. Right-click Red Hat Vir under Disk drivers, and then select Update Driver Software.
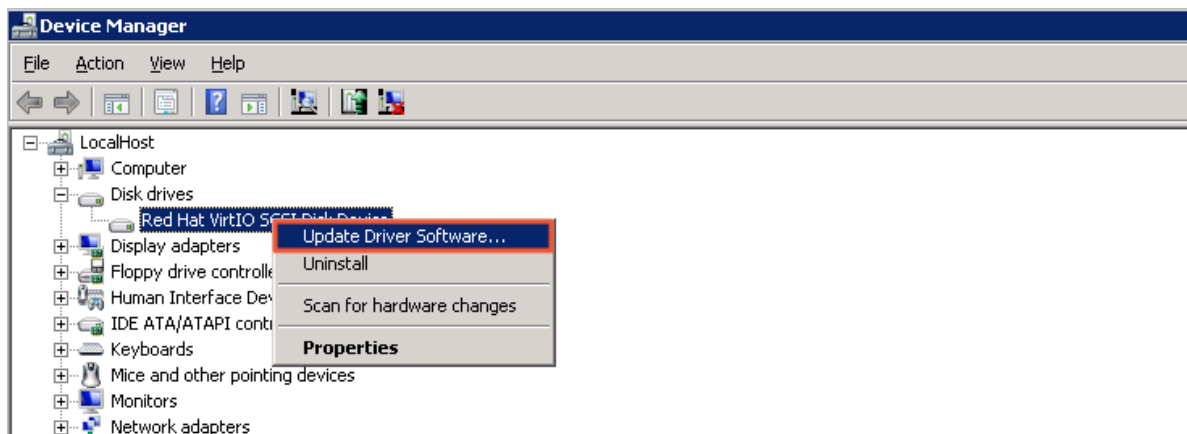
3. Select Browse my computer for driver software.

4. Select Let me pick from a list of device drivers on my computer.

5. Click Hard Disk.

6. Select the driver file `viostor` from the target folder, and then update the driver by following the wizard.

7. Log on to the ECS console and restart the instance, or call the API RebootInstance to restart the instance to validate the updated driver.

What to do next

Resize a disk online.

3.9 Change the operating system

3.9.1 Replace the system disk (non-public image)

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

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

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

After you replace the system disk, note that:

· A new system disk with a new disk ID is allocated to your instance, and the original ID is released.

· The cloud type of the cloud disk cannot be replaced.

· The IP address and the MAC address remain unchanged.

· We recommend that you delete snapshots or automatic snapshot policies to ensure sufficient snapshot quota for executing automatic snapshot policies of the new system disk.

This topic describes how to replace an existing image with a non-public image. If you need to use a public image, see Replace the system disk by using a 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 unidentifi able.

- If no important data exists on the data disk, we recommend that you reinitialize the disk and format it to the default file system of your OS.
- If important data exists in your data disk, perform the following actions as required:
  - From Windows to Linux, you must install a software application, for example, NTFS-3G, because NTFS cannot be identified by Linux.
  - From Linux to Windows, you must install a software application, for example , Ext2Read or Ext2Fsd, because ext3, ext4, and XFS cannot be identified by Windows.

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

Prerequisites

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

· If the target image is a custom image:

   - If you want to use an image of a specified ECS instance, you must create a snapshot for the system disk of the specified instance and create a custom image using a snapshot. If the specified instance and the one whose system disk you want to change are located in different regions, you need to copy the images.

   - To use a local physical image file, import it on the ECS console or use Packer to create and import the local image. The region where the image is located must be the same as that of your instance.

   - To use an image in a region other than that of your instance, you must copy the image first.

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

· To use an image owned by another Alibaba Cloud account, the account must first share the image.

· If you want to replace the OS to Linux and use an SSH key pair for authentication, you must first create an SSH key pair.

· Replacing the system disk may cause data loss or service interruption. To minimize impact to your business services, we recommend that you create snapshots for the original system disk before replacement.

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

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

Procedure

1. Log on to the ECS console.

2. In the left-side navigation pane, choose Instances & Images > Instances.

3. In the Actions column of the target instance, choose More > Instance Status > Stop and follow the instructions in the prompt to stop the instance.

   The action is successful when the instance status is Stopped.

4. In the Actions column, choose More > Disk and Image > Replace System Disk.

5. In the displayed dialog box, read the precautionary statement about system disk replacement, and then click OK.

6. On the Replace System Disk page, complete the following settings:

   a. Image Type: Select Custom Image, Shared Image, or Marketplace Image, and then select the image version.

   b. System Disk: Unchangeable. However, you can expand the disk space to meet the requirements of your system disk and services. The maximum disk space is 500 GiB. The minimum space of the system disk you can configured is determined by the current disk space and image type.

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

   Note:

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

   c. Security enhancement:

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

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

   e. Check the configuration and click Confirm to change.

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

What to do next

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

- (Optional) Apply automatic snapshot policies to disks. Automatic snapshot policies are bound to the disk ID. After the system disk is replaced, automatic snapshot policies applied on the original disk automatically fail. You need to configure automatic snapshot policies for the new system disk.
- If the OS is Linux before and after disk replacement, and if a data disk is mounted to the instance and the partition is set to be mounted automatically at instance startup, then all mounting information will be lost. In this case, you need to write the new partition information to the /etc/fstab file of the new system disk and mount the partition, but do not need to partition or format the data disk again. The

steps are described as follows. For more information, see Format a data disk of a Linux instance.

1. (Recommended) Back up the /etc/fstab file.

2. Write information about the new partition into the /etc/fstab file.

3. Check the information in the /etc/fstab file.

4. Run `mount` to mount the partition.

5. Run `df - h  - h` to check the file system space and usage.

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

Related API

ReplaceSystemDisk

3.9.2 Replace the system disk by using a public image

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

Scenarios

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

· You selected an incorrect OS when creating an ECS instance.

· You need to replace the current OS.

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

After you replace the system disk:

· A new system disk with a new ID is allocated to your instance, and the original system disk is released.

· The cloud disk type remains unchanged.

· The IP address and the MAC address remain unchanged.

· We recommend that you delete snapshots or automatic snapshot policies to ensure that you have sufficient snapshots available for the automatic snapshot policies of the new system disk.

Precautions

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

Risks

Replacing the system disk may involve the following risks:

· The original system disk will be released. Therefore, we recommend that you create a snapshot to back up your data before replacing the system disk.

· Your instance will be stopped and your services interrupted.

· You must redeploy the service environment on the new system disk. However, this may result in an interruption to your services.

· The disk ID will be changed. Therefore, snapshots of the original system disk cannot be used to roll back the new system disk.

> **Note:**
>
> After you replace the system disk, the snapshots you have manually created are not affected. You can still use them to create custom images.
>
> If you have configured automatic snapshot policies for the original system disk to allow automatic snapshots to be released along with the disk, the snapshot policies will no longer apply. All automatic snapshots of the original system disk will be automatically deleted.

Precautions during cross-OS disk replacements

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

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

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

- In case no important data exists on the data disk, we recommend that you reinitialize the disk and format it to the default file system of your OS.

- In case important data exists in your data disk, perform the following operations as required:

    - For switches from Windows to Linux, you must install a software application, such as NTFS-3G, because Linux cannot identify NTFS.

    - For switches from Linux to Windows, you also must install a software applicatio n, such as Ext2Read or Ext2Fsd, because Windows cannot identify ext3, ext4, or XFS.

Precautions during Windows OS replacements

- Windows OS only supports password authentication.

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

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

OS version	Image ID
Windows Server 2008 R2 Enterprise Edition (English)	win2008r2_64_ent_sp1_en-us_40G_ali base_20170915.vhd
Windows Server 2008 R2 Enterprise Edition (Chinese)	win2008r2_64_ent_sp1_zh-cn_40G_ali base_20170915.vhd
Windows Server 2012 R2 Datacenter Edition (Chinese)	win2012r2_64_dtc_17196_zh-cn_40G_ali base_20170915.vhd
Windows Server 2012 R2 Datacenter Edition (English)	win2012r2_64_dtc_17196_en-us_40G_ali base_20170915.vhd
Windows Server 2016 Datacenter Edition (Chinese)	win2016_64_dtc_1607_zh-cn_40G_alibase_20170 915.vhd
Windows Server 2016 Datacenter Edition (English)	win2016_64_dtc_1607_en-us_40G_alibase_20170 915.vhd

Preparations

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

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

Procedure

To replace the system disk, follow these steps:

1. Log on to the ECS console.
2. In the left-side navigation pane, choose Instances & Images > Instances.
3. In the Actions column of the target instance, choose More > Instance Status > Stop and stop the instance as prompted.

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

4. From the Actions column, choose More > Disk and Image > Replace System Disk.
5. In the displayed dialog box, read the precautionary statement about system disk replacement and then click OK.
6. On the Replace System Disk page, set the following parameters:

   a. Image Type: Select Public Image and then select the image version.

   b. System Disk: Unchangeable. However, you can expand the disk space to meet the requirements of your system disk and services. The maximum disk space is 500

GiB. The minimum disk space you can configure depends on the current disk space and image type.

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

📋 Note:

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

c. Security enhancement:

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

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

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

e. Check your settings and click Confirm to change.

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

What to do next

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

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

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

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

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

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

Related API

[ReplaceSystemDisk](#)

3.10 Roll back a cloud disk

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

Limits

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

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

Prerequisites

Before rolling back a cloud disk, check that:

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

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

Procedure

To roll back a cloud disk , follow these steps:

1. Log on to the ECS console.

2. In the left-side navigation pane, choose Instances & Images > Instances.

3. Find the target instance and click its ID to go to the Instance Details page.

4. In the left-side navigation pane, click Instance Snapshots.

5. Find the target snapshot and then in the Actions column, click Disk Rollback.

6. In the dialog box, click OK.

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

Related API

ResetDisk

What to do next

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

· Linux _ Resize a data disk
· Windows _ Resize a data disk

3.11 Reinitialize a cloud disk

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

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

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

- Depending on how the data disk was created, it is restored to the following initial status:

  - Restored to an empty disk if it was an empty disk

  - Restored to a disk with only the data of the source snapshot if it was created from a snapshot

- If an automatic snapshot policy is applied to a cloud disk, the policy is retained and does not need to be applied again after reinitialization.

- If an automatic snapshot policy is applied to a cloud disk, the policy is retained and does not need to be applied again after reinitialization.

- After a cloud disk is reinitialized, all the snapshots, both automatically and manually created, are retained. You can use them to roll back a cloud disk.

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

Reinitialize a system disk

Prerequisites

If an SSH key pair is used as the authentication method, check that you have created an SSH key pair or import ed an SSH key pair.

Procedure

To reinitialize a system disk, follow these steps:

1. Log on to the ECS console.

2. In the left-side navigation pane, choose Instances & Images > Instances.

3. Find the target ECS instance and click its ID to go to the Instance Details page.

4. 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 instance   ✕

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

Stopped By: ◉ Stop
                   ○ Force Stop

Stopped By: ☑ Keep Stopped Instances and Continue Billing

Stopping Subscription based instances does not change their expiration date.

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

OK   Cancel

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

7. 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.

      · 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.

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.

8. For Linux instances: If you have attached a data disk to the instance, connect to the instance and create a mounting point for the partitions of data disks, because the mounting points are lost after the system disk is reinitialized.

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

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

Reinitialize a data disk

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

· For a Windows instance, the data disk is ready to use without any additional operations required.

· For a Linux instance:

  - If the data disk was an empty disk after it was created, then all the data and partitions on the disk are lost. You must partition and format the disk, and mount the partitions again.

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

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

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

Prerequisites

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

Procedure

To reinitialize a data disk, follow these steps:

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

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

 a. Connect to the Linux instance.

 b. Run `vim   / etc / fstab`.

 c. Press the `i` key to enter the Insert mode.

 d. Locate the mounting configuration lines and type # before the lines. For example:

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

 e. Press the `Esc` key to exit the Insert mode, and then run `:wq` to save and exit.

2. Log on to the ECS console.

3. In the left-side navigation pane, choose Instances & Images > Instances.

4. Find the target ECS instance and click its ID to go to the Instance Details page.

5. Click Stop.

> **Note:**
>
> For a Pay-As-You-Go VPC-Connected ECS instance, if the 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.
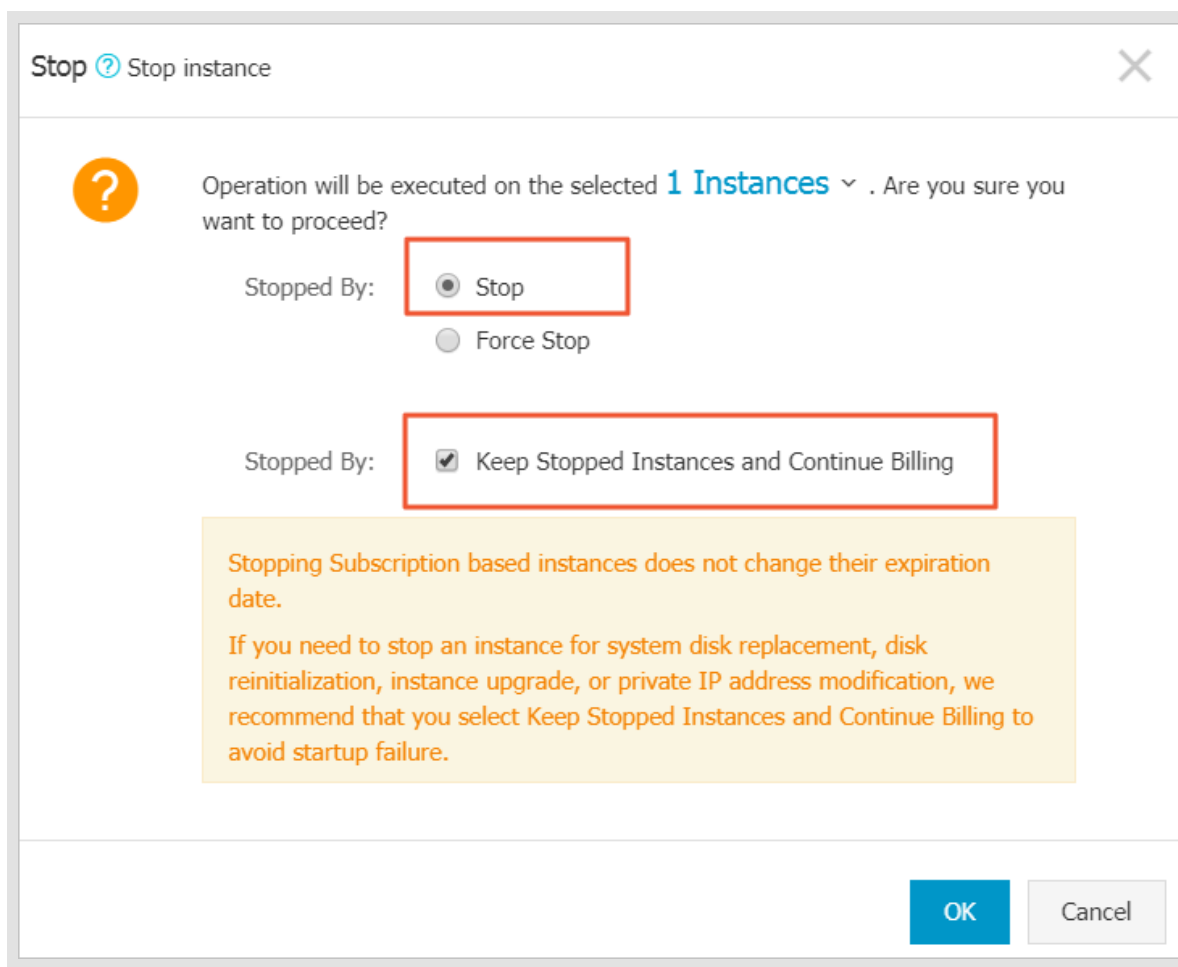
6. After the instance is stopped, click Disks in the left-side navigation pane.

7. Find the target data disk and in the Actions column, click Reinitialize Disk.

8. In the Reinitialize Disk dialog box, read the notes and click Confirm.

9. In the left-side navigation pane, click Instance Details.

10.Click Start.

11.For Linux instances: If the data disk was an empty disk after it was created, format and mount data disks for Linux instances.

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

Related API

   ReInitDisk

3.12 Convert billing methods of cloud disks

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

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

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

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

3.13 Modify the performance level of an ESSD cloud disk

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

Prerequisites

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

- Your account does not have overdue payment.
- If you attach an ESSD cloud disk to a Pay-As-You-Go ECS instance, make sure that the instance is not in the Expired state.

· You can change the performance level of a new ESSD cloud disk only after it enters the Unmounted (`Available`) state.

Context

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

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

You can also call the API action ModifyDiskSpec to complete this task.

Procedure
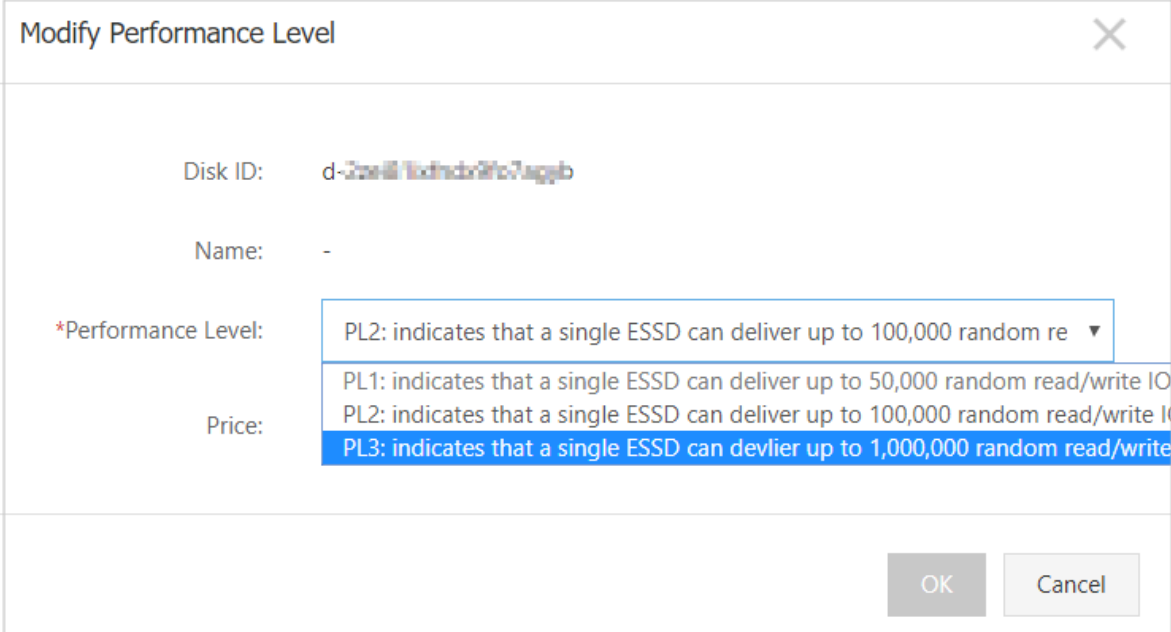
1. Log on to the ECS console.

2. In the left-side navigation pane, choose Storage & Snapshots > Disks.

3. In the upper-left corner, select the target region.
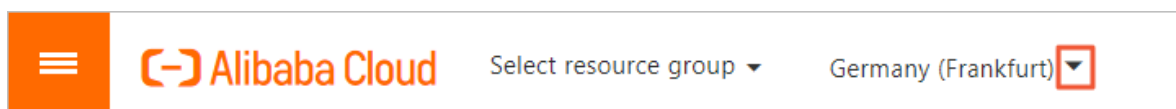
4. Find the target ESSD cloud disk. In the Actions column, choose More > Modify Performance Level.

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

The performance level you can select for an ESSD cloud disk is determined by its storage capacity. If the performance level of an ESSD cloud disk cannot be upgraded, resize the ESSD cloud disk and then modify its performance level.
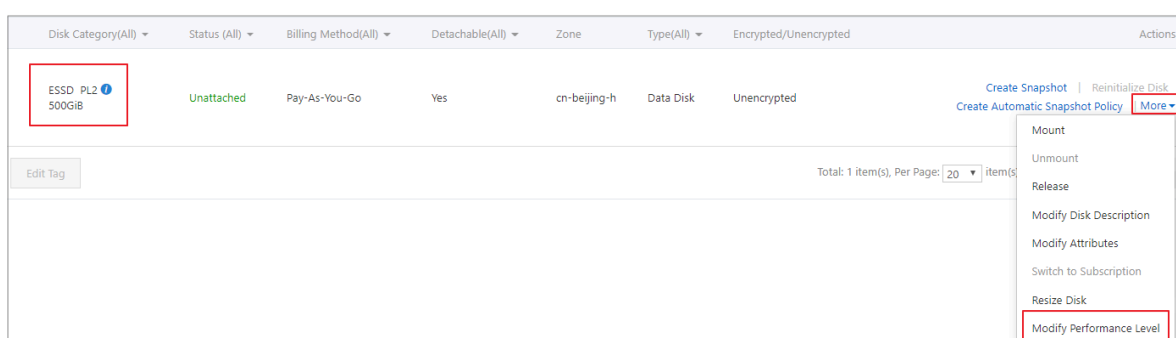
3.14 Monitor a cloud disk

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

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

1. Log on to the ECS console.

2. In the left-side navigation pane, choose Storage & Snapshots > Disks.

3. Find a cloud disk and click its ID to go to the Details page.

4. In the left-side navigation pane, click Disk Monitoring.

5. On the Monitoring Information page, click the ⬛ icon and set the start time and

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

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

📋 **Note:**

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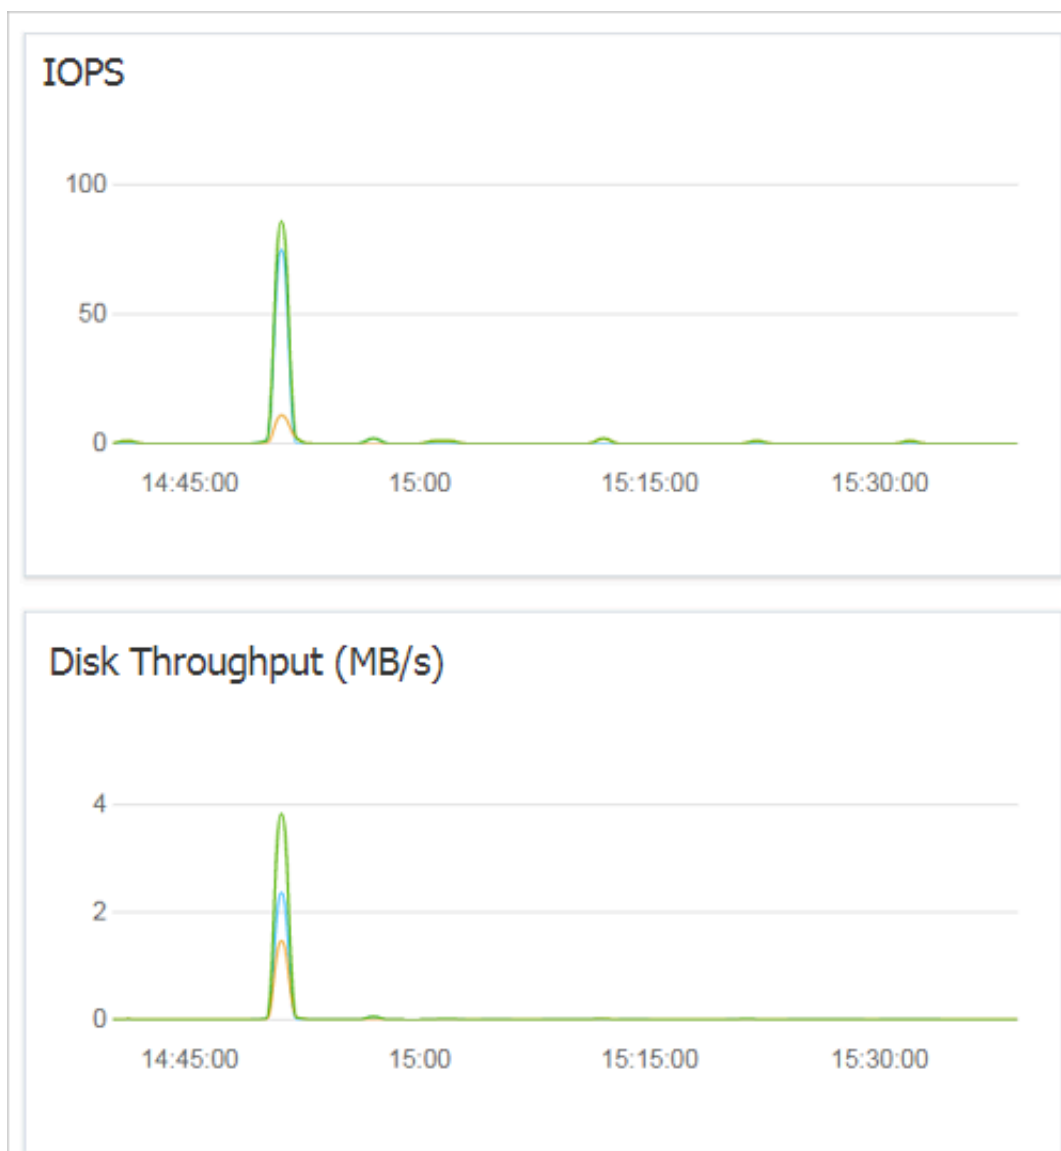
3.15 Detach a cloud disk

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

When detaching a cloud disk, consider the following:

· Only the Pay-As-You-Go cloud disks in the In Use state and used as a Data Disk can be detached.

· You cannot detach a local disk.

· For a Windows instance:

- To guarantee data integrity, we recommend that you stop writing or reading the files on the cloud disk. Otherwise, data may be lost.

- Before detaching a cloud disk in the ECS console, you must connect to the instance and set its status to Offline in Disk Management.

· For a Linux instance:

- Before detaching a cloud disk in the ECS console, you must connect to the instance and run `umount` to unmount the partitions.

- If you have configured the /etc/fstab file to automatically mount the partitions at the startup of the instance, before detaching it, you must delete the 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 are attached to the instance and its status is In Use.

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

Procedure

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

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

Operating system	Steps
Linux	Run `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, choose Instances & Images > Instances.

4. Find the target instance and click its ID to go to the Instance Details page.

5. In the left-side navigation pane, click Disks.

6. 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.

7. In the dialog box, click Confirm.

8. 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 Disks page

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

Prerequisites

The cloud disk is attached to the instance and is in the In Use state.

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

Procedure

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

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

Operating system	Steps
Linux	Run `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, choose Storage & Snapshots > Disks.

4. 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.

5. In the dialog box, click Confirm.

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

Related API

   DetachDisk

What to do next

   If you no longer need the disk, you can release it.

3.16 Release a 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 state can be released independently. Other cloud disks, such as those used as system disks, or Subscription billed cloud disks used as data disks, can only be released together with ECS instances. If a cloud disk is in the In Use state, you must first Detach it from the instance.

· By default, automatic snapshots are released together with their cloud disks. However, manually created snapshots are not. You can change the snapshot release configuration when attaching a cloud disk.

> **Note:**
>
> 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, choose Storage & Snapshots > Disks.

3. Select the disk that you want to release and ensure that it is in the Unmounted state. Then, in the Actions column, choose More > Release.

4. In the Release Disk dialog box, read the note and click Confirm Release.

Related API

DeleteDisk

4 Shared Block Storage

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

> Note:
>
> Currently, Shared Block Storage is available only in public beta release. You can open a ticket to submit an application to join the testing group.

Shared Block Storage replicates data three times by default, and stores the data copies across different servers, to provide 99.9999999% data reliability for ECS instances. This means that if service disruptions occur (for example, due to hardware failure), Shared Block Storage can automatically copy your data within the target zone to help ensure data availability. You can partition and format Shared Block Storage attached to an ECS instance, and create a file system and store data on Shared Block Storage. For more information, see Triplicate technology.

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

Background information

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

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

Limits

Before you use Shared Block Storage, note the following:

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

If you use a traditional file system to manage Shared Block Storage that is attached to multiple ECS instances, the following two errors may occur:

· Disk space allocation conflict

  If a Shared Block Storage device is attached to multiple instances and one of these instances (for example, Instance A) writes data to a file, Instance A checks the file system and available disk space. After the write operation is completed, the space allocation record of Instance A is changed, but the records of other instances are not updated. In this case, when another instance (for example, Instance B) tries to write data to the file, this instance may allocate the disk space that has been already allocated by Instance A, resulting in a disk space allocation conflict.

· Data file inconsistency

  After an instance (for example, Instance A) reads and caches data, another process that accesses the same data will directly read the data from the cache. However, if a copy of the same data that is stored on another instance (for example, Instance B) is modified during this period, and Instance A does not update according to the latest data change, Instance A still reads the data from the cache. As a result, data inconsistency occurs.

Types of Shared Block Storage

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

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

Performance

For information about the performance indexes of Shared Block Storage, see Block storage performance.

Billing

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

5 Local disks

This topic provides information about local disks that you can purchase along with instances.

Background information

Local disks are disks attached to a machine on which ECS instances are hosted, and are designed to provide local storage and access for instances. Features of local disks include low latency, high random IOPS, high throughput, and cost-effective performance. For more information, see Instance type families. If you are using phased-out local SSDs, see Ephemeral SSDs.

Disk types

Currently, Alibaba Cloud provides the following two local disks:

· NVMe SSD: A local disk type that can be used by such instance type families as i2, i1, and gn5. NVMe SSDs are suitable for the following scenarios:

   - Online services such as online gaming, e-commerce, live video, and media, where instance type families such as i1 and i2 meet low latency and high I/O performance required by I/O intensive applications in block storage.
   - Service scenarios with high requirements for I/O performance and high-availability architecture at the application layer, such as NoSQL non-relational databases, MPP data warehouses, and distributed file systems.

· SATA HDD: A local disk type that can be used by such instance type families as d1ne and d1, which are applicable to multiple industries with high requiremen ts for big data computing, storage analysis, massive data storage, and offline computing (for example, e-commerce and financial industries). SATA HDDs are designed to meet the requirements of distributed computing service types (for example, Hadoop) for instance storage performance, storage capacity, and intranet bandwidth.

NVMe SSD performance

The following table describes the performance indexes of the NVMe SSD disk type.

NVMe SSD index	Single disk performance	Overall instance performance
Maximum capacity	1456 GiB	2,912 GiB
Maximum IOPS	240,000	480,000
Write IOPS *	min{165 * capacity, 240,000}	2 * min{165 * capacity, 240,000}
Read IOPS *		
Maximum read throughput	2 Gbit/s	4 Gbit/s
Read throughput *	min{1.4 * capacity, 2,000} Mbit/s	2 * min{1.4 * capacity, 2,000} Mbit/s
Maximum write throughput	1.2 Gbit/s	2.4 Gbit/s
Write throughput *	min{0.85 * capacity, 1,200} Mbit/s	2 * min{0.85 * capacity, 1,200} Mbit/s
Access latency	Microsecond (µs)	

* The performance of a single disk is calculated as follows:

· Write IOPS: 165 IOPS for each GiB, up to 240,000 IOPS

· Write throughput: 0.85 Mbit/s for each GiB, up to 1,200 Mbit/s

To obtain the standard performance data and measure the Quality of Service (QoS) of Alibaba Cloud local disks, you can test the bandwidth, IOPS, and latency of the NVMe SSD by using the method described in the NVMe SSD performance test.

SATA HDD performance

The following table describes the performance indexes of the SATA HDD disk type.

SATA HDD index	Single disk performance	Overall instance performance
Maximum capacity	5,500 GiB	154,000 GiB
Maximum throughput	190 Mbit/s	5,320 Mbit/s
Access latency	Millisecond (ms)	

Billing

The fees billed for local disks are included in the fees billed for the instances to which local disks are attached. For more information about instance billing methods, see Subscription and Pay-As-You-Go.

Limits

- A single point of failure (SPOF) may be evident at the application layer due to each local disk being attached to a single host machine, meaning data reliability is dependent on the reliability of the host machine. We recommend that you implement data redundancy at the application layer to guarantee data availability.

> ⚠️ Warning:
>
> When you use a local disk to store data, there is a risk of data loss (for example, if a hardware fault occurs to the host machine). Therefore, we recommend you do not use a local disk to store any data that you need to keep for a long term. If your applications have no data reliability architecture, we recommend that you use cloud disks or Shared Block Storage in your instances to increase data reliability.

- After you purchase an instance along with a local disk, you need to log on to the instance to partition and format the local disk.
- Regarding disk operations, you are not able to:

    - Create a separate local disk.

    - Use a snapshot to create a local disk.

    - Attach a local disk.

    - Detach and release a local disk separately.

    - Resize a local disk.

    - Reinitialize a local disk.

    - Create a snapshot for a local disk.

    - Use a snapshot to roll back a local disk.

Disk initialization sequence

When you create an instance along with a local disk, all disks are initialized according to the following rules:

- Rule 1: If the specified image does not have a data disk snapshot, the local disk is initialized before the cloud disks created along with the instance.

· Rule 2: If the specified image has a data disk snapshot, the sequence of the data disks recorded in the snapshot is retained, subject to Rule 1.

For example, for a Linux image that includes a snapshot of two data disks, the disks are initialized according to the following sequence.

· If the original device names of the two data disks are `/dev/xvdb` and `/dev/xvdc` respectively, these names are allocated to the specified data disks in the image. The disk initialization starts from the system disk, then proceeds to data disk 1 specified in the image, data disk 2 specified in the image, local disk 1, local disk 2, cloud disk 1, cloud 2, and in such a sequence until initialization is completed.
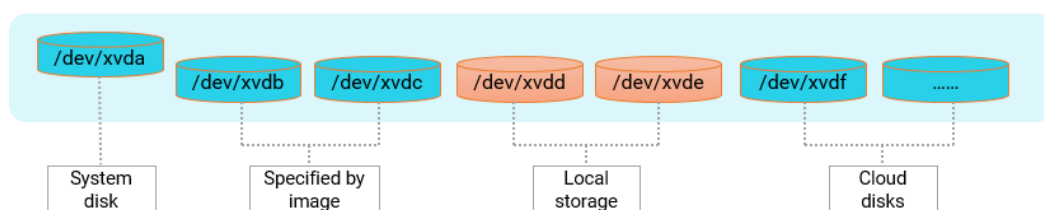
· If the original device names of the two data disks are `/dev/xvdc` and `/dev/xvdd` respectively, these names are allocated to the specified data disks in the image. Then, the remaining device names are allocated to the local disks. The disk initialization starts from the system disk, then proceeds to local disk 1, data disk 1 specified in the image, data disk 2 specified in the image, local disk 2, cloud disk 1, cloud 2, and in such a sequence until initialization is completed.
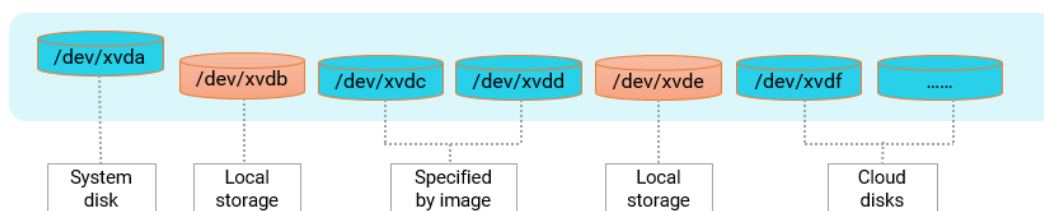
Lifecycle

The lifecycle of a local disk is the same as that of the instance to which the local disk is attached. For more information, see ECS instance lifecycle.

Impact of instance operations on the data stored on local disks

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

Instance operation	Data on the local disk	Local disk
Restart the operating system, restart the instance by using the ECS console, or forcibly restart the instance.	Retained	Retained
Shut down the operating system, stop the instance by using the ECS console, or forcibly stop the instance.	Retained	Retained
Automatically restore the instance.	Erased	Released
Release the instance.	Erased	Released
A Subscription instance is stopped upon expiry or your account has an overdue payment, but the instance is not released .	Retained	Retained
A Subscription instance is stopped upon expiry or your account has an overdue payment, and the instance is released.	Erased	Released

NVMe SSD performance test

We recommend that you use the fio tool to test the performance of a local disk for Linux instances and Windows instances. The following example describes how to test the block storage performance of a bare disk `/ dev / vdx` . For a description of the test command parameters, see performance indexes of block storage.

> ⚠️ Warning:
> A direct bare disk test destroys the file system structure. Therefore, you must back up your data before the test. Additionally, the write operation overwrites the data on local disks. We recommend that you only test local disk performance on newly purchased ECS instances. If you decide to test a bare disk, we recommend that you exercise caution when performing this action.

· NVMe SSD bandwidth performance test

  - To test the read bandwidth, run the following command:

    ```
    fio  - direct = 1  - iodepth = 128  - rw = read  - ioengine =
    libaio  - bs = 128k  - numjobs = 1  - time_based = 1  - runtime =
    1000  - group_repo  rting  - filename =/ dev / vdx
    ```

  - To test the write bandwidth, run the following command:

    ```
    fio  - direct = 1  - iodepth = 128  - rw = write  - ioengine =
    libaio  - bs = 128k  - numjobs = 1  - time_based = 1  - runtime =
    1000  - group_repo  rting  - filename =/ dev / vdx
    ```

· NVMe SSD IOPS performance test

  - To test the random read IOPS, run the following command:

    ```
    fio  - direct = 1  - iodepth = 32  - rw = randread  - ioengine =
    libaio  - bs = 4k  - numjobs = 4  - time_based = 1  - runtime =
    1000  - group_repo  rting  - filename =/ dev / vdx
    ```

  - To test the random write IOPS, run the following command:

    ```
    fio  - direct = 1  - iodepth = 32  - rw = randwrite  - ioengine
    = libaio  - bs = 4k  - numjobs = 4  - time_based = 1  - runtime =
    1000  - group_repo  rting  - filename =/ dev / vdx
    ```

· NVMe SSD latency performance test

  - To test the random read latency, run the following command:

    ```
    fio  - direct = 1  - iodepth = 1  - rw = randread  - ioengine =
    libaio  - bs = 4k  - numjobs = 1  - time_based = 1  - runtime =
    1000  - group_repo  rting  - filename =/ dev / vdx
    ```

  - To test the random write latency, run the following command:

    ```
    fio  - direct = 1  - iodepth = 1  - rw = randwrite  - ioengine =
    libaio  - bs = 4k  - numjobs = 1  - time_based = 1  - runtime =
    1000  - group_repo  rting  - filename =/ dev / vdx
    ```

  - To test the sequential read latency, run the following command:

    ```
    fio  - direct = 1  - iodepth = 1  - rw = read  - ioengine =
    libaio  - bs = 4k  - numjobs = 1  - time_based = 1  - runtime =
    1000  - group_repo  rting  - filename =/ dev / vdx
    ```

  - To test the sequential write latency, run the following command:

    ```
    fio  - direct = 1  - iodepth = 1  - rw = write  - ioengine =
    libaio  - bs = 4k  - numjobs = 1  - time_based = 1  - runtime =
    1000  - group_repo  rting  - filename =/ dev / vdx
    ```

6 Block storage FAQ

What is an ESSD cloud disk?

An Enhanced SSD (ESSD) cloud disk is an ultra-high performance cloud disk provided by Alibaba Cloud. ESSD cloud disks are based on a 25 GE network and remote direct memory access (RDMA) technology. They deliver a maximum random IOPS of 1 million per disk and low one-way latency.

When are ESSD cloud disks available for purchase?

ESSD cloud disks are commercially available from the end of June 2019 after a beta test started on July 14, 2018. You can log on to the ECS console or call CreateDisk to create ESSD cloud disks.

What is the difference of ESSD cloud disks before and after the beta test?

The difference of ESSD cloud disks before and after the beta test lies in their specificat ions. After the beta test, ESSD cloud disks are commercially available in three specifications, as described in the following table.

For ESSD cloud disks, larger capacity means higher data-processing performance. The capacity, IOPS, and throughput provided by ESSD cloud disks vary depending on their performance levels (PLs), as described in the following table.

PL	Capacity (GiB)	Maximum IOPS	Maximum throughput (MBps)
PL1	10-32,768	50,000	350
PL2	461-32,768	100,000	750
PL3	1,261-32,768	1,000,000	4,000

What are the similarities and differences among ESSD, SSD, and Ultra cloud disks?

ESSD, SSD, and Ultra cloud disks have the following similarities and differences:

· Similarities: These three types of cloud disks are based on a distributed block storage architecture. They provide high reliability and scalability and support snapshots and data encryption.

· Differences: Compared with SSD and Ultra cloud disk, ESSD cloud disks provide better performance. For more information, see Block storage performance.

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

The performance level of an ESSD cloud disk positively correlates with its capacity. Specifically, larger capacity delivers higher performance.

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

We recommend that you use the FIO tool to test the performance of ESSD cloud disks. For more information, see How can I test the performance of an ESSD disk.

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

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

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

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

· If the total storage performance of ESSD cloud disks exceeds that of the instance, the storage performance of the instance takes precedence.

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

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

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

Instance type	vCPU (Core)	Memory (GiB)	IOPS (10 thousand)	Throughput (Mbit/s)
ecs.g5se. 6xlarge	24	96	34	1350
ecs.g5se. 8xlarge	32	128	45	1800
ecs.g5se. 16xlarge	64	256	90	3600
ecs.g5se. 18xlarge	70	336	100	4000

How are ESSD cloud disks billed?

ESSD cloud disks support Subscription or Pay-As-You-Go billing methods. For more information, see Prices.

In which regions and zones are ESSD cloud disks available?

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

· China North 2 Zone G

· China North 3 Zone A

· China East 2 Zone E and Zone F

· China East 1 Zone H and Zone I

· China South 1 Zone E

On what instance types can I mount ESSD cloud disks?

Currently, ESSD cloud disks can be mounted to 25 GE network instance type families (c5, ic5, g5, r5, and g5se), ECS Bare Metal Instance type families (ebmhfg5, ebmc4, and ebmg5), and heterogeneous computing instance type families (vgn5i, gn6i, gn6v, gn5 , gn5i, gn4, ga1, f1, and f3).

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

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

· ../DNA0011854887/EN-US_TP_9605.dita#concept_a3f_mg2_wdb

· ../DNA0011854887/EN-US_TP_9604.dita#concept_jl1_qzd_wdb

· Detach a cloud disk

· EN-US_TP_9679.dita#concept_stg_xd3_ydb

- EN-US_TP_9683.dita#concept_vbb_ckj_ydb

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

For detailed information, see Block storage performance.

What scenarios are SSD cloud disks suitable for?

SSD cloud disks apply to the following scenarios:

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

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

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

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

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

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

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

For detailed information, see Prices.

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

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

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

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

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

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

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

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

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

For detailed information, see Block storage performance.

What is Shared Block Storage?

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

What are the benefits of Shared Block Storage?

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

How is Shared Block Storage billed?

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

How can I apply for Shared Block Storage?

To use Shared Block Storage, you must first apply for trial use.

What scenarios are Shared Block Storage suitable for?

Shared Block Storage is designed for the high availability architecture required by enterprise-level applications. It provides shared access to block storage devices in a Shared-everything architecture, such as the high availability server cluster architecture and the Oracle database with Oracle RAC architecture. The Oracle RAC architecture is common among government departments, enterprises, and financial customers.

How can I use Shared Block Storage?

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

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

If you attach Shared Block Storage devices to multiple ECS instances but use a conventional file system to manage them, disk space allocation conflicts and data file inconsistencies may occur:

· Disk space allocation conflict

If a Shared Block Storage device is attached to multiple instances and one of these instances (for example, Instance A) writes data to a file, Instance A checks the file system and available disk space. After the write operation is completed, the space allocation record of Instance A is changed, but the records of other instances are not updated. In this case, when another instance (for example, Instance B) tries to write data to the file, this instance may allocate the disk space that has been already allocated by Instance A, resulting in a disk space allocation conflict.

· Data file inconsistency

After an instance (for example, Instance A) reads and caches data, another process that accesses the same data will directly read the data from the cache. However, if a copy of the same data that is stored on another instance (for example, Instance B) is modified during this period, and Instance A does not update according to the latest data change, Instance A still reads the data from the cache. As a result, data inconsistency occurs.

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

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

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

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

What are the specifications and performance of Shared Block Storage?

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

How can I test the performance of Shared Block Storage?

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

📋 Note:

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

Number of instances involved	Test item	Command
2	Random write IOPS	```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```
	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```
	Random write IOPS	```FIO - direct = 1 - iodepth = 64 - rw = write - ioengine = libaio - bs = 64k - size = 1G - numjobs = 1 - runtime = 1000 - group_repo rting - filename = iotest - name = Write_PPS_ Testing```
	Random read throughput	```FIO - direct = 1 - iodepth = 64 - rw = read - ioengine = libaio - bs = 64k - size = 1G - numjobs = 1 - runtime = 1000 - group_repo rting - filename = iotest - name = Read_PPS_T esting```
4	Random write IOPS	```FIO - direct = 1 - iodepth = 96 - rw = randwrite - ioengine = libaio - bs = 4k - size = 1G - numjobs = 1 - runtime = 1000 - group_repo rting - filename = iotest - name = Rand_Write _Testing```
	Random read IOPS	```FIO - direct = 1 - iodepth = 96 - rw = randread - ioengine = libaio - bs = 4k - size = 1G - numjobs = 1 - runtime = 1000 - group_repo rting - filename = iotest - name = Rand_Read_ Testing```
	Write throughput	```FIO - direct = 1 - iodepth = 64 - rw = write - ioengine = libaio - bs = 64k - size = 1G - numjobs = 1```

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

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

What is an independent cloud disk?

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

Can I mount one cloud disk to multiple ECS instances?

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

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

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

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

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

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

Yes.

Can I detach a data disk from a Subscription instance?

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

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

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

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

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

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

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

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

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

Why are some automatic snapshots on my cloud disk missing?

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

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

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

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

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

- In Windows, we recommend that you stop the read and write operations on all file systems of the cloud disk to ensure data integrity. Otherwise, the data being read or written will be lost.

- In Linux, you must log on to the ECS instance and run the `umount` command. After the command is executed successfully, log on to the ECS console to unmount the disk.

Can I unmount the system disk?

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

What is a mount point?

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

Are my snapshots retained if I reinitialize a cloud disk?

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

Are my snapshots retained if I replace the system disk?

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

> 📋 Note:
>
> The disk ID is changed when you replace the system disk. This means that the snapshots of the original system disk cannot be used for the new system disk. However, you can use the manual snapshots preserved on the original system disk to create a custom image.

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

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

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

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

1. Locate the disk where the data is stored and check whether the data disk is mounted to the corresponding ECS instance by using either of the following methods:

   · Perform the check in the ECS console. For more information, see Monitor a cloud disk.

   · Log on to the instance, run the `fdisk - l` command to check whether the data disk partition information is correct, and run the `df - h` and `mount | grep "< devpath >"` commands to view the mount information.

2. Run the `cat` command to view the `/ etc / fstab` file and check whether you mount two cloud disks to the same directory.

   · If two cloud disks are mounted to the same directory, the previously mounted cloud disk will be replaced by the new one, resulting in data inaccessibility. We recommend that you mount one of the cloud disks to a different directory.

   · If they are mounted to different directories but the mount information shows that they are in the same directory, run the `ll` command to check whether there are links between the two directories. If there is a link, run the `mkdir` command to create a new directory, mount one of the cloud disks to this directory, and check whether data can be accessed.

How can I resize a system disk?

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

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

Before replacing a system disk, you must:

· Create snapshots and custom images of the current system disk.

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

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

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

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

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

What is the storage capacity range of a system disk?

The capacity of a system disk varies depending on the operating system. For more information, see Overview.

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

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

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

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

Can I partition a data disk for data storage?

Yes, you can divide a data disk into multiple partitions according to your specific needs. We recommend that you use the system tool for partitioning.

For a disk with multiple partitions, are snapshots created for the entire disk or only for a specific partition?

Snapshots are created for the entire disk, rather than for a specific partition.

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

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

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

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

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

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

1. Create a snapshot of the current data disk.

2. Go to the disk purchase page to purchase a Pay-As-You-Go cloud disk, select the same region and zone as the ECS instance, and click Create from snapshot.

3. Log on to the ECS console and then mount the data disk you purchased to an ECS instance.

4. Log on to the ECS instance, run the `mount` command to mount the purchased disk to the ECS instance, and then check whether files in this disk are the same as those in the original data disk.

5. Run the `fdisk` command to delete the original partition table, and then run commands such as `fdisk` and `mkfs . ext3` to format and partition the original data disk again, so that the available space of the original data disk is the same as that after the scale-out action.

6. Run the `cp - R` command to copy all the data in the purchased disk back to the original data disk. You can add the `-- preserve = all` parameter to preserve the file's properties when copying files.

7. Run the `umount` command to unmount the purchased data disk.

8. Log on to the ECS console, unmount the purchased data disk from the ECS instance , and release it.

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

When you add a data disk to the Linux system and add the partition information according to Format a data disk of a Linux instance, note that the mount point of a basic cloud disk is `/ dev / xvdb1` and that of an SSD cloud disk is `/ dev / Vdb1` . If you add invalid information, disk mounting will fail. To avoid this issue, follow these steps:

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

2. Check whether the information added to `/ etc / fstab` is valid. Do not add mount information repeatedly because this will cause a system startup failure.

3. Run the `vim` command to modify the `/ etc / fstab` file.

4. Comment out or delete invalid information, add valid mount information, and then run the `mount  - a` command to check whether the data disk is successfully mounted.

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

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

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

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

   The command output does not contain information related to `/ dev / vdb1` .

2. Run the `fdisk  - l` command to view partitioned data disks.

3. Run the `mkdir  / InitTest` command to re-create the mount point of the data disk partition. The mount point name must be the same as that of the mount point `/ dev / vdb1` before the system disk initialization. You can view the original mount point name by running the `cat  / etc / fstab` command.

4. Run the `mount  / dev / vdb1  / InitTest` command to remount the data disk partition.

5. Run the `df  - h` command to view the mount results.

6. Check whether `/ dev / vdb1` can be mounted automatically.

   a. Run the `umount  / dev / vdb1` command to unmount `/ dev / vdb1` .

   b. Run the `mount` command to check whether the partition is successfully unmounted. If this is the case, the returned result does not contain information related to `/ dev / vdb1` .

   c. Run the `mount  - a` command to automatically mount `/ dev / vdb1` .

   d. Run the `mount` command to check whether the auto-mount succeeded. If this is the case, the returned result contains information related to `/ dev / vdb1` .

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

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

Cause:

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

Solution:

1. Run the `mount  / dev / xvdb1` command to mount the data disk.

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

   In this example, the file format is ext3.

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

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

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

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

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

Solution:

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

1. Run the `fdisk  - l` command.

   The output shows that one disk is not mounted.

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

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

What can I do if the error message "Bad magic number in super-block while trying to open /dev/ xvdb1" is returned when I resize the cloud disk of a Linux instance?

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

Cause: There is no disk partition

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

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

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

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

Solution:

· Check whether the operating system of instances associated with related snapshots has been replaced. For more information, see EN-US_TP_9683.dita#concept_vbb_ckj_ydb.

  If the operating system has been replaced, the system disk of the instance will be automatically re-created from the new image and the system disk ID will change. Therefore, original snapshots cannot be used for the rollback. However, you can ../ DNECS19100339/EN-US_TP_9696.dita#concept_gpg_t5l_xdb from related snapshots and then specify the custom image by replacing the system disk (non-public image) to switch the instance to the corresponding snapshot state.

· Check whether the instance has stopped.

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

· Check whether the disk has an ongoing snapshot task.

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

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

How can I copy data across instances?

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

· Data copy between Linux instances

  - Using the lrzsz tool

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

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

  - Using the FTP tool

    If you use the SFTP tool, we recommend that you use the root account for instance logon and file upload and download.

  - Using the `wget` command

    After compressing a file or a folder, save it to the `web` directory to generate a download URL. Then run the `wget` command on another instance to download the file or folder.

· Data copy between Linux and Windows

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

· Data copy between Windows

  - FTP tool

  - Alibaba TradeManager

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

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

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

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

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

Local SSD disks are no longer available for purchase. If you are still using a local SSD disk, see EN-US_TP_9561.dita#concept_g3w_qzv_tdb and FAQ about SSD cloud disks.

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

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

> ⚠️ **Warning:**
> You can obtain accurate test results by testing a bare disk partition, but you may destroy the file system structure by testing it directly. Before such a test, you must create a snapshot of the disk to back up your data. We recommend that you use newly purchased ECS instances that contain no data for the test to prevent any data loss.

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

We also recommend that you use the FIO for the performance test.

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

1. Connect to the Linux instance.
2. Run the following commands to install the libaio and FIO tools.

```
sudo  yum  install  libaio  - y
```

```
sudo   yum   install   libaio – devel  – y
sudo   yum   install   fio  – y
```

3. Run the `cd  / tmp` command to change the directory.

4. Run the `vim    test100w . sh` command to create a script *test100w . sh*

   and paste the following code to the script to test the random write `randwrite`

   IOPS.

```
function   RunFio
{
  numjobs =$ 1   #  The   number   of   the   tested   threads ,
such   as   8   in   this   example .
  iodepth =$ 2   #  The   maximum   concurrent   I / O   requests ,
such   as   64   in   this   example .
  bs =$ 3        #  The   size   of   the   data   block   for   one
  I / O ,  such   as   4k   in   this   example .
  rw =$ 4        #  The   read   and   write   policy ,  such   as
randwrite   in   this   example .
  filename =$ 5 #  The   name   of   the   tested   file ,  such
as  / dev / vdb   in   this   example .
 nr_cpus =` cat  / proc / cpuinfo  | grep  " processor " | wc  – l
`
  if  [ $ nr_cpus  – lt  $ numjobs  ]; then
      echo  " Numjobs   is   more   than   cpu   cores ,  exit !"
      exit  – 1
  fi
  let   nu =$ numjobs + 1
  cpulist =""
  for  (( i = 1 ; i < 10 ; i ++))
  do
      list =` cat  / sys / block / vdb / mq /*/ cpu_list  |  awk  '{
if ( i <= NF )  print  $ i ;}'  i ="$ i " |  tr  – d  ',' |  tr  '\
n ' ','`
      if  [ – z  $ list  ]; then
         break
      fi
      cpulist =${ cpulist }${ list }
  done
  spincpu =` echo  $ cpulist  |  cut  – d  ',' – f   2 -${ nu }`
  echo  $ spincpu
  fio  -- ioengine = libaio  -- runtime = 30s  -- numjobs =${
numjobs } -- iodepth =${ iodepth } -- bs =${ bs } -- rw =${ rw }
-- filename =${ filename } -- time_based = 1  -- direct = 1  --
name = test  -- group_repo  rting  -- cpus_allow  ed =$ spincpu  --
cpus_allow  ed_policy = split
}
echo   2  > / sys / block / vdb / queue / rq_affinit  y
sleep   5
RunFio   8   64   4k   randwrite  / dev / vdb
```

📋 **Note:**

- You must modify the following commands according to your test environment.

  - `vdb` in the command line `list =` cat / sys / block / vdb / mq /` `*/ cpu_list | awk '{ if (i <= NF) print $ i ;}' i ="$ i "` `| tr - d ',' | tr '\ n ' ','`

  - `8` , `64` , `4k` , `randwrite` , and / dev / vdb in the command line

    `RunFio  8   64   4k   randwrite  / dev / vdb`

- The file system structure may be damaged if you test the bare disk partition directly. If you choose to proceed with this action, set `filename` as the device name, such as / dev / vdb . If you do not want to risk data loss, set `filename` as a file path, such as / mnt / test . image .

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

```
[root@             ~]# sh test100w.sh
4,8,12,16,20,24,28,1
test: (g=0): rw=randwrite, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B, ioengine=libaio, iodepth=64
...
fio-3.0
Starting 8 processes
Jobs: 8 (f=8): [w(8)][100.0%][r=0KiB/s,w=3965MiB/s][r=0,w=1015k IOPS][eta 00m:00s]
test: (groupid=0, jobs=8): err= 0: pid=9565: Thu Jan 11 17:37:00 2018
  write: IOPS=1006k, BW=3931MiB/s (4122MB/s)(115GiB/30047msec)
    slat (usec): min=2, max=52179, avg= 4.22, stdev= 9.78
    clat (usec): min=82, max=624296, avg=500.05, stdev=814.37
     lat (usec): min=86, max=624300, avg=504.46, stdev=814.43
    clat percentiles (usec):
     |  1.00th=[   174],  5.00th=[   210], 10.00th=[   235], 20.00th=[   273],
     | 30.00th=[   306], 40.00th=[   338], 50.00th=[   383], 60.00th=[   441],
     | 70.00th=[   523], 80.00th=[   660], 90.00th=[   930], 95.00th=[ 1139],
     | 99.00th=[ 1582], 99.50th=[ 1909], 99.90th=[ 3458], 99.95th=[ 4424],
     | 99.99th=[17957]
   bw (  KiB/s): min=256936, max=526637, per=12.53%, avg=504445.98, stdev=33299.37, samples=480
```

Script details

- Block device parameter

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

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

  - If the value of the `rq_affinit  y` parameter is `2` , the I/O Completions are forced to run on the requesting vCPU. As a result, performance of each vCPU is fully used in the case of concurrent processing I/O requests by multiple threads.

- Binding threads to corresponding vCPU cores

  - Generally, a device has only one Request-Queue. This unique Request-Queue
    is a bottleneck in the case of concurrent processing I/O requests by multiple
    threads.

  - In the latest Multi-Queue mode, one device can have multiple Request-Queues
    that process I/O requests, which deliver full performance of the back-end
    storage. If you have four I/O threads, you must bind them to the vCPU core
    corresponding to each Request-Queue. This allows you to fully use the Multi-
    Queue mode to improve the storage performance.

To fully output the performance of the device, you must assign the I/O requests to
different Request-Queues. The following command in the script can be used to bind
`jobs` to different vCPU cores. `/ dev / vd *` is the device name of your ESSD cloud
disk, for example, `/ dev / vdb` .

```
fio — ioengine = libaio  — runtime = 30s  — numjobs =${ numjobs
} — iodepth =${ iodepth } — bs =${ bs } — rw =${ rw } — filename
=${ filename } — time_based = 1  — direct = 1  — name = test
— group_repo  rting  — cpus_allow  ed =$ spincpu  — cpus_allow
ed_policy = split
```

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

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

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