Alibaba Cloud

Elastic Compute Service Block Storage

Document Version: 20220712

(-) Alibaba Cloud

Legal disclaimer

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

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized 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 because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should 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 this document based on the "status quo", "being defective", and "existing functions" of its products and services. 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 take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, 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 contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, 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 this document 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 directly contact Alibaba Cloud for any errors of this document.

> Document Version: 20220712

Document conventions

Style	Description	Example	
<u> Danger</u>	A danger notice 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.	
<u> </u>	A warning notice 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 restart an instance.	
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.	
? Note	A note indicates supplemental instructions, best practices, tips, and other content.	? Note: You can use Ctrl + A to select all files.	
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type.	
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click OK.	
Courier font is used for commands		Run the cd /d C:/window command to enter the Windows system folder.	
Italic formatting is used for parameters and variables.		bae log listinstanceid Instance_ID	
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]	
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}	

Table of Contents

1.	Block Storage overview	07
	1.1. Elastic Block Storage devices	07
	1.2. Disks	11
	1.3. ESSDs	14
	1.4. Local disks	17
	1.5. Triplicate storage	25
2.	Cloud disks	27
	2.1. Create a cloud disk	27
	2.1.1. Create a disk	27
	2.1.2. Create a disk from a snapshot	30
	2.2. Attach a data disk	34
	2.3. Format a data disk	36
	2.3.1. Partition and format a data disk on a Linux instance	37
	2.3.2. Partition and format a data disk on a Windows instan	44
	2.3.3. Partition and format a data disk larger than 2 TiB in	48
	2.3.4. Create a file system on a raw data disk	55
	2.4. Replace the operating systems of instances	56
	2.4.1. Replace the operating system of an instance by using	56
	2.4.2. Replace the operating system of an instance by using	60
	2.5. Reinitialize a cloud disk	61
	2.5.1. Re-initialize a system disk	61
	2.5.2. Re-initialize a data disk	64
	2.6. Roll back a disk by using a snapshot	67
	2.7. Change the category of a disk	68
	2.8. Modify the performance levels of ESSDs	71
	2.9. View the monitoring data of a disk	72

2.10. Query the serial number of a disk	75
2.11. Change the billing methods of a disk	77
2.12. Detach or attach a system disk	78
2.13. Edit the tags of a disk	81
2.14. Detach a data disk	82
2.15. Release a disk	86
3.Encrypt disks	90
3.1. Encrypt a system disk	90
3.2. Encrypt a data disk	95
4.Resize cloud disks	- 98
4.1. Overview	98
4.2. Resize disks online for Linux instances	99
4.3. Resize disks offline for Linux instances	108
4.4. Resize disks online for Windows instances	109
4.5. Resize disks offline for Windows instances	115
4.6. Update Red Hat virtio drivers of Windows instances	121
5.Support for the NVMe protocol and multi-attach feature	124
5.1. Overview of disks that support NVMe	124
5.2. Enable multi-attach	126
6.Performance	132
6.1. EBS performance	132
6.2. Storage I/O performance	134
6.3. Test the performance of EBS devices	143
6.4. Test the IOPS performance of an ESSD	147
7. Dedicated block storage clusters	151
7.1. Overview	151
7.2. Billing	153
7.3. Create a dedicated block storage cluster	154

7.4. View the information of a dedicated block storage cluster	156
7.5. Create disks in a dedicated block storage cluster	157
8.Async replication	161
8.1. Overview	161
8.2. Create a replication pair	165
8.3. Enable async replication for replication pairs	166
8.4. Implement disaster recovery	167
8.5. Delete a replication pair	170
9.Replication pair-consistent groups	172
9.1. Overview	172
9.2. Create a replication pair-consistent group	175
9.3. Add replication pairs	176
9.4. Remove replication pairs	177
9.5. Activate async replication for replication pair-consistent gr	177
9.6. Disable the async replication feature for replication pair-c	179
9.7. Use replication pair-consistent groups to implement disaste.	180
9.8. Delete a replication pair-consistent group	182
10.Storage capacity units	184
10.1. Overview	184
10.2. Usage rules	186
10.3. Create an SCU	188
11.FAO	190

1.Block Storage overview1.1. Elastic Block Storage devices

Elastic Block Storage (EBS) is a high-performance, low-latency block storage service provided by Alibaba Cloud for Elastic Compute Service (ECS). EBS supports random read and write to meet the data storage requirements of your business. EBS devices can be used as physical hard disks. You can partition and format EBS devices and create file systems for them.

EBS devices

Alibaba Cloud provides a variety of EBS devices for ECS, such as cloud disks based on a distributed storage architecture and local disks located on the same physical machine that hosts the ECS instance.

EBS device		Description	References
Cloud disk	Enhanced SSD (ESSD)	ESSDs utilize a next-generation distributed block storage architecture and provide 25 Gibabit Ethernet and remote direct memory access (RDMA). Each ESSD has low latency and can deliver up to 1,000,000 random read/write IOPS. We recommend that you use ESSDs for scenarios such as online transactional processing (OLTP) databases, NoSQL databases, and Elasticsearch, Logstash, and Kibana (ELK) distributed logs.	ESSDs
	Standard SSD	Standard SSDs are high-performance disks that offer consistent high random IOPS and high data reliability. We recommend that you use standard SSDs for scenarios such as I/O-intensive applications, small and medium-sized relational databases, and NoSQL databases.	Disks
	Ultra disk	Ultra disks are cost-effective and offer medium random IOPS and high data reliability. We recommend that you use ultra disks as system disks or for scenarios such as development and testing.	Disks
	Basic disk	Basic disks are the previous generation of disks and are unavailable for purchase.	None

EBS device	Description	References
	Local disks are physical disks attached to physical machines that host ECS instances. Local disks provide local storage for ECS instances. They are suitable for business scenarios that require high storage I/O performance, a large amount of storage, and high cost-effectiveness. Local disks provide low latency, high random IOPS and throughput, and high cost performance.	
Local disk	Note All of the local disks for an instance reside on a single physical machine. This increases the risks of single points of failure (SPOFs). The durability of data stored in a local disk is determined by the reliability of the associated physical machine. To ensure data availability, we recommend that you implement data redundancy at the application layer. You can use deployment sets to distribute ECS instances across multiple physical machines for high availability and disaster recovery.	Local disks

Service performance

Different EBS devices deliver different performance.

- For more information about the performance of cloud disks, see EBS performance.
- For more information about the performance of local disks, see Local disks.

Data security

Note The following features except for the data erasure mechanism are unique to cloud disks.

• Data durability during read and write operations

Three copies of your business data are stored in the block storage cluster within the same zone to ensure 99.999999% (nine 9's) data durability during read and write operations. For more information, see Triplicate storage.

Proactive backup

You can create snapshots for disks on a regular basis to enhance data security. Snapshots are a backup service provided by Alibaba Cloud. They back up the data stored on cloud disks, such as logs and customer transactions, for query purposes. For more information, see Snapshot overview.

Dat a erasure mechanism

When you delete a piece of data, the data is completely erased from the distributed block storage system and can no longer be accessed by other users. The following mechanisms are used to ensure that deleted data is completely erased:

- The storage system performs sequential writes to append data to existing files at the underlying layer of cloud disks. This mechanism makes full use of high-bandwidth and low-latency sequential writes to physical disks. If you delete a logical space from a cloud disk after data is appended at the underlying layer of the cloud disk, this operation is recorded as metadata. The storage system returns only zero for all requests of reading data from the logical space. Similarly, when you overwrite the data in a logical space of a cloud disk, the storage system does not immediately overwrite the data in the corresponding physical space but modifies the mapping between the logical space and the physical space. This ensures that data that has been overwritten is no longer readable. Data fragments that result from delete or overwrite operations are forcibly and permanently deleted from the underlying physical disks.
- When an EBS device (cloud disk) is released, the storage system immediately destroys the
 metadata of the device to ensure that the data is no longer accessible. At the same time, the
 physical storage space that corresponds to the disk is recycled. The physical storage space must
 be cleared before it is re-assigned to store data. Before data is written to a new cloud disk, the
 system returns only zero for all read requests.

• Dat a encryption

We recommend that you encrypt the storage devices that you use if your applications are data-sensitive. Cloud disks and their snapshots are encrypted by using keys based on the standard AES-256 algorithm. Data is automatically encrypted when it is transmitted from ECS instances to cloud disks. Encrypted data is automatically decrypted when it is read. For more information, see 加密概述.

Billing

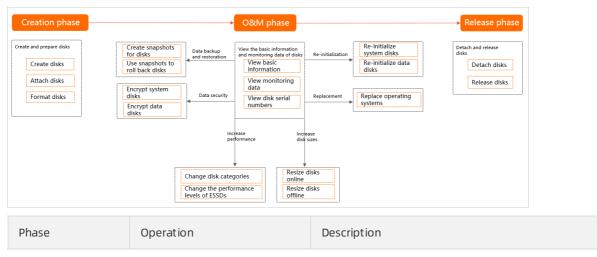
For information about the billing methods and billable items of EBS, see Overview and EBS devices.

Limits

For information about the limits and quotas of EBS devices, see the "Elastic Block Storage (EBS) limits" in Limits.

Operations

The following figure shows common operations that can be performed on a cloud disk throughout its lifecycle.



Phase	Operation	Description		
Creation phase	Create an empty cloud disk or a cloud disk that contains backup data for an ECS instance.	 Create an empty cloud disk. Create a disk Attach a data disk Initialize the data disk and create file systems for it. Partition and format a data disk on a Linux instance Partition and format a data disk on a Windows instance Create a cloud disk that contains backup data. Create a disk from a snapshot Attach a data disk 		
	View the basic information and monitoring data of cloud disks.	 Query the serial number of a disk View the monitoring data of a disk		
	Use the snapshot feature to back up and restore data stored in cloud disks.	Create a snapshot of a diskRoll back a disk by using a snapshot		
	Upgrade the specifications of cloud disks that do not meet your performance requirements.	Change the category of a diskModify the performance levels of ESSDs		
	Resize the storage capacity of cloud disks that do not meet your performance requirements.	 Resize disks online for Linux instances Resize disks offline for Linux instances Resize disks online for Windows instances Resize disks offline for Windows instances 		
O&M phase	If you want to restore cloud disks to their initial states, you can initialize the disks.	Re-initialize a system diskRe-initialize a data disk		
	If you want to use a new operating system, you can change the operating system of an instance by replacing the system disk.	 Replace the operating system of an instance by using a public image Replace the operating system of an instance by using a non-public image 		

Phase	Operation	Description
Release phase	If cloud disks are no longer needed, you can release them to minimize costs.	 Release a pay-as-you-go disk. i. Detach a data disk ii. Release a disk Release a subscription disk. i. Change the billing method to pay-as-you-go: Change the billing methods of a disk ii. Detach a data disk iii. Release a disk

Local disks can be created and released only along with instances of specific instance types, and can be used as data disks and formatted. For more information, see the following topics:

- Partition and format a data disk on a Linux instance
- Partition and format a data disk on a Windows instance

Other Alibaba Cloud storage services

In addition to EBS, Object Storage Service (OSS) and Apsara File Storage NAS are provided by Alibaba Cloud. For more information, see What is OSS? and What is NAS?.

1.2. Disks

Disk categories

Disks are classified into the following categories based on their performance:

- Enhanced SSDs (ESSDs): ESSDs are based on the next-generation distributed block storage
 architecture and utilize the 25 Gigabit Ethernet and remote direct memory access (RDMA)
 technologies. Each ESSD has low latency and can deliver up to 1,000,000 random read/write IOPS. For
 more information, see ESSDs.
 - We recommend that you use ESSDs for scenarios such as online transactional processing (OLTP) databases, NoSQL databases, and Elasticsearch, Logstash, and Kibana (ELK) distributed logs.
- Standard SSDs: Standard SSDs are high-performance disks that provide consistent high random IOPS and high data reliability.
 - We recommend that you use standard SSDs for scenarios such as I/O-intensive applications, small and medium-sized relational databases, and NoSQL databases.
- Ultra disks: Ultra disks are cost-effective and provide medium random IOPS and high data reliability. We recommend that you use ultra disks as system disks or for scenarios such as development and testing.
- Basic disks: Basic disks are the previous generation of disks and are unavailable for purchase.

The following table describes the performance of disks of different categories.

Category	ESSD			Standard	Ultra disk	Basic disk	
caregory	PL3	PL2	PL1	PL0	SSD	Ollia uisk	basic disk
Capacity range per disk (GiB)	1261~327 68	461~3276 8	20~32768	40~32768	20~32768	20~32768	5~2000
Maximum IOPS	1000000	100000	50000	10000	25000	5000	Several hundreds
Maximum throughp ut (MB/s)	4000	750	350	180	300	140	30~40
Formula for calculatin g the IOPS per disk	min{1800 + 50 × Capacity, 1000000}	min{1800 + 50 × Capacity, 100000}	min{1800 + 50 × Capacity, 50000}	min{1800 + 12 × Capacity, 10000}	min{1800 + 30 × Capacity, 25000}	min{1800 + 8 × Capacity, 5000}	N/A
Formula for calculatin g the throughp ut per disk (MB/s)	min{120 + 0.5 × Capacity, 4000}	min{120 + 0.5 × Capacity, 750}	min{120 + 0.5 × Capacity, 350}	min{100 + 0.25 × Capacity, 180}	min{120 + 0.5 × Capacity, 300}	min{100 + 0.15 × Capacity, 140}	N/A
Average single- channel random write latency in millisecon ds (block size = 4 K)	0.2		0.3~0.5	0.5~2	1~3	5~10	
API paramete r value	cloud_essd				cloud_ssd	cloud_effi ciency	cloud

For more information about disk performance, see EBS performance.

Disks are classified into system disks and data disks based on their purposes.

- System disks contain operating systems and can be created only along with instances. The lifecycle of a system disk is the same as that of the instance to which the system disk is attached.
- Data disks are used to store application data and can be created separately or along with instances.

Note When a disk is created, the disk capacity displayed in the ECS console includes the capacity occupied by the operating system, and the remaining available capacity may be less than the disk capacity displayed in the ECS console. For example, if the capacity of a system disk displayed in the ECS console is 40 GiB, the remaining available capacity may be less than 40 GiB because the operating system occupies some of the disk capacity.

Limits

A disk can be attached only to a single ECS instance within the same zone.

The following limits also apply.

Billing methods

For more information, see the Pricing tab on the Elastic Compute Service product page.

Disks support the subscription and pay-as-you-go billing methods. For more information, see Subscription and Pay-as-you-go.

- Disks created along with subscription instances or separately created for subscription instances use the subscription billing method.
- Disks created along with pay-as-you-go instances or separately created use the pay-as-you-go billing method. Storage capacity units (SCUs) can be used to offset the bills of pay-as-you-go disks.

After a disk is created, you can change its billing method. For more information, see Change the billing methods of a disk.

Related operations

The following table describes the operations that you can perform on disks.

Operation	References
Attach an idle pay-as-you-go disk to an ECS instance	 Attach a data disk Format the data disk based on the operating system: Partition and format a data disk larger than 2 TiB in size Partition and format a data disk on a Linux instance Partition and format a data disk on a Windows instance
Create and use a data disk	 Create a disk Attach a data disk Format the data disk based on the operating system: Partition and format a data disk on a Linux instance Partition and format a data disk on a Windows instance

Operation	References
Encrypt data stored on a disk	For more information, see 加密概述. For more information, see the following topics: Encrypt a system disk Encrypt a data disk
Resize a system disk or data disk	For more information, see Overview. For more information, see the following topics: Resize disks online for Linux instances Resize disks online for Windows instances
Replace the operating system of an instance	Replace the operating system of an instance by using a public image
Back up data stored on a disk	Create a snapshot of a diskApply or disable an automatic snapshot policy
Restore a disk to its initial state	Re-initialize a system disk
Restore a disk to its state at a specific point of time	Roll back a disk by using a snapshot
Detach a damaged system disk from an instance and attach the disk back to the instance after the disk is repaired	Detach or attach a system disk
Release an instance but retain data stored on its system disk	 Detach or attach a system disk Enable or disable the Release Disk with Instance feature for disks when you create an instance Enable or disable the Release Disk with Instance feature for a disk on the Disks page
Release a subscription disk that is no longer needed to minimize costs	 Change the billing methods of a disk Detach a data disk Release a disk
Release a pay-as-you-go disk that is no longer needed to minimize costs	 Detach a data disk Release a disk

1.3. ESSDs

Alibaba Cloud enhanced SSDs (ESSDs) use 25 Gigabit Ethernet and remote direct memory access (RDMA) to reduce latency and deliver up to 1,000,000 random read/write IOPS per disk. This topic provides general information about ESSDs, such as performance levels (PLs), use scenarios, and performance specifications.

Specifications

The API parameter value *cloud_essd* indicates ESSDs. ESSDs are divided into four PLs based on the maximum performance per disk.

FCCD attribute	PL				
ESSD attribute	PL3	PL2	PL1	PLO	
Performance description	Ultra-high maximum concurrent I/O performance and ultra-low I/O latency	High maximum concurrent I/O performance and low I/O latency	Moderate maximum concurrent I/O performance and low I/O latency	Moderate maximum concurrent I/O performance and low I/O latency	
Capacity range (GiB)	1,261~32,768	461~32,768	20~32,768	40~32,768	
Data durability	99.999999%	99.999999%	99.999999%	99.999999%	
Maximum IOPS per disk	1,000,000	100,000	50,000	10,000	
Maximum throughput per disk (MB/s)	4,000	750	350	180	
Formula for calculating the IOPS per disk	min{1800 + 50 × Capacity, 1,000,000}	min{1800 + 50 × Capacity, 100,000}	min{1800 + 50 × Capacity, 50,000}	min{1,800 + 12 × Capacity, 10,000}	
Formula for calculating the throughput per disk (MB/s)	min{120 + 0.5 × Capacity, 4,000}	min{120 + 0.5 × Capacity, 750}	min{120 + 0.5 × Capacity, 350}	min{100 + 0.25 × Capacity, 180}	
Example scenario	Large and medium-sized relational databases for core business and NoSQL databases, and large SAP and Oracle databases	Medium-sized relational databases and NoSQL databases, medium-sized Elasticsearch, Logstash, and Kibana (ELK) log clusters, and enterprise-level commercial software such as SAP and Oracle	Small and medium-sized MySQL and SQL Server databases, small and medium-sized ELK log clusters, enterprise-level commercial software such as SAP and Oracle, and container applications	Small and medium-sized MySQL and SQL Server databases, small and medium-sized ELK log clusters, enterprise-level commercial software such as SAP and Oracle, and container applications	

ESSD attribute	PL				
ESSD attribute	PL3	PL2	PL1	PLO	
Recommended system or data disks to be replaced with ESSDs in recommended business scenarios	Data disks of instance families that are equipped with local SSDs and have 16 or more vCPUs (i1, i2, and i2g)	Standard SSDs and data disks of instance families that are equipped with local SSDs (i1, i2, and i2g)	Standard SSDs	System disks	

Note For information about how to test the IOPS performance of an ESSD, see Test the IOPS performance of an ESSD.

Billing methods

ESSDs support the pay-as-you-go and subscription billing methods.

For information about the pricing of ESSDs at different PLs, see the Pricing tab on the Elastic Compute Service page.

Scenarios

ESSDs are suitable for the following latency-sensitive applications or I/O intensive business scenarios:

- Large online transaction processing (OLTP) databases: relational databases such as MySQL, PostgreSQL, Oracle, and SQL Server databases
- NoSQL databases: non-relational databases such as MongoDB, HBase, and Cassandra databases
- Elasticsearch distributed logs: ELK log analysis

Capacity and PLs

The performance of a storage device is closely related to the capacity of the device. A storage device that has a larger capacity provides higher data processing capabilities. All ESSDs have the same I/O performance per unit of capacity. However, the performance of ESSDs increases linearly with its capacity until the maximum performance per disk at the PL is reached.

PL	ESSD capacity range (GiB)	Maximum IOPS	Maximum throughput (MB/s)
PLO	40~32,768	10,000	180
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 Alex selects 20 GiB of storage capacity when Alex creates an ESSD in the ECS console, PL1 is the only available option. The PL1 ESSD delivers a maximum of 50,000 IOPS.



• Example 2: If Alex selects 32,000 GiB of storage capacity when Alex creates an ESSD in the ECS console, all PLs are available. A maximum of 10,000 IOPS is delivered for PL0, 50,000 IOPS for PL1, 100,000 IOPS for PL2, and 1,000,000 IOPS for PL3.



Storage I/O performance of instance types

The storage I/O performance of some new-generation instance families is proportional to the specifications of instance types. For example, in the g7se storage enhanced instance family, the higher the specifications of an instance type, the higher storage IOPS and throughput the instance can deliver. For more information, see Storage I/O performance.

The following section describes the relationship between instance types and ESSDs in terms of performance:

- If the total performance of ESSDs of an instance does not exceed the storage I/O performance of the instance type, the total performance of the ESSDs prevails.
- If the total performance of ESSDs of an instance exceeds the storage I/O performance of the instance type, the storage I/O performance of the instance type prevails.

The actual storage performance differs when instances of different instance types are attached with ESSDs at different PLs. In the following examples, g7se is used:

- Example 1: Alex creates a 16 GiB instance of the ecs.g7se.xlarge instance type. The maximum IOPS that this instance type can deliver is 60,000. Then, Alex attaches a PL2 ESSD to the instance. The ESSD has a capacity of 2,000 GiB and a maximum of 100,000 IOPS. The maximum IOPS of the instance is limited by the maximum IOPS of the instance type to 60,000.
- Example 2: Alex creates a 64 GiB instance of the ecs.g7se.4xlarge instance type. The maximum IOPS that this instance type can deliver is 150,000. Then, Alex attaches three PL2 ESSDs to the instance. Each ESSD has a capacity of 2,000 GiB and a maximum of 100,000 IOPS. The total maximum IOPS that these ESSDs can deliver is 300,000. The maximum IOPS of the instance is limited by the maximum IOPS of the instance type to 150,000.
- Example 3: Alex creates a 64 GiB instance of the ecs.g7se.4xlarge instance type. The maximum IOPS that this instance type can deliver is 150,000. Then, Alex attaches a PL3 ESSD to the instance. The ESSD has a capacity of 2,000 GiB and a maximum of 101,800 IOPS. The maximum IOPS of the instance is not limited by the maximum IOPS of the instance type, but is limited by the maximum IOPS of the ESSD to 101,800.

Instance families supported by ESSDs

For information about instance families supported by ESSDs at PLO, PL1, PL2, and PL3, see Instance family.

1.4. Local disks

Local disks are located on the physical server that hosts their associated Elastic Compute Service (ECS) instance. Local disks provide local storage for ECS instances. Local disks are cost-effective and offer high random IOPS, high throughput, and low latency.

Limits

• Local disks are physically located on a single physical server. This increases the risks of single points of failure (SPOFs). The durability of data stored on local disks is determined by the reliability of the associated physical server.

• Warning For example, data stored on local disks may be lost when a hardware failure occurs. We recommend that you make sure to store only temporary data on local disks.

- To ensure data availability, we recommend that you implement data redundancy at the
 application layer. You can use deployment sets to distribute ECS instances across multiple
 physical servers for high availability and disaster recovery. For more information, see
 Create a deployment set.
- If your applications do not utilize an architecture that prioritizes data reliability, we recommend that you use cloud disks or a backup service with ECS instances to improve data reliability. For more information, see <u>Disks</u> or <u>What is Hybrid Backup Recovery?</u>.
- After you purchase an ECS instance that has local disks attached, you must log on to the instance to partition and format the local disks. For more information, see Partition and format a data disk on a Linux instance or Partition and format a data disk on a Windows instance.
- Local disks do not support the following operations:
 - o Create a separate local disk.
 - Use a snapshot to create a local disk.
 - o Attach a local disk.
 - Detach and release a local disk.
 - Resize a local disk.
 - o Re-initialize a local disk.
 - o Create a snapshot for a local disk.
 - Use a snapshot to roll back a local disk.

Disk categories

Note This topic provides information about local disks that are purchased together with ECS instances. For more information about the performance of instance families that are equipped with local SSDs and big data instance families, see Instance families.

Local disks are suited for business scenarios that require high storage I/O performance, mass storage, and high cost efficiency. Alibaba Cloud provides two categories of local disks. The following table describes the categories.

Category	Instance family	Use scenario	
----------	-----------------	--------------	--

Category	Instance family	Use scenario
Local NVMe SSD	 The following instance families use local NVMe SSDs: Instance families equipped with local SSDs: i3, i3g, i2, i2g, i2ne, i2gne, and i1 GPU-accelerated compute-optimized instance family: gn5 	Instance families equipped with local NVMe SSDs are suited for the following scenarios: I/O-intensive applications that require high I/O performance and low latency, such as online gaming, e-commerce, live streaming, and media Applications that require high storage I/O performance and a high-availability architecture at the application layer, such as NoSQL databases (including Cassandra, MongoDB, and HBase), MPP data warehouses, and distributed file systems
Local SATA HDD	The d1ne and d1 instance families use local SATA HDDs.	Local SATA HDDs are the preferred storage media for industries such as Internet and finance that have high requirements for big data computing, storage, and analysis. These disks are suited for mass storage and offline computing scenarios and can meet the high requirements of distributed computing services such as Hadoop in terms of storage performance, storage capacity, and internal bandwidth.

Local NVMe SSDs

? Note You can test the bandwidth, IOPS, and latency of local NVMe SSDs to obtain the standard performance data and measure the Quality of Service (QoS) of Alibaba Cloud local disks. For more information, see Commands used to test the performance of local disks.

• The following table describes the performance metrics of the local NVMe SSDs that the i3 instance family uses.

Performanc e metric	ecs.i3.xlarg e	ecs.i3.2xlar ge	ecs.i3.4xlar ge	ecs.i3.8xlar ge	ecs.i3.13xla rge	ecs.i3.26xla rge
Maximum read IOPS	250,000	500,000	1,000,000	2,000,000	3,000,000	6,000,000
Maximum read throughput (GB/s)	1.5	3	6	12	18	36

Performanc	ecs.i3.xlarg	ecs.i3.2xlar	ecs.i3.4xlar	ecs.i3.8xlar	ecs.i3.13xla	ecs.i3.26xla
e metric	e	ge	ge	ge	rge	rge
Maximum write throughput (GB/s)	1	2	4	8	12	24

Note The performance data in the preceding table represents the highest performance levels of local storage in the i3 instance family. We recommend that you use images that contain Linux kernel 4.10 or later, such as Alibaba Cloud Linux 2 and CentOS 8.x images, to obtain optimal performance.

• The following table describes the performance metrics of the local NVMe SSDs that the i3g instance family uses.

Performance metric	ecs.i3g.2xlarg e	ecs.i3g.4xlarg e	ecs.i3g.8xlarg e	ecs.i3g.13xlar ge	ecs.i3g.26xlar ge
Maximum read IOPS	250,000	500,000	1,000,000	1,500,000	3,000,000
Maximum read throughput (GB/s)	1.5	3	6	9	18
Maximum write throughput (GB/s)	1	2	4	6	12

Note The performance data in the preceding table represents the highest performance levels of local storage in the i3g instance family. We recommend that you use images that contain Linux kernel 4.10 or later, such as Alibaba Cloud Linux 2 and CentOS 8.x images, to obtain optimal performance.

• The following table describes the performance metrics of the local NVMe SSDs that the i2 and i2g instance families use.

	Single disk performance	Overall instance	
Performance metric	ecs.i2.xlarge and ecs.i2g.2xlarge	Other i2 and i2g instance types	performance ^①
Maximum capacity (GiB)	894	1,788	8 × 1,788
Maximum read IOPS	150,000	300,000	1,500,000
Maximum read throughput (GB/s)	1	2	16

	Single disk performance	Overall instance	
Performance metric	ecs.i2.xlarge and ecs.i2g.2xlarge	Other i2 and i2g instance types	performance ^①
Maximum write throughput (GB/s)	0.5	1	8
Access latency	Within microseconds		

^① Overall instance performance data in the preceding table applies only to the ecs.i2.16xlarge instance type and represents the highest performance levels of local storage in the i2 instance family.

• The following table describes the performance metrics of the local NVMe SSDs that the i2ne and i2gne instance families use.

Performance metric	ecs.i2ne.xlarg e and ecs.i2gne.2xla rge	ecs.i2ne.2xlar ge and ecs.i2gne.4xla rge	ecs.i2ne.4xlar ge and ecs.i2gne.8xla rge	ecs.i2ne.8xlar ge and ecs.i2gne.16xl arge	ecs.i2ne.16xla rge
Maximum capacity (GiB)	894	1,788	2 × 1,788	4 × 1,788	8 × 1,788
Maximum read IOPS	250,000	500,000	1,000,000	2,000,000	4,000,000
Maximum read throughput (GB/s)	1.5	3	6	12	24
Maximum write throughput (GB/s)	1	2	4	8	16
Access latency	Within microseconds				

Note To obtain the maximum throughput performance of disks for Linux instances, we recommend that you use the latest versions of Alibaba Cloud Linux 2 images. Otherwise, Linux instances may be unable to deliver their maximum IOPS.

• The following table describes the performance metrics of the local NVMe SSDs that the i1 instance family uses.

Performance metric	Single disk performance	Overall instance performance ^②
Maximum capacity (GiB)	1,456	2,912

Performance metric	Single disk performance	Overall instance performance ^②	
Maximum IOPS	240,000	480,000	
Write IOPS ^①	min{165 × Capacity, 240,000}	2 v min(165 v Canacity, 240 000)	
Read IOPS ^①	111111(103 × Capacity, 240,000)	2 × min{165 × Capacity, 240,000}	
Maximum read throughput (GB/s)	2	4	
Read throughput (MB/s) ^①	min{1.4 × Capacity, 2,000}	2 × min{1.4 × Capacity, 2,000}	
Maximum write throughput (GB/s)	1.2	2.4	
Write throughput (MB/s) ^①	min{0.85 × Capacity, 1,200}	2 × min{0.85 × Capacity, 1,200}	
Access latency	Within microseconds		

① Items in the formulas used to calculate the performance specifications of a single local NVMe SSD:

- In the formula used to calculate the write IOPS, each GiB of capacity produces a write IOPS of 165 for a maximum of 240,000 IOPS.
- In the formula used to calculate the write throughput, each GiB of capacity produces a write throughput of 0.85 MB/s for a maximum of 1,200 MB/s.

Local SATA HDDs

Note You can test the bandwidth, IOPS, and latency of local NVMe SSDs to obtain the standard performance data and measure the QoS of Alibaba Cloud local disks. For more information, see Commands used to test the performance of local disks.

The following table describes the performance metrics of local SATA HDDs.

	d1 and d1ne		d2c		d2s	
Performanc e metric	Single disk performanc e	Overall instance performanc e	Single disk performanc e	Overall instance performanc e	Single disk performanc e	Overall instance performanc e
Maximum capacity (GiB)	5,500	154,000	3,700	44,400	7,300	219,000

^② Overall instance performance data in the preceding table applies only to the ecs.i1.14xlarge instance type and represents the highest performance levels for local storage in the i1 instance family.

	d1 and d1ne		d2c		d2s	
Performanc e metric	Single disk performanc e	Overall instance performance	Single disk performanc e	Overall instance performanc e	Single disk performanc e	Overall instance performanc e
Maximum throughput (MB/s)	190	5,320	190	2,280	190	5,700
Access latency	Within milliseconds					

? Note Overall instance performance data in the preceding table applies only to the ecs.d1.14xlarge, ecs.d1ne.14xlarge, ecs.d2c.24xlarge, and ecs.d2s.20xlarge instance types and represents the highest performance levels for local storage in the corresponding instance families.

Billing methods

Local disks are billed along with the instances to which they are attached. For more information, see Subscription and Pay-as-you-go.



? Note

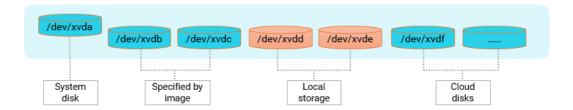
Disk initialization sequence

When you create an ECS instance that has local disks attached, all disks of the created instance are initialized based on the following rules:

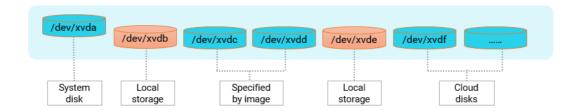
- Rule 1: If the image used to create the instance does not contain data disk snapshots, the local disks are initialized prior to the cloud disks that were created together with the instance.
- Rule 2: If the image used to create the instance contains data disk snapshots, the data disks created from these snapshots are initialized based on the sequence of data disks in the image. The remaining disks are initialized in the order that was specified in Rule 1.

For example, assume that a Linux image which contains the snapshots of two data disks is used to create an ECS instance. The disks on the created instance are initialized in the following sequence:

• If the device names of the two data disks in the image are /dev/xvdb and /dev/xvdc, Alibaba Cloud first allocates /dev/xvdb and /dev/xvdc as device names to the data disks created from the image. The system disk is initialized first. Then, the data disks are initialized in the following sequence: data disk 1 created from the image, data disk 2 created from the image, local disk 1, local disk 2, cloud disk 1, cloud disk 2, and cloud disk N. The following figure shows the sequence in which the disks are initialized.



• If the device names of the two data disks in the image are /dev/xvdc and /dev/xvdd, Alibaba Cloud first allocates /dev/xvdc and /dev/xvdd as device names to the data disks created from the image. Then, remaining available device names are allocated to the local disks first and then to other disks in ascending alphabetic order. The system disk is initialized first. Then, the data disks are initialized in the following sequence: local disk 1, data disk 1 created from the image, data disk 2 created from the image, local disk 2, cloud disk 1, cloud disk 2, and cloud disk N. The following figure shows the sequence in which the disks are initialized.



Lifecycle

A local disk shares the same lifecycle as the instance to which it is attached. For more information, see Instance lifecycle.

Impacts of instance operations on data stored on local disks

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

Instance operation	Data stored on local disks	Local disk
Restart the operating system, restart an instance by using the ECS console, or forcibly restart an instance.	Retained	Retained
Shut down the operating system, stop an instance by using the ECS console, or forcibly stop an instance.	Retained	Retained
Automatically recover an instance.	Erased	Released
Release an instance.	Erased	Released
When a subscription instance expires or when you have an overdue payment for a pay-as-you-go instance, the instance is stopped but has not been released.	Retained	Retained

Instance operation	Data stored on local disks	Local disk
When a subscription instance expires or when you have an overdue payment for a pay-as-you-go instance, the instance is stopped and then released.	Erased	Released
Renew an expired subscription instance.	Retained	Retained
Reactivate a pay-as-you-go instance that was stopped due to an overdue payment.	Retained	Retained

References

For information about retired local SSDs, see Ephemeral SSDs.

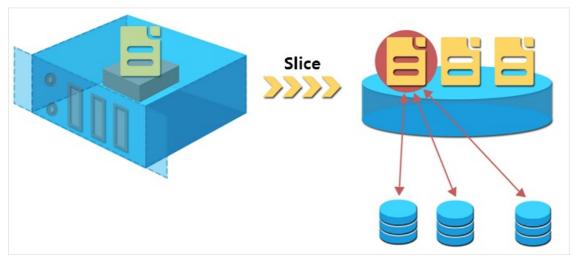
For information about how to handle system events on ECS instances that are equipped with local disks, see O&M scenarios and system events for instances equipped with local disks.

1.5. Triplicate storage

Alibaba Cloud disks use the triplicate storage technology to provide stable, efficient, and reliable access to data in ECS instances based on a distributed file system, achieving a data reliability of 99.999999%.

Overview

When you perform read and write operations on disks, the operations are translated into the corresponding processes on the files stored in Alibaba Cloud data storage system. Alibaba Cloud uses a flat design in which a linear address space is divided into chunks. Each chunk is replicated into three copies stored on different data nodes of the storage cluster to ensure data reliability.



All user level operations on disk data (including adding, modifying, and deleting data) are synchronized across three chunk copies at the underlying layer. This mechanism ensures the reliability and consistency of your data.

How triplicate storage works

Triplicate storage involves three types of key component: the master, chunk server, and client. Chunk servers are data nodes where chunk copies are stored. Each write operations is executed by the client in the following manner:

- 1. The client receives your write request and determines which chunk corresponds to the write operation.
- 2. The client queries the master to find the chunk servers where the three copies of the chunk are stored
- 3. The client sends write requests to the chunk servers returned from the master.
- 4. If the write operation succeeds on all three chunk copies, the client returns a success. Otherwise, the client returns a failure.

The master ensures that the copies of each chunk are distributed to different chunk servers cross different racks. This prevents data unavailability caused by the failure of a single chunk server or rack. The distribution strategy of the master takes many factors of the storage system into account, such as chunk server disk usage, chunk server distribution across racks, power distribution conditions, and node workloads.

Data protection

When a data node is damaged or disk faults occur on a data node, the number of valid copies of some chunks within the cluster may become less than three. In such cases, the master will initiate synchronization tasks and replicate data between chunk servers to ensure that there are three valid copies of each chunk in the cluster.



Notice To prevent data loss caused by viruses, accidental deletion, or attacks, we recommend that you use other protection methods such as data backup and periodic snapshots in addition to triplicate storage. Take all appropriate measures to ensure the security and availability of your data. For information about snapshots, see Snapshot overview.

2.Cloud disks 2.1. Create a cloud disk

2.1.1. Create a disk

You can create a subscription or pay-as-you-go data disk to increase the storage space of your Elastic Compute Service (ECS) instance.

Context

The following table describes the limits that apply to disks that use different billing methods.

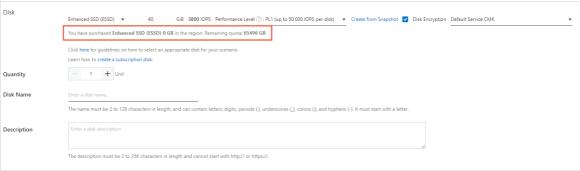
Disk	Limits
Pay-as-you-go disk	 You can have only a limited number of pay-as-you-go disks. You can view resource quotas in the ECS console. For more information, see View and increase resource quotas. For more information about the pay-as-you-go billing method, see Pay-as-you-go.
Subscription disk	 When you create a subscription disk, you must attach it to a subscription instance. You cannot separately create subscription disks. Subscription disks cannot be separately detached or released. Subscription disks expire and are released along with the instances to which they are attached. If you want to release a subscription disk, you can convert it to a pay-as-you-go disk, and then detach and release it.

In addition, you must take note of the following items:

- Disks are independent of each other and cannot be merged by formatting. Therefore, we recommend that you determine the number and capacity of disks that you need before you create them.
- We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes on multiple disks. This is because a snapshot can back up the data only of a single disk. If you create a logical volume on multiple disks by using LVM, data discrepancies may occur when you roll back these disks.

Procedure

- 1.
- 2.
- 3. On the Disks page, click Create Disk.
- 4. On the Disk page, configure the parameters described in the following table.



Parameter	Description		
	Specify whether to attach the disks to ECS instances.		
Attach	• Not Attach: creates disks without attaching them to ECS instances.		
	If you select Not Attach, only pay-as-you-go disks can be created. The disks and the ECS instances to which you want to attach the disks must reside within the same zone. Proceed with caution when you specify Region.		
	Note If you want to partition a disk in the ECS console after the disk is created, select Not Attach. After the disk is created, you can partition the disk when you attach the disk to a Linux instance. For more information, see Attach a data disk.		
	 Attach to ECS Instance: creates disks and attaches the disks to a specified ECS instance. 		
	If you select Attach to ECS Instance, you must specify ECS Instance by selecting a destination region and an ECS instance.		
	Specify the billing method of the disks.		
Billing Method	 Pay-as-you-go: Pay-as-you-go disks can be attached to subscription or pay-as-you-go instances. 		
-	 Subscription: Subscription disks must be attached to subscription instances. 		
	Select a disk category and specify the disk capacity.		
	Description of parameters in this section:		
Storage	 Performance Level: You can set performance levels for only enhanced SSDs (ESSDs). You can select a performance level based on the capacity o an ESSD. The performance of an ESSD varies with the capacity of the ESSD and the performance level. For more information, see ESSDs. 		
	 Create from Snapshot: You can select a snapshot to create disks. Therefore the created disks contain data from the snapshot. 		
	For more information, see Create a disk from a snapshot.		
	 Disk Encryption: Disk encryption is applicable to scenarios that require data security and regulatory compliance. If you select this option, data stored on the created disks is automatically encrypted. 		

Parameter	Description		
	Specify the number of disks that you want to create.		
Quantity	 Note When you create disks, take note of the number of data disks on the instance to which you want to attach the disks. The number of data disks that can be attached to an ECS instance is limited. For more information, see the "Elastic Block Storage (EBS) limits" section in Limits. Pay-as-you-go disks have capacity quotas. The total capacity that you purchase cannot exceed your remaining capacity quota. You can use the following formula to calculate the total disk capacity: Total capacity = Capacity of a single disk × Number of disks. After you configure the disk category, the purchased capacity and the remaining capacity quota are displayed on the disk buy page. 		
Release	Specify whether to release the automatic snapshots of the disks along with the disks or release the disks along with the instance to which they are attached. This parameter is available only when you set Attach to Attach to ECS Instance and set Billing Method to Pay-as-you-go.		
Terms of Service	Read and select ECS Terms of Service.		
Others (Optional)	Specify the Disk Name, Description, Tags, and Resource Group parameters for easy management.		

- 5. Confirm the configurations and fees and click Preview.
- 6. In the Preview message, confirm the purchase information and click **Create**.

 After the disks are created, you can view the created disks on the **Disks** page. However, these disks cannot be directly used on ECS instances. You must perform the following operations before you can use the disks on ECS instances.

What's next

The following table describes the operations that you must perform before you can use the created disks.

Scenario	What to do next
You set Attach to Attach to ECS Instance when you create the disks	 Perform the following operations to partition and format the disks: If the disks are larger than 2 TiB in size, partition them to GUID Partition Table (GPT). For more information, see Partition and format a data disk larger than 2 TiB in size. If the disks are smaller than or equal to 2 TiB in size, partition them to Master Boot Record (MBR). For more information, see Partition and format a data disk on a Linux instance or Partition and format a data disk on a Windows instance.

Scenario	What to do next	
	Perform the following operations to attach, partition, and format data disks: 1. Attach a data disk	
	? Note If you partition the disks when you attach the disks to instances, skip this step.	
You set Attach to Not Attach when you create the disks	(Optional) Perform the following operations to partition and format the disks:	
	 If the disks are larger than 2 TiB in size, partition them to GUID Partition Table (GPT). For more information, see Partition and format a data disk larger than 2 TiB in size. 	
	 If the disks are smaller than or equal to 2 TiB in size, partition them to Master Boot Record (MBR). For more information, see Partition and format a data disk on a Linux instance or Partition and format a data disk on a Windows instance. 	

Related information

- CreateDisk
- RunInstances
- CreateInstance

2.1.2. Create a disk from a snapshot

If you want to use data stored on an existing or a released disk to perform data extraction or fault analysis, you can use a snapshot of the disk to create disks. You can create disks from snapshots of system disks or data disks.

Prerequisites

A snapshot is created from a system disk or data disk. The snapshot ID is obtained. For more information, see Create a snapshot of a disk.

Context

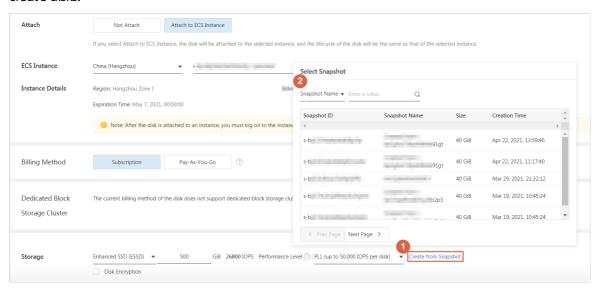
When you create a disk from a snapshot, take note of the following items:

- Disks are independent of each other and cannot be merged by formatting. Therefore, we recommend that you determine the number and capacity of disks that you need before you create them
- We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes on
 multiple disks. This is because a snapshot can back up the data only of a single disk. If you create a
 logical volume on multiple disks by using LVM, data discrepancies may occur when you roll back these
 disks.
- If you want to attach the new disk to the Elastic Compute Service (ECS) instance to which the disk from which the snapshot was created was attached, you must change the UUID of the new disk. For more information, see Modify the UUID of a disk.

• You can use a snapshot to create disks across zones.

Procedure

- 1.
- 2.
- 3. In the upper-left corner of the Disks page, click Create Disk.
- 4. On the Disk page, click **Create from Snapshot** and select the snapshot from which you want to create disks.



5. Configure other parameters described in the following table for the disks.

Parameter	Description
Attach	Specify whether to attach the created disks to ECS instances. Not Attach: creates disks without attaching them to ECS instances. If you select this option, only pay-as-you-go disks can be created. The disks and the ECS instances to which you want to attach the disks must reside within the same zone. The region and zone of a disk cannot be changed after the disk is created. Proceed with caution when you specify Region. Attach to ECS Instance: creates disks and attaches the disks to the specified ECS instance. If you select this option, you must specify ECS Instance by selecting a destination region and an ECS instance.
Billing Method	 Specify the billing method of the disks. Pay-As-You-Go: Pay-as-you-go disks can be attached to subscription or pay-as-you-go instances. Subscription: Subscription disks must be attached to subscription instances.

Parameter	Description
raidilletei	Description
	Select a disk category and specify the disk capacity. You must specify a disk capacity of at least the same size as the snapshot. If the disk capacity that you specify is greater than the snapshot size, you must re-partition the disk to ensure that the excessive disk capacity can be used.
Storage	Notice If the snapshot is less than 2,048 GiB in size and you want to specify a disk capacity of greater than 2,048 GiB, make sure that the source disk of the snapshot uses the GUID Partition Table (GPT) partition format. If the source disk does not use GPT, we recommend that you specify a disk capacity of less than 2,048 GiB to prevent data loss during partitioning. For more information, see Partition and format a data disk larger than 2 TiB in size
	The following section describes other parameters:
	 Performance Level: You can set performance levels only for enhanced SSDs (ESSDs). You can select a performance level based on the capacity of an ESSD. The disk performance varies with the capacity of the ESSD and the performance level. For more information, see ESSDs.
	 Disk Encryption: Disk encryption is suitable for scenarios that require data security and regulatory compliance. If you select this option, data stored on the created disks is automatically encrypted.
	Specify the number of disks that you want to create.
	 Note When you create disks, take note of the number of data disks on the instance to which you want to attach the disks.
	 Pay-as-you-go disks have capacity quotas. The total capacity that you purchase cannot exceed your capacity quota balance. You can use the following formula to calculate the total disk capacity: Total capacity = Capacity of a single disk × Number of disks. After you configure the disk category, the purchased capacity and the capacity quota balance are displayed on the disk buy page.
Quantity	

Parameter	Description
Release	Specify whether to release the automatic snapshots of the disks along with them or release the disks along with the instance to which they are attached. This parameter is available only when you select Attach to ECS Instance and set Billing Method to Pay-as-you-go .
Terms of Service	Read and select ECS Terms of Service.
Others (Optional)	Specify the Disk Name, Description, Tags, and Resource Group parameters for easy management.

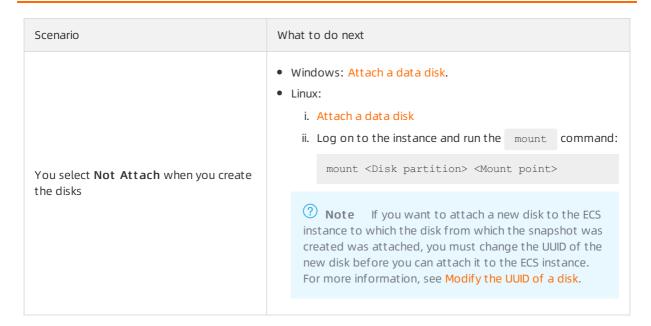
- 6. Confirm the configurations and fees and click **Preview**.
- 7. In the Preview message, confirm the purchase information and click **Create**.

 After the disks are created, you can view the created disks on the **Disks** page. However, these disks cannot be directly used on ECS instances. You must perform the following operations before you can use the disks on ECS instances.

What's next

The following table describes the operations that you must perform before you can use the created disks.

Scenario	What to do next
You select Attach to ECS Instance when you create the disks	 Windows: No operations are required. Linux: Log on to the instance and run the mount command: mount <disk partition=""> <mount point=""></mount></disk>
	Notice If you want to attach a new disk to the ECS instance to which the disk from which the snapshot was created was attached, you must change the UUID of the new disk before you can attach it to the ECS instance. For more information, see Modify the UUID of a disk.



Related information

CreateDisk

2.2. Attach a data disk

You can separately create a pay-as-you-go disk and manually attach the disk to an Elastic Compute Service (ECS) instance as a data disk.

Prerequisites

- The disk resides within the same zone as the instance to which you want to attach the disk.
- The instance is in the **Running** (*Running*) or **Stopped** (*Stopped*) state and is not in the **Locked** (*Lock ed*) state.
- The disk is in the **Unattached** (Available) state.
- You have no overdue payments within your Alibaba Cloud account.

Precautions

Before you attach a disk, take note of the following items:

- Disks that are created along with instances and subscription data disks that are created for subscription instances are automatically attached to the corresponding instances.
- A maximum of 64 data disks can be attached to an ECS instance. The maximum number of disks that can be attached to an instance varies with instance types. For more information, see Instance family.
 - **Note** You can attach one system disk and a maximum of 16 data disks to an instance when you create the instance. If the instance requires more data disks, attach more data disks after the instance is created.
- Disks that do not support the multi-attach feature and that have the multi-attach feature disabled can be attached only to a single instance at a time. Disks that support the multi-attach feature can be attached to multiple instances at a time. For more information, see Overview of disks that support NVMe.

Note The attach operation refers to the process of attaching disks to ECS instances in the ECS console, instead of mounting file systems in the operating systems of ECS instances by running the mount command. For information about how to format data disks and mount file systems, see Partition and format a data disk on a Linux instance and Partition and format a data disk on a Windows instance.

Attach a data disk

This section describes how to attach a data disk on the Instances page. You can also choose **Storage** & **Snapshots** > **Disks** in left-side navigation pane. Find the disk that you want to attach and choose **More** > **Attach** in the **Actions** column to attach the disk to an ECS instance.

- 1.
- 2.
- 3.
- 4.
- 5. Find the instance to which you want to attach disks and click the ID of the instance.
- 6. Click the Cloud Disk tab. In the upper-left corner of the Cloud Disk tab, click Attach Disk.
- 7. In the Attach Disk dialog box, configure parameters to attach a disk.
 - i. Select the disk that you want to attach and specify the release settings for the disk.

Parameter	Description	
Disk	Select the disk that you want to attach.	
Release Disk with Instance	If you select this option, the disk is automatically released when the instance to which it is attached is released. If you do not select this option, the disk is retained when the instance to which it is attached is released.	
	Note If the disk that you want to attach was detached from another instance on which it was previously used as the system disk, the instance specified by Release Disk with Instance refers to the instance from which the disk was detached, not the current instance.	
Delete Automatic Snapshots While Releasing Disk	If you select this option, automatic snapshots of the disk are released when the disk is released. To retain backup data, do not select this option.	

ii. Click Attach.

If the disk is attached, the state of the disk becomes In Use.

8. Create partitions and file systems before you can use the disk.

Note When the disk is attached to the instance but has no partitions or file systems, the disk is unusable on the instance. For example, you cannot view the mount information of the disk attached to a Linux instance by running the df -h command. Perform the following operations based on the disk status and the operating system of the instance.

Disk status	Operating system of the instance	What to do next
New empty disk	Linux	 For more information about the operations for disks that are smaller than 2 TiB in size, see Partition and format a data disk on a Linux instance. For more information about the operations for disks that are larger than 2 TiB in size, see Partition and format a data disk larger than 2 TiB in size.
	Windows Server	 For more information about the operations for disks that are smaller than 2 TiB in size, see Partition and format a data disk on a Windows instance. For more information about the operations for disks that are larger than 2 TiB in size, see Partition and format a data disk larger than 2 TiB in size.
Disk created from a snapshot	Linux	Log on to the instance and run the following command to mount the disk partitions that have file systems available: mount <data disk="" partition=""> <mount point=""> Notice If you want to attach a new disk to the ECS instance to which the disk from which the snapshot was created was attached, you must change the UUID of the new disk before you can attach it to the ECS instance. For more information, see Modify the UUID of a disk.</mount></data>
	Windows Server	N/A

Related information

References

AttachDisk

•

2.3. Format a data disk

2.3.1. Partition and format a data disk on a Linux instance

After you attach a new data disk to an Elastic Compute Service (ECS) instance, you must create and mount one or more file systems on the disk. This topic describes how to partition and format a new data disk on a Linux instance.

Prerequisites

A data disk that is created together with an instance is automatically attached to the instance. An independent data disk is created and attached to an instance. For more information about how to attach data disks to ECS instances, see Attach a data disk.

Note The attach operation refers to attaching disks to ECS instances in the ECS console, instead of mounting file systems by running the **mount** command within the operating systems of ECS instances.

Context

By default, the device names of data disks are assigned by the system based on the following naming conventions:

- I/O optimized instances:
 - The device names of data disks that are attached by using the Non-Volatile Memory Express
 (NVMe) protocol use the /dev/nvmeXn1 format. Examples: /dev/nvme1n1, /dev/nvme2n1, and /d
 ev/nvme3n1. For more information about disks that support NVMe, see Overview of disks that
 support NVMe.
 - The device names of data disks that are attached not by using the NVMe protocol use the /dev/vd [b-z] format. Examples: /dev/vdb, /dev/vdc, and /dev/vdd.
- The device names of data disks that are attached to non-I/O optimized instances use the /dev/xvd[b -z]format. Examples: /dev/xvdb, /dev/xvdc, and /dev/xvdd.

Data disks support the GUID Partition Table (GPT) and Master Boot Record (MBR) partition formats.

- GPT: can recognize partitions larger than 2 TiB in size and allows an unlimited number of partitions to be created on each disk.
- MBR: can recognize partitions up to 2 TiB in size and allows up to four partitions to be created on each disk.

Precautions

When you partition and format a data disk, take note of the following items:

- Disk partitioning and formatting are high-risk operations. Proceed with caution when you partition and format disks. This topic applies only to new data disks.
 - If a data disk that you want to format contains data, you must create snapshots for the disk before you format it. This can prevent data loss. For more information, see Create a snapshot of a disk.
- In ECS, only data disks can be partitioned, whereas system disks cannot.

If you use a third-party tool to forcefully partition a system disk, unforeseen risks such as system failures and data loss may occur. You can only extend or create partitions on system disks that have been extended. For more information, see Resize disks online for Linux instances.

Procedure

If you want to create GPT partitions, perform the following steps:

- 1. Step 1: Connect to an ECS instance and view the data disk
- 2. Step 2: Create a GPT partition for the data disk
- 3. Step 3: Create a file system for the partition
- 4. Step 4: Configure the /etc/fstab file and mount the partition

If you want to create MBR partitions, perform the following steps:

- 1. Step 1: Connect to an ECS instance and view the data disk
- 2. Step 2: Create an MBR partition for the data disk
- 3. Step 3: Create a file system for the partition
- 4. Step 4: Configure the /etc/fstab file and mount the partition

The following table describes the resources used in the examples of this topic.

Resource	Description	
Image used by the instance	Alibaba Cloud Linux 3.2104 64-bit public image	
Data disk	 Device name of the data disk: /dev/vdb Capacity: 40 GiB 	

Step 1: Connect to an ECS instance and view the data disk

Connect to an ECS instance and check whether the data disk is attached to the ECS instance.

1. Connect to the ECS instance.

For information about how to connect to an instance, see Connect to a Linux instance by using a password.

2. Run the following command to view the information of the data disk attached to the ECS instance:

fdisk -l

A command output similar to the following one is returned.

```
[root@ecs ~]# fdisk -1
Disk /dev/vda: 40 GiB, 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
Disklabel type: dos
Disk identifier: 0x3565eeef

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

Disk /dev/vdb: 40 GiB, 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
```

If no /dev/vd* device name is displayed, check whether the data disk is attached to the instance. / dev/vd* indicates the new data disk. In this example, /dev/vdb is used. For information about how to attach a data disk in the ECS console, see Attach a data disk.

- 3. Create a partition for the data disk.
 - For information about how to create GPT partitions, see Step 2: Create a GPT partition for the data disk.
 - For information about how to create MBR partitions, see Step 2: Create an MBR partition for the data disk.

Step 2: Create a GPT partition for the data disk

Perform the following steps to create a GPT partition that is larger than 2 TiB in size for the data disk.

- 1. If Parted and e2fsprogs are not installed, run the following commands to install them.
 - Run the following command to install Parted:

```
yum install -y parted
```

• Run the following command to install e2fsprogs:

```
yum install -y e2fsprogs
```

Note The yum command in this step is applicable to Linux distributions such as Cent OS.
For other Linux distributions, modify the command based on your package management software. For example, run the <a href="https://apt-get_install <a href="https://apt-

- 2. Use Parted to partition the data disk.
 - i. Run the following command to start partitioning the data disk:

```
parted /dev/vdb
```

ii. Run the following command to set the partition format to GPT:

```
mklabel gpt
```

iii. Run the following command to create a primary partition and specify the start and end sectors for the partition:

```
mkpart primary 1 100%
```

iv. Run the following command to check whether the partitions are aligned:

```
align-check optimal 1
```

A command output similar to the following one is returned:

```
Note If the partition is not aligned, 1 not aligned is returned. For more information about how to troubleshoot this issue, see FAQ.
```

v. Run the following command to view the partition table:

```
print
```

vi. Run the following command to exit Parted:

```
quit
```

A command output similar to the following one is returned.

```
[root@ecs ~]# parted /dev/vdb

GNU Parted 3.2

Using /dev/vdb

Welcome to GNU Parted! Type 'help' to view a list of commands.

(parted) mklabel spt

(parted) mkpart primary 1 100%

(parted) align-check optimal 1

1 aligned

(parted) print

Model: Virtio Block Device (virtblk)

Disk /dev/vdb: 42.9GB

Sector size (logical/physical): 512B/512B

Partition Table: gpt

Disk Flags:

Number Start End Size File system Name Flags

1 1049kB 42.9GB 42.9GB primary

(parted) ouit

Information: You may need to update /etc/fstab.
```

3. Run the following command to re-read the partition table:

```
partprobe
```

4. Run the following command to view the new partition:

```
fdisk -lu /dev/vdb
```

A command output similar to the following one is returned. If the new partition is created, the information of gpt is displayed.

Step 2: Create an MBR partition for the data disk

Perform the following steps to create an MBR partition that is up to 2 TiB in size for the data disk.

Note If your data disk is larger than 2 TiB in size or if you may need to extend your data disk to larger than 2 TiB in size, we recommend that you use the GPT partition format. For more information, see Step 2: Create a GPT partition for the data disk.

- 1. Create an MBR partition.
 - i. Run the following command to partition the data disk:

```
fdisk -u /dev/vdb
```

- ii. Enter $\it p$ to view the partition information of the data disk.
 - In this example, the data disk is not partitioned.
- iii. Enter *n* to create a partition.
- iv. Enter p to set the partition as a primary partition.
 - **?** Note When you create a single partition for a data disk, the partition must be the primary partition. If you want to create four or more partitions, enter *e* (extended) at least once to create at least one extended partition.
- v. Enter the partition number and press the Enter key.
 - In this example, only one partition is created. Press the Enter key to use the default value 1.
- vi. Enter the number of the first available sector and press the Enter key.
 - In this example, press the Enter key to use the default value 2048.
- vii. Enter the number of the last sector and press the Enter key.
 - In this example, only one partition is created. Press the Enter key to use the default value.
- viii. Enter p to view the partition information of the data disk.
- ix. Enter w to start partitioning. Then, exit after the disk is partitioned.

A command output similar to the following one is returned.

2. Run the following command to view the new partition:

```
fdisk -lu /dev/vdb
```

A command output similar to the following one is returned. If the new partition is created, the information of /dev/vdb1 is displayed.

```
[root@ecs ~]# fdisk -lu /dev/vdb
Disk /dev/vdb: 40 G1B, 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
Disklabel type: dos
Disk identifier: 0x37ec7546

Device Boot Start End Sectors Size Id Type
/dev/vdb1 2048 83886079 83884032 40G 83 Linux
```

Step 3: Create a file system for the partition

Create a file system for the new partition. The following commands are applicable to ext4 and XFS file systems. You can run one of the following commands to create the corresponding file system.

Note If the data disk is larger than 16 TiB, you must use the specified version of e2fsprogs to format the data disk. For more information, see Appendix 1: Update e2fsprogs on a Linux instance.

• Run the following command to create an ext4 file system:

```
mkfs -t ext4 /dev/vdb1
```

• Run the following command to create an XFS file system:

```
mkfs -t xfs /dev/vdb1
```

A command output similar to the following one is returned. In this example, an ext4 file system is created

```
[root@ecs ]# mkfs -t ext4 /dev/vdb1
mke2fs 1.45.6 (20-Mar-2020)
Creating filesystem with 10485504 4k blocks and 2621440 inodes
Filesystem UUID: 3caacl7' 3609a5
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000, 7962624

Allocating group tables: done
Writing inode tables: done
Creating journal (65536 blocks): done
Writing superblocks and filesystem accounting information: done
```

Step 4: Configure the /etc/fstab file and mount the partition

Write the information of the new partition to /etc/fstab to enable this partition to be automatically mounted on instance startup.

Notice We recommend that you use a universally unique identifier (UUID) to reference the new partition in /etc/fstab. The device names of other disks may change due to operations such as releasing a disk. If you use the device name of the data disk in /etc/fstab, your stored data may be affected if the device name changes.

1. Run the following command to back up etc/fstab:

```
cp /etc/fstab /etc/fstab.bak
```

- 2. Write the information of the new partition to /etc/fstab.
 - If you are a root user, you can run the following command to modify /etc/fstab:

```
echo `blkid /dev/vdb1 | awk '{print $2}' | sed 's/\"//g'` /mnt ext4 defaults 0 0 >> / etc/fstab
```

Description of the parameters in this command:

- /dev/vdb1: the data disk partition that has a file system created. Replace it with your actual partition name.
- /mnt: the directory to which the partition is mounted. Replace it with the actual directory of your partition.
- ext4: the file system type of the partition. Replace it with the file system type of the created file system.

```
? Note The Ubuntu 12.04 operating system does not support barriers. You must run the echo '`blkid /dev/vdb1 | awk '{print $3}' | sed 's/\"//g'` /mnt ext4 barrier=0 0 0' >> /etc/fstab command to write the information of the new partition to /etc/fstab. If your instance does not run an Ubuntu 12.04 operating system, ignore the note.
```

- If you are a common user, you can manually modify /etc/fstab. For more information, see
 Configure UUIDs in the fstab file to automatically attach data disks.
- 3. Run the following command to check the information of the new partition in /etc/fstab:

```
cat /etc/fstab
```

A command output similar to the following one is returned.

4. Run the following command to mount the file system configured in the /etc/fstab file:

```
mount -a
```

5. Run the following command to check the mount result:

```
df -h
```

A command output similar to the following one is returned. If the file system is mounted, the information of the new file system is displayed.

```
[root@ecs ~] # df -h
Filesystem Size Used Avail Use% Mounted on
devtmpfs 441M 0 441M 0% /dev
tmpfs 459M 0 459M 0% /dev/shm
tmpfs 459M 476K 459M 1% /run
tmpfs 459M 0 459M 0% /sys/fs/cgroup
/dev/vda1 40G 2.6G 35G 7% /
tmpfs 92M 0 92M 0% /run/user/0
/dev/vdb1 40G 49M 38G 1% /mnt
```

FAO

Issue: When GPT partitions are created, the partitions are not aligned.

Solution:

1. We recommend that you run the following commands:

```
cat /sys/block/vdb/queue/optimal_io_size
cat /sys/block/vdb/queue/minimum_io_size
cat /sys/block/vdb/alignment_offset
cat /sys/block/vdb/queue/physical_block_size
```

 ${\bf 2.} \ \ {\bf Calculate\ the\ start\ sector\ number\ of\ the\ optimal\ partition\ format\ by\ using\ the\ following\ formula:$

```
(<optimal io size>+<alignment offset>) /<physical block size> .
```

For example, if the start sector number is 1024, you can run the mkpart primary 1024s 100% command to create another primary partition.

2.3.2. Partition and format a data disk on a Windows instance

Before data can be stored in a new data disk that is attached to an Elastic Compute Service (ECS) instance, you must create and mount one or more file systems on the disk. This topic describes how to partition and format a new data disk on a Windows instance.

Prerequisites

An independent data disk is created and attached to an instance. For information about how to attach data disks to instances, see Attach a data disk.

Context

Data disks support the GUID Partition Table (GPT) and Master Boot Record (MBR) partition formats.

- GPT: can recognize partitions larger than 2 TiB in size and allows up to 128 partitions to be created on each disk.
- MBR: can recognize partitions up to 2 TiB in size and allows up to four primary partitions to be created on each disk.

Note The data disks created along with an instance are automatically partitioned and formatted. You can use them directly without the need to manually partition or format them as described in this topic.

Precautions

When you partition and format a data disk, take note of the following items:

• Disk partitioning and formatting are high-risk operations. Exercise caution when you partition and format disks. This topic applies only to new data disks.

If you want to format a data disk to which data has been written, you must create a snapshot for the disk before you format it to prevent data loss. For more information, see Create a snapshot of a disk.

• In ECS, only data disks can be partitioned, whereas system disks cannot.

If you use a third-party tool to forcibly partition a system disk, unforeseen risks such as system failures and data loss may occur. You can only extend or create partitions on system disks that have been extended. For more information, see Resize disks online for Windows instances.

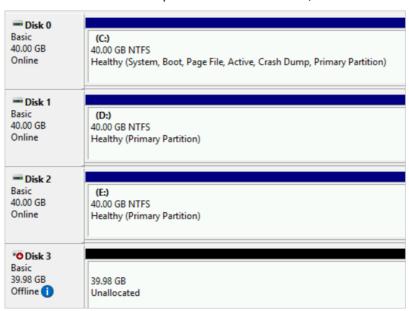
Procedure

In the following examples, an instance that runs a Windows Server 2012 R2 64-bit operating system is used.

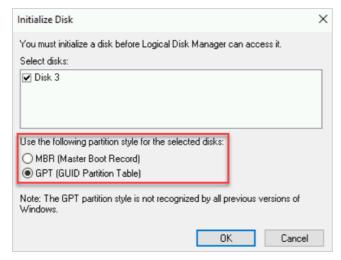
- 1. Open the Disk Management window and find a data disk that is not partitioned or formatted.
 - i. Connect to an instance that runs a Windows Server operating system.
 For more information, see Connect to a Windows instance by using a password.
 - ii. On the Windows Server desktop, click the **Server Manager** icon in the lower-left corner.



- iii. In the **Server Manager** window, choose **Tools** > **Computer Management** in the upper-right corner.
- iv. In the left-side navigation pane, choose Computer Management (Local) > Storage > Disk Management.
- v. Find a data disk that is not partitioned or formatted, which is in the Offline state.



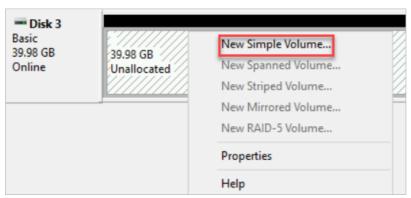
- 2. Right-click the blank area around the disk and select **Online**. When the disk goes online, it enters the **Not Initialized** state.
- 3. Right-click the blank area around the disk and select **Initialize Disk**.
- 4. In the Initialize Disk dialog box, select the disk, select a partition format, and then click OK.



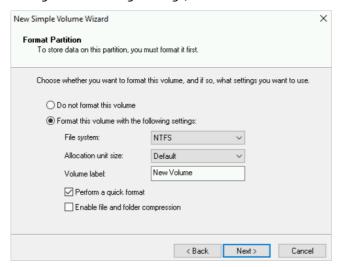
You can select GPT or MBR as the partition format.

- o GPT is a new partition format that can define data disks larger than 2 TiB in size. The maximum data disk size that GPT supports is determined based on the operating system and the file systems. In Windows, GPT supports up to 128 primary partitions on each disk.
 - **? Note** If your data disk is larger than 2 TiB in size or may need to be extended to larger than 2 TiB, we recommend that you select GPT.

- MBR can define data disks that are up to 2 TiB in size and supports up to four primary partitions on each disk. If you want to create more partitions on an MBR disk, you must use a primary partition as an extended partition and create logical partitions in the primary partition.
- 5. Right-click the **Unallocated** section and select **New Simple Volume**.



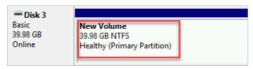
- 6. In the New Simple Volume Wizard dialog box, perform the following operations:
 - i. Click Next.
 - ii. In the Specify Volume Size step, set Simple volume size in MB and click Next.
 If you want to create only a single primary partition on the data disk, use the default value.
 - iii. In the **Assign Drive Letter or Path** step, select **Assign the following drive letter**, select a drive letter, and then click **Next**.
 - iv. In the Format Partition step, select Format this volume with the following settings, configure formatting settings, and then click Next.



v. Check the settings and click Finish.

Result

After the data disk is partitioned and formatted, it enters the Healthy state, as shown in the following figure.



In This PC, you can find a new drive named New Volume (E:). The data disk is ready for use.

2.3.3. Partition and format a data disk larger than 2 TiB in size

Use GPT to partition and format a data disk larger than 2 TiB in size

This topic describes how to partition and format a data disk larger than 2 TiB in size for different operating systems.

Prerequisites

- A data disk is attached to an ECS instance. For more information, see Attach a data disk.
- A remote connection is established to the ECS instance. For more information about how to connect to an ECS instance, see Connection methodsGuidelines on instance connection.

Context

- The amount of time required to create a snapshot of a data disk is proportional to the volume of data stored on the data disk. The more data stored on the disk, the longer it takes to create a snapshot.
- Alibaba Cloud Elastic Block Storage (EBS) supports the master boot record (MBR) and globally unique identifier partition table (GPT) partition formats. MBR is applicable to data disks up to 2 TiB in size and allows up to four primary partitions to be created per disk. If you want to partition a data disk larger than 2 TiB in size, you must use the GPT partition format.
 - Notice Conversion between MBR and GPT may cause data loss. When you create a disk from a snapshot or resize a disk and you expect the resulting disk size to exceed 2 TiB, we recommend that you check whether the disk uses the MBR partition format. If the disk uses the MBR partition format and you want to retain the disk data when you change the partition format, we recommend that you create and attach a new data disk to the instance, partition and format the new data disk in the GPT partition format, and then copy data from the original data disk to the new data disk.
- To partition and format a data disk larger than 2 TiB in size, we recommend that you use the partition tools, partition format, and file systems described in the following table.

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

Partition and format a data disk larger than 2 TiB in size on a Windows instance

This section describes how to partition and format a data disk larger than 2 TiB in size on a Windows instance. In this example, Windows Server 2012 R2 64-bit is used.

- 1. On the Windows Server desktop, press Win+R.
- 2. In the Run dialog box, enter *diskmgmt.msc* and click **OK** to open Disk Management.
- 3. Find the disk to be partitioned and formatted. In this example, Disk 1 is used. The disk is in the

Offline state.

- 4. Right-click the blank area next to Disk 1 and select Online.
 - After Disk 1 comes online, it enters the Not Initialized state.
- 5. Right-click the blank area next to Disk 1 and select Initialize Disk.
- 6. In the Initialize Disk dialog box, select Disk 1 and select GPT (GUID Partition Table) as the disk partition format.
- 7. In the **Disk Management** window, right-click the **Unallocated** section of **Disk 1**, and then select **New Simple Volume** to create a 4 TiB volume in the NTFS format.
- 8. In the New Simple Volume Wizard window, click Next and perform the following operations:
 - i. Specify Volume Size: Specify the size of the simple volume. If you want to create only one primary partition, use the default value. Click Next. You can also divide Disk 1 into multiple partitions.
 - ? Note In theory, an NTFS volume can contain up to 2^{64} -1 clusters. However, in Windows XP Pro, an NTFS volume can contain up 2^{32} -1 clusters. For example, if the cluster size is 64 KiB, the maximum NTFS volume size is 256 TiB. If the cluster size is 4 KiB, the maximum NTFS volume size is 16 TiB. NTFS automatically selects a cluster size based on the disk capacity.
 - ii. Assign Drive Letter or Path: Select a drive letter. D is used in this example. Click Next.
 - iii. Format Partition: Configure the formatting settings such as the file system, allocation unit size, and volume label, and then determine whether to select **Perform a quick format** and **Enable file and folder compression**. In this example, **Perform a quick format** is selected.

 Click **Next**.
 - iv. After a new simple volume is created, click **Finish** to close the **New Simple Volume Wizard** window.

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

Notice Conversion between partition formats may cause data loss. Make sure that you have backed up the data on the disk before you convert the disk to a different partition format.

This section describes how to convert the partition format of a 3 TiB data disk on a Windows instance. In this example, Windows Server 2012 R2 64-bit is used.

- 1. On the Windows Server desktop, right-click the **Start** icon and select **Disk Management**.
- 2. Find the disk for which you want to convert the partition format. In this example, Disk 2 is used.
- 3. Right-click a simple volume and select **Delete Volume**.
- 4. Right-click the blank area next to Disk 2, and select Convert to GPT Disk.
- 5. In the **Disk Management** window, right-click the **Unallocated** section of Disk 2, and then select **New Simple Volume** to create a 3 TiB volume in the NTFS format.
- 6. In the New Simple Volume Wizard window, click Next and perform the following operations:

- i. Specify Volume Size: Specify the size of the simple volume. If you want to create only one primary partition, use the default value. Click Next. You can also divide Disk 2 into multiple partitions.
 - ? Note In theory, an NTFS volume can contain up to 2^{64} -1 clusters. However, in Windows XP Pro, an NTFS volume can contain up 2^{32} -1 clusters. For example, if the cluster size is 64 KiB, the maximum NTFS volume size is 256 TiB. If the cluster size is 4 KiB, the maximum NTFS volume size is 16 TiB. NTFS automatically selects a cluster size based on the disk capacity.
- ii. Assign Drive Letter or Path: Select a drive letter. E is used in this example. Click Next.
- iii. Format Partition: Configure the formatting settings such as the file system, allocation unit size, and volume label, and then determine whether to select **Perform a quick format** and **Enable file and folder compression**. In this example, **Perform a quick format** is selected. Click **Next**.
- iv. After a new simple volume is created, click Finish to close the **New Simple Volume Wizard** window.

Partition and format a data disk larger than 2 TiB in size on a Linux instance

This section describes how to use the Parted and e2f sprogs tools to partition and format a data disk larger than 2 TiB in size on a Linux instance. In this example, Alibaba Cloud Linux 2.1903 LTS 64-bit is used. Assume that the data disk is a new 3 TiB empty disk and its device name is /dev/vdb.

Prerequisites: The Parted and e2fsprogs tools are installed on your Linux instance.

Run the following command to install the Parted tool:

```
yum install -y parted
```

• Run the following command to install the e2fsprogs tool:

```
yum install -y e2fsprogs
```

Perform the following operations to partition and format a data disk larger than 2 TiB in size and mount the file system.

1. Connect to the instance.

For information about how to connect to an ECS instance, see Connect to a Linux instance by using a password.

2. Check whether a data disk exists.

Run the following command:

```
fdisk -l
```

A command output similar to the following one is returned. The command output includes information about the data disk. If no data disk information is returned, the instance does not have data disks attached.

```
Disk /dev/vdb: 3221.2 GB, 3221225472000 bytes, 6291456000 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes
```

- 3. Use the Parted tool to partition the data disk.
 - i. Run the following command to start partitioning:

```
parted /dev/vdb
```

ii. Run the following command to convert the partition format from MBR to GPT:

```
mklabel gpt
```

iii. Run the following command to create a primary partition and specify the start and end sectors for the partition:

```
mkpart primary 1 100%
```

iv. Run the following command to check whether the partition is aligned:

```
align-check optimal 1
```

A command output similar to the following one is returned:

```
1 aligned
```

Note If 1 not aligned is returned, the partition is not aligned. We recommend that you run the following commands and use the (<optimal_io_size> + <alignment_off set>)/<physical_block_size> formula to obtain the start sector number to align the partition for optimal performance. For example, if the start sector number is 1024, you can then run the mkpart primary 1024s 100% command to create another primary partition.

```
cat /sys/block/vdb/queue/optimal_io_size
cat /sys/block/vdb/queue/minimum_io_size
cat /sys/block/vdb/alignment_offset
cat /sys/block/vdb/queue/physical_block_size
```

v. Run the following command to view the partition table:

```
print
```

vi. Run the following command to exit the Parted tool:

```
quit
```

The following figure shows the result of partitioning by using the Parted tool.

```
[root@ecs ] # parted /dev/vdb

GNU Parted 3.2

Using /dev/vdb

Welcome to GNU Parted! Type 'help' to view a list of commands.

(parted) mklabel gpt

(parted) align-check optimal 1

1 aligned

(parted) print

Model: Virtio Block Device (virtblk)

Disk /dev/vdb: 42.9GB

Sector size (logical/physical): 512B/512B

Partition Table: gpt

Disk Flags:

Number Start End Size File system Name Flags

1 1049kB 42.9GB 42.9GB primary

(parted) ouit

Information: You may need to update /etc/fstab.
```

4. Run the following command to enable the system to re-read the partition table:

```
partprobe
```

- 5. Run one of the following commands to create a file system for the /dev/vdb1 partition. Run one of the following commands to create a file system based on your needs:
 - o Create an ext4 file system.

```
mkfs -t ext4 /dev/vdb1
```

o Create an XFS file system.

```
mkfs -t xfs /dev/vdb1
```

? Note

- If the capacity of the data disk is 16 TiB, you must format the disk by using the specified version of e2fsprogs. For more information, see the Appendix 1: Update e2fsprogs on a Linux instance section of this topic.
- If you want to disable the lazy init feature of an ext4 file system to avoid impacts on the I/O performance of data disks, see the Appendix 2: Disable the lazy init feature on a Linux instance section of this topic.
- 6. Run the following command to create a mount point named /test:

```
mkdir /test
```

7. Run the following command to mount the /dev/vdb1 partition to /test:

```
mount /dev/vdb1 /test
```

8. Run the following command to view the current disk space and usage:

```
df -h
```

If the command output shows the information of the new file system, the mount operation is successful. You can use the new file system.

```
[root@ecs~]# df -h
Filesystem Size Used Avail Use% Mounted on
devtmpfs 457M 0 457M 0% /dev
tmpfs 467M 0 467M 0% /dev/shm
tmpfs 467M 476K 467M 1% /run
tmpfs 467M 0 467M 0% /sys/fs/cgroup
/dev/vda1 40G 2.3G 36G 6% /run/user/0
/dev/vdb1 3.0T 89M 2.9T 1% /test
```

- 9. (Recommended) Write the information of the new partition to /etc/fstab to enable this partition to be automatically mounted on system startup.
 - i. Run the following command to back up etc/fstab:

```
cp /etc/fstab /etc/fstab.bak
```

ii. Run the following command to write information of the new partition to /etc/fstab:

```
echo `blkid /dev/vdb1 | awk '{print $2}' | sed 's/\"//g'` /test ext4 defaults 0 0 >
    /etc/fstab
```

? Note

- You must run the command as the root user. If you are a common user, you can run the su command to switch to the root user, and then run this command.

 Alternatively, you can run the sudo vi /etc/fstab command to edit /etc/fstab.
- We recommend that you use a universally unique identifier (UUID) to reference the new partition in /etc/fstab. You can run the blkid command to obtain the UUID of the new partition.
- iii. Run the following command to check the information of /etc/fstab:

```
cat /etc/fstab
```

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

The 3 TiB data disk is partitioned and formatted.

Appendix 1: Update e2fsprogs on a Linux instance

If the disk capacity is 16 TiB, you must use e2fsprogs 1.42 or later to format the partitions of the disk to ext4. If e2fsprogs of a version earlier than 1.42 is used, the following error occurs:

```
mkfs.ext4: Size of device /\text{dev}/\text{vdb} too big to be expressed in 32 bits using a blocksize of 4096.
```

Perform the following operations to install a later version of e2fsprogs. In this example, e2fsprogs 1.42.8 is used.

1. Run the following command to check the current e2fsprogs version:

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

A command output similar to the following one is returned:

```
$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 version is earlier than 1.42, perform the following operations to update the software.

2. Run the following command to download e2fsprogs 1.42.8. You can visit e2fsprogs to obtain the latest software package.

```
wget https://www.kernel.org/pub/linux/kernel/people/tytso/e2fsprogs/v1.42.8/e2fsprogs-1
.42.8.tar.gz
```

- 3. Run the following commands in sequence to compile the tool of a later version.
 - i. Run the following command to decompress the software package:

```
tar xvzf e2fsprogs-1.42.8.tar.gz
```

ii. Run the following command to go to the package directory:

```
cd e2fsprogs-1.42.8
```

iii. Run the following command to generate the Makefile file:

```
./configure
```

iv. Run the following command to compile e2fsprogs:

```
make
```

v. Run the following command to install e2fsprogs:

```
make install
```

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

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

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

By default, the lazy init feature of an ext4 file system is enabled. When this feature is enabled, the instance initiates a thread to continuously initialize the metadata of the ext4 file system. If you partition and format a data disk when this feature is enabled, the IOPS of the disk may be affected for a short period after the disk is formatted.

If you want to test the performance of a data disk immediately after the disk is formatted, run the following command to disable the lazy init feature when you format the file system:

```
mke2fs -0 64bit,has_journal,extents,huge_file,flex_bg,uninit_bg,dir_nlink,extra_isize -E la
zy_itable_init=0,lazy_journal_init=0 /dev/vdb1
```

Note If the lazy init feature is disabled, it may take an extended period of time to format the file system. For example, the file system of a 32 TiB data disk may take 10 to 30 minutes to format. Choose whether to disable the lazy init feature.

2.3.4. Create a file system on a raw data disk

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

Prerequisites

A cloud disk has been created and attached to an ECS instance. For more information, see Create a disk and Attach a data disk.

Procedure

1.

2. Run the following command to view the name of the attached cloud disk.

```
fdisk -1
```

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

```
dzdZ:~# fdisk -l
root@iZuf6em
root@iZuf6em | dzdZ:~# fdisk -1 | Disk /dev/vda: 80 GiB, 85899345920 bytes, 167772160 sectors Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x6c740fc2
Device
            Boot Start
                              End Sectors Size Id Type
                  2048 83886046 83883999
                                             40G 83 Linux
/dev/vda1 *
Disk /dev/vdb: 100 GiB, 107374182400 bytes, 209715200 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes
```

- 3. Create a file system for the /dev/vdb data disk. Example:
 - o Create an ext4 file system

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

o Create an ext3 file system

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

Create an xfs file system

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

Create a btrfs file system

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

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

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

```
mkdir /media/vdb
```

5. Attach the disk to the mount directory.

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

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

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

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

Related information

- Partition and format a data disk on a Windows instance
- Partition and format a data disk on a Linux instance

2.4. Replace the operating systems of instances

2.4.1. Replace the operating system of an instance by using a public image

Replace the image of an instanceReplace the image of an instance

You can replace the operating system of an Elastic Compute Service (ECS) instance by replacing its system disk. This topic describes how to use a public image to replace the operating system of an instance.

Prerequisites

- Snapshots are created to back up instance data to prevent data loss caused by the automatic release of the system disk to be replaced. For more information, see Create a snapshot of a disk.
- The instance whose operating system you want to replace is in the **Stopped** state. For information about how to stop an instance, see **Stop an instance**.



Context

This topic describes how to replace the operating system of a single instance in the ECS console. If you want to batch replace the operating systems of multiple instances, you can use the ACS-ECS-BulkyReplaceSystemDisk public template provided by Operation Orchestration Service (OOS). To learn more about the template, go to the OOS console.

For information about how to use a non-public image such as a custom image to replace the operating system of an instance, see Replace the operating system of an instance by using a non-public image. For information about image types, see Image overview.

Precautions

Risks may arise when you replace the operating system of an instance. Before you perform this operation, take note of the following items:

- When you replace the operating system of an instance, Alibaba Cloud assigns a new system disk to
 the instance. The system disk ID changes, and the original system disk is released. The category of the
 system disk, the IP addresses of the instance, and the MAC addresses of the elastic network
 interfaces (ENIs) bound to the instance remain unchanged.
- After you replace the operating system of an instance, you must redeploy environments on the new system disk to run services. This may interrupt the services that are running on the instance for an extended period of time.
- Snapshots of the original system disk cannot be used to roll back the new system disk, but can be used to create custom images. Take note of the following items about the snapshots based on their types:
 - The manual snapshots of the original system disk are not automatically released.
 - If Delete Automatic Snapshots While Releasing Disk is enabled for the original system disk, the automatic snapshots of the disk are automatically released along with the disk. If Delete Automatic Snapshots While Releasing Disk is not enabled for the original system disk, the automatic snapshots of the disk are not automatically released until they expire.
- If you want to make replacements between Windows Server and Linux operating systems, take note of the following items:
 - You can make replacements between Windows Server and Linux operating systems only within regions inside mainland China.
 - Make sure that the hostname of the instance meets the requirements of the replacement operating system, and delete invalid characters. For information about how to modify the hostname, see Modify the properties of an instance or ModifyInstanceAttribute.
- If you want to replace the operating system of an instance with Windows Server, take note of the following items:
 - The system disk must have at least 1 GiB space reserved. Otherwise, the instance may not be able to start after the operating system is replaced.
 - Microsoft no longer provides support for Windows Server 2003. To ensure security for your data, we recommend that you do not use Windows Server 2003 for your instances. Alibaba Cloud no longer provides Windows Server 2003 images. For more information, see End of support for Windows Server 2003 public images.

Procedure

1. Go to the Instances page in the ECS console.

- i.
- ii.
- iii.
- 2. Find the instance whose operating system you want to replace and choose **More > Disk and Image > Replace System Disk** in the **Actions** column.
- 3. In the Replace System Disk message, read the notes and click OK.
- 4. On the Change Operating System page, configure parameters.
 - i. In the **Image** section, select **Public Image** and then select a public image and an image version.
 - For information about how to use a non-public image to replace the operating system of an instance, see Replace the operating system of an instance by using a non-public image.
 - ii. In the **System Disk** section, specify the capacity of the new system disk based on your business needs and the requirements of your selected image.

You cannot replace the current system disk with a disk of a different category. The minimum capacity of the new system disk is determined by the image type and the capacity of the current system disk. Your specified disk capacity must be smaller than or equal to 500 GiB.

lmage	Size range of a resized system disk (GiB)
CoreOS and FreeBSD	[Max{30, Original size of the system disk}, 500]
Other Linux distributions	[Max{20, Original size of the system disk}, 500]
Windows Server	[Max{40, Original size of the system disk}, 500]

? Note If you have renewed your instance and downgraded its configurations, the system disk capacity cannot be modified until the next subscription period starts.

iii. In the Security Settings section, configure the authentication method.

Authentication method	Description
Key Pair	Select an existing key pair or click Create Key Pair to create a key pair. After a key pair is created, go back to the ECS instance creation wizard and click the icon to query the most recent key pair list. For more information, see Create an SSH key pair. Note Key pairs can be used to log on only to Linux instances.
Password	Enter and confirm a password. When you log on to an instance by using a username and a password, the default username for Linux is root and that for Windows is administrator.
Set after Change	After the operating system is replaced, bind the key pair or reset the password for the instance. For more information, see Bind an SSH key pair to an instance and Reset the logon password of an instance.

- 5. Check the fees, read and select ECS Service Terms, and then click Create Order.
 - For more information about billing details of the image, see Image overview.
 - For more information about prices of system disks, visit the Pricing tab of the Elastic Compute Service product page.
- 6. Complete the payment (if any) as prompted.
 It takes about 10 minutes to replace the operating system. After the operating system is replaced, you can go back to the ECS console and find that the instance automatically enters the Running state. You can click the instance ID to go to the Instance Details page and view the information of the new image.

What to do next

After you replace the operating system of an instance, you can perform the following operations based on the actual conditions:

- Redeploy environments on the new operating system to run services.
- If the new operating system is a Linux operating system and the instance has data disks attached, you must write the information of the data disk partitions to the /etc/fstab file and run the mount command to mount the partitions. For more information, see Configure UUIDs in the fstab file to automatically attach data disks.
- Optional. The automatic snapshot policy applied to the original system disk automatically becomes invalid, and you must configure an automatic snapshot policy for the new system disk. For more information, see Apply or disable an automatic snapshot policy.
- Optional. You can delete the snapshots of the original system disk that are no longer needed. For more information, see Delete a snapshot.

FAO

Question	Answer
How do I use the data disks of an instance after I replace its Linux operating system with Windows Server?	 By default, Windows Server operating systems cannot recognize ext3, ext4, or xfs file systems. We recommend that you perform one of the following operations on your data disks: If your data disks do not contain important data, re-initialize the disks and format them to file systems that are supported by the Linux operating system. For more information, see Re-initialize a system disk and Partition and format a data disk on a Windows instance. If your data disks contain important data, separately install software (such as Ext2Read and Ext2Fsd) that allows the Windows Server operating system to work with ext3, ext4, or xfs file systems.
How do I use the data disks of an instance after I replace its Windows Server operating system with Linux?	By default, Linux operating systems cannot recognize New Technology File System (NTFS) file systems. We recommend that you perform one of the following operations on your data disks: If your data disks do not contain important data, re-initialize the disks and format them to file systems that are supported by the Linux operating system. For more information, see Re-initialize a system disk and Partition and format a data disk on a Linux instance. If your data disks contain important data, separately install software (such as NTFS-3G) that allows the Linux operating system to work with NTFS file systems.

2.4.2. Replace the operating system of an instance by using a non-public image

You can replace the operating system of an Elastic Compute Service (ECS) instance by replacing its system disk. This topic describes how to use a non-public image such as a custom image to replace the operating system of an instance.

Prerequisites

Context

For information about how to use a public image to replace the operating system of an instance, see Replace the operating system of an instance by using a public image. For information about image types, see Image overview.

Precautions

Risks may arise when you replace the system disk of an instance. Before you perform this operation, take note of the following items:

- •
- •
- •
- •
- If you want to replace the operating system to Windows Server, take note of the following items:

0

0

Procedure

1.

2.

3.

- 4. On the **Change Operating System** page, configure parameters.
 - i. In the Image section, select Custom Image, Shared Image, or Market place Image and then select an image and an image version.

ii.

iii. In the Security Settings section, configure the authentication method.

Authentication method	Description	
Key Pair	Select an existing key pair or click Create Key Pair to create a key pair. After a key pair is created, go back to the ECS instance creation wizard and click the icon to query the most recent key pair list. For more information, see Create an SSH key pair. Note Key pairs can be used to log on only to Linux instances.	
Use Predefined Password	Use the predefined password of your selected image. Note Only custom images or shared images support this authentication method.	
Password	Enter and confirm a password. When you log on to an instance by using a username and a password, the default username for Linux is root and that for Windows is administrator.	
Set after Change	After the operating system is replaced, bind the key pair or reset the password for the instance. For more information, see Bind an SSH key pair to an instance and Reset the logon password of an instance.	

5.

6.

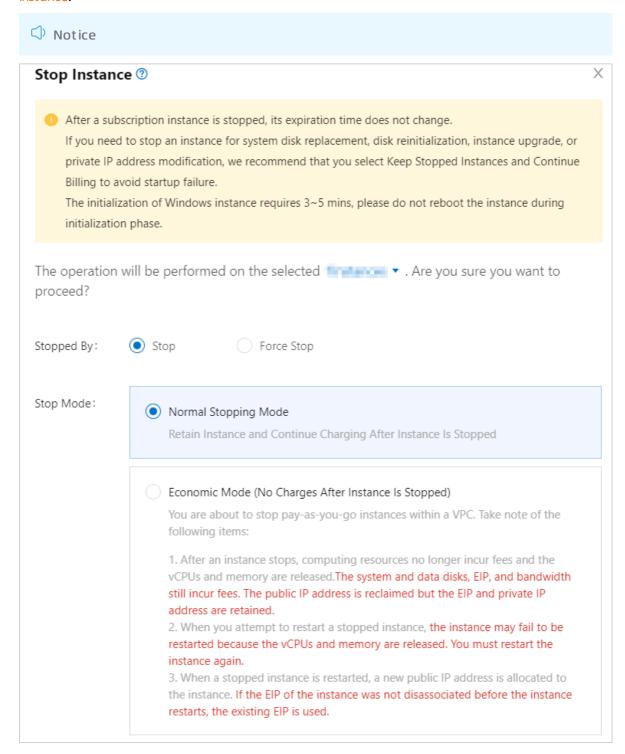
2.5. Reinitialize a cloud disk

2.5.1. Re-initialize a system disk

This topic describes how to re-initialize a system disk in the Elastic Compute Service (ECS) console. After a system disk is re-initialized, it is restored to the state when it was created.

Prerequisites

- Snapshots are created for the disk. For more information, see Create a snapshot of a disk.
- The instance to which the disk is attached is in the Stopped state. For more information, see Stop an instance.



• The instant access feature is disabled for the disk that is an enhanced SSD (ESSD). For more

information, see Enable or disable the instant access feature.

• An SSH key pair is created or imported for authentication if you re-initialize the system disk of a Linux instance. For more information, see Create an SSH key pair and Import an SSH key pair.

Context

When a system disk is attached to an ECS instance, you can re-initialize the system disk to restore it to the state when it was created. The following section describes the changes of a system disk after the disk is re-initialized:

• The system disk is restored to the state when it was created.

Warning When a disk is re-initialized, all data stored in the disk is deleted. We recommend that you create snapshots for the disk to back up data before you re-initialize the disk.

For example, if you use a Windows Server 2012 R2 public image to create an instance, the instance still runs Windows Server 2012 R2 after you re-initialize the system disk of the instance. However, the applications that were installed and data that was generated after the instance was created are deleted.

Note If the system disk of an instance is replaced with a new system disk, the new system disk instead of the previous system disk can be re-initialized.

- The automatic snapshot policy applied to the system disk before the re-initialization remains valid.
- The IP addresses and disk IDs of the instance remain unchanged.
- The automatic and manual snapshots of the system disk remain available. You can use these snapshots to roll back the disk. For more information, see Roll back a disk by using a snapshot.

Procedure

1.

2.

3.

- 4. Find the ECS instance whose system disk you want to re-initialize and click the instance ID to go to the Instance Details page.
- 5. Click the Cloud Disk tab.
- 6. Find the system disk and click **Reinitialize Disk** in the **Actions** column.
- 7. In the Reinitialize Disk dialog box, configure the parameters described in the following table.

Parameter	Description	
	Windows instances: Select Set Password. You can use the old password or specify a new password.	
Security	Linux instances: Select Set SSH Key or Set Password . • Set SSH Key : Bind an SSH key pair to the instance. You can log on to the instance by using the specified SSH key pair.	
	 Set Password: Specify a logon password. You can use the old password or specify a new password. 	

Parameter	Description	
Security Hardening	If you select Activate , security components are loaded to the instance free of charge. These components provide security features such as webshell detection, alerts of unusual logons, and protection against brute-force attacks.	
Instance Startup Policy	If you select Start Instance after Resetting Disk , the instance automatically starts after you re-initialize the system disk.	

8. Click Confirm.

What's next

• If data disks have been attached to a Linux instance before you re-initialize its system disk, you must re-create the mount points for the data disk partitions and mount file systems to the data disk partitions. For more information, see the "How do I re-attach data disks after I re-initialize the system disk of a Linux instance?" section in How do I re-attach data disks after I re-initialize the system disk of a Linux instance?

? Note After you re-initialize the system disk of a Linux instance, data stored in the data disks attached to the instance remains unchanged, but the mount information of the data disks is lost. You must re-create the mount points for the data disk partitions and mount file systems to the data disk partitions.

- After the system disk is re-initialized, you must redeploy applications and configurations to restore your business.
- If you have created snapshots for the system disk before the system disk is re-initialized, you can use the snapshots to create data disks and attach the data disks to the instance to obtain data originally stored in the system disk. For more information, see Create a disk from a snapshot.

Related information

- ReInit Disk
- Re-initialize a data disk

2.5.2. Re-initialize a data disk

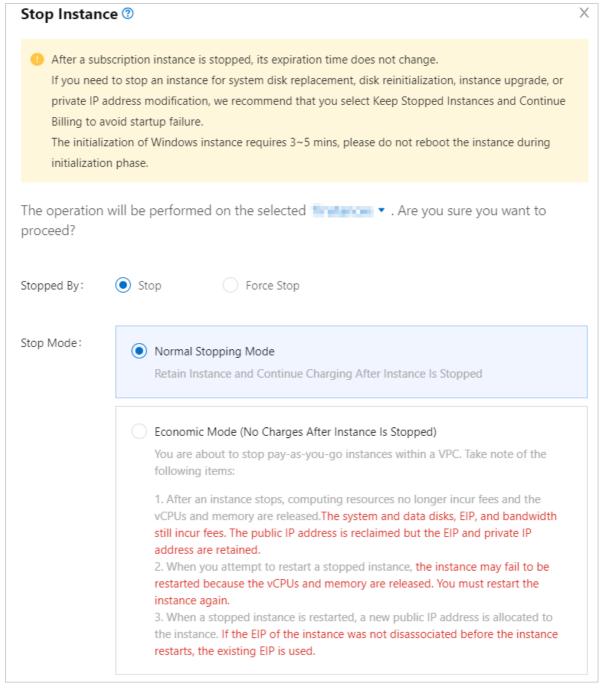
When a data disk is attached to an ECS instance, you can re-initialize the disk to restore it to the status when it was created.

Prerequisites

- Snapshots are created for the disk. For more information, see Create a snapshot of a disk.
- The data disk is attached to an instance. For more information about how to attach a data disk to an instance, see Attach a data disk.
- The instance to which the disk is attached is in the Stopped state. For more information, see Stop an instance.



64



- The instant access feature is disabled for the disk that is an enhanced SSD (ESSD). For more information, see Enable or disable the instant access feature.
- (Applicable only to Linux instances) If you create an empty data disk and add a command in the /etc/fstab file to mount partitions of the data disk on system startup, the command is not executed and the instance cannot start up as expected after you re-initialize the data disk. We recommend that you comment out the command in the /etc/fstab file. Perform the following steps:
 - i. Connect to the instance. For more information, see Overview.
 - ii. Run the vim /etc/fstab command.
 - iii. Press the I key to enter the edit mode.
 - iv. Find the command used to mount data disk partitions and comment it out by using #, as shown

in the following line.

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

Note /dev/vdb1 is a partition and /InitTest is a mount point used in this example. You
```

v. Press the Esc key to exit the edit mode, and enter :wq to save the file and exit.

Context

The status of a data disk after it is re-initialized varies based on its original status when it was created and the operating system that the instance runs.

- The data disk is restored to the initial status when it was created:
 - The data disk becomes an empty disk if it was originally an empty disk.

can modify the command based on your actual conditions.

- The data disk retains the data recorded in the source snapshot if it was created from a snapshot.
- For a Windows instance, after you re-initialize a data disk, the data disk is ready for use without additional operations regardless of its original status.
- For a Linux instance:
 - If a data disk was created from a snapshot, the data disk stores only the data recorded in the source snapshot after it is re-initialized. You do not need to mount the partitions again, but all the data generated after the disk was created is lost.
 - If a data disk was created as an empty disk, all the data and file systems on the disk are lost. You must reformat and partition the disk, and then mount the partitions again.

Procedure

- 1.
- 2.
- 3.
- 4. Find the ECS instance for which you want to re-initialize a data disk, and click the instance ID to go to the Instance Details page.
- 5. Click the Cloud Disk tab.
- 6. Find the data disk that you want to re-initialize, and click **Reinitialize Disk** in the **Actions** column.
- 7. In the Reinitialize Disk dialog box, read the notes and click Confirm.

What's next

- (Applicable only to Linux instances) If the data disk was created as an empty disk, you must format the data disk after you re-initialize it. For more information, see Partition and format a data disk on a Linux instance.
- After the data disk is re-initialized, you must redeploy the applications and reconfigure the parameters on the disk to restore your business as soon as possible.

Related information

- ReInit Disk
- Re-initialize a system disk

2.6. Roll back a disk by using a snapshot

This topic describes how to roll back a disk by using a snapshot. When the system does not respond or an incorrect operation is performed, you can roll back a disk to a previous version by using its snapshots. Before you roll back a disk, you must have created at least one snapshot for the disk. You can roll back both system disks and data disks. When you roll back a system disk, the instance is automatically bound with the current SSH key pair or the username and password combination.

Prerequisites

Before you roll back a disk by using a snapshot, take note of the following items:

• A snapshot of the disk to be rolled back is created, and no new snapshot is being created for the disk. For more information, see Create a snapshot of a disk.

Warning The rollback operation is irreversible. After a disk is rolled back, the data that you added, removed, or changed from the time when the snapshot is created to the time when the disk is rolled back is lost. To avoid data loss caused by accidental changes, we recommend that you create a snapshot for the disk to back up data before you roll back the disk.

- The disk is not released.
- After the operating system is replaced, snapshots of the previous system disk cannot be used to roll back the new system disk.
- If you use a disk to create a dynamic extended volume or redundant array of independent disks (RAID), you must stop all I/O operations on the disk before you roll back the disk.
- The disk is attached to an ECS instance and the instance is in the Stopped state. For more information, see Attach a data disk and Stop an instance.



Roll back a disk by using a snapshot

You can roll back a disk on the **Snapshots** or **Instances** page. The following example shows how to roll back a disk on the **Instances** page:

- 1.
- 2.
- 3.
- 4. Find the instance whose disk you want to roll back and click Manage in the Actions column.
- 5. The Instance Details tab appears. Click the Snapshot tab.
- 6. Find the snapshot and click **Roll Back Disk** in the **Actions** column.
 - Note You can roll back only one disk at a time. Other disks that are attached to the instance are unaffected. After the disk is rolled back, the disk is restored to the state of the entire disk at a certain point in time, instead of the state of a partition or a directory.

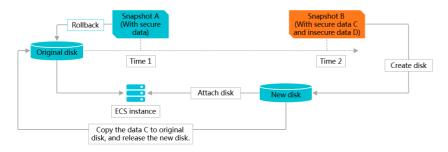
7. In the Roll Back Disk dialog box, click **OK**.



- Before you click OK, we recommend that you click Create Snapshot to back up the latest data
- If you select **Start Instance After Disk Rollback**, the instance automatically starts after the disk is rolled back.

(Optional) Synchronize data after a disk is rolled back

If you roll back a disk by using Snapshot A that was created at the point in time T1 and you want to synchronize the data that you added, removed, and changed after T1, you can perform the following operations:



- 1. Create Snapshot B for the original disk at the point in time T2.
- 2. Roll back the original disk by using Snapshot A.
- Create a disk by using Snapshot B.For more information, see Create a disk from a snapshot.
- 4. Attach the new disk to the same ECS instance.
 - For more information, see Attach a data disk.

5.

- 6. View the new disk.
 - o Windows instance: The new disk is displayed in the system.
 - Linux instance: Run the mount command to mount the partitions of the new disk.
- 7. Copy useful data from the new disk to the original disk.
- 8. (Optional)Release the new disk.

Related information

Reset Disk

2.7. Change the category of a disk

Alibaba Cloud offers different categories of disks to meet the storage performance and cost requirements for a variety of scenarios. You can change the categories of your disks based on your needs. For example, assume that you have created a standard SSD. If you want the standard SSD to deliver a higher IOPS, you can change this standard SSD into an enhanced SSD (ESSD).

Prerequisites

Snapshots of the disk whose category you want to change are created. For more information, see Create a snapshot of a disk.

Context

For information about the performance of different disk categories, see EBS performance.

Before you change the category of a disk, take note of the following items:

• Only the following disk category change operations are supported. Other operations such as disk category downgrade are not supported.

Category of the original disk		Description of supported category change
Ultra disk		You can convert ultra disks into standard SSDs, PL0 ESSDs, PL1 ESSDs, PL2 ESSDs, or PL3 ESSDs.
		Note You cannot convert ultra disks into standard SSDs in Hangzhou Zone D.
Standard SSD		You can convert standard SSDs into PL1, PL2, or PL3 ESSDs.
ESSD	PL0	You can convert PL0 ESSDs into PL1, PL2, or PL3 ESSDs.
	PL1, PL2, and PL3	You can convert PL1, PL2, and PL3 ESSDs based on their billing methods. If ESSDs use the pay-as-you-go billing method, you can convert the ESSDs between PL1, PL2, and PL3.
		 If ESSDs use the subscription billing method, you can upgrade the ESSDs only from the low performance level to the high performance level.
		You can upgrade PL1 ESSDs to PL2 ESSDs.
		You can upgrade PL1 ESSDs to PL3 ESSDs.
		■ You can upgrade PL2 ESSDs to PL3 ESSDs.

? Note

- Before you convert a disk on an Elastic Compute Service (ECS) instance into an ESSD, check whether the instance type supports ESSDs. For more information, see Instance family.
 - If the instance type does not support ESSDs, you must upgrade the instance type to support ESSDs and then convert the disk. For more information, see Instance families that support instance type changes.
- You can select a performance level for an ESSD based on the capacity of the ESSD. If you
 cannot select a higher performance level for an ESSD, resize the ESSD and then upgrade
 its performance level. For more information, see Overview.
- When you change the category of a disk, the performance of the disk may also change. We

recommend that you change the categories of your disks during off-peak hours.

- It may take an extended period of time to change the category of a disk. The amount of time required is determined based on the storage capacity, original category, and throughput of the disk during the change. After you submit the request to change the category of a disk, we recommend that you view the change progress by going to the Task logs page in the ECS console or by calling the DescribeTaskAttribute operation.
- Disk categories may be unable to change if resources are insufficient. If a failure occurs, try again later.

Limits

Phase	Limits
Before the category change	 Before you change the category of a disk, make sure that the following requirements are met: Up to five disk category change tasks can be concurrently performed within the same region and account. You cannot change the category of a disk attached to an instance that is included in a migration plan.
During the category change	 The following limits apply when the category of a disk is being changed: You cannot cancel the disk category change task after it is initiated. You cannot create snapshots for the disk. You cannot resize the disk. You cannot partition or format the disk. You cannot re-initialize the disk. You cannot roll back the disk by using snapshots. You cannot attach the disk to an instance or detach it from an instance. If the disk is an ESSD, you cannot change its performance level. If the disk is a system disk, you cannot replace the operating system of the instance.
After the category change	 The following limits apply after the category of a disk is changed: If only the performance level of an ESSD is changed, you can make another category change to the ESSD anytime. A disk can have its category changed only once within seven days.

Billing

After the category of a disk is changed, the billing of the disk has the following changes:

- After the category of a pay-as-you-go disk is changed, you are charged for the disk based on the new disk category.
- After the category of a subscription disk is changed, the price difference is calculated based on the remaining subscription duration and the new disk category. You must pay for the difference.

For information about disk pricing, see EBS devices.

Procedure

- 1.
- 2.
- 3.
- 4.
- 5. Find the disk whose category you want to change and choose More > Change Category in the Actions column.
- 6. In the Change Disk Category dialog box, set Category after Change.
- 7. Confirm the price and click **OK**.

Results

Use one of the following methods to view the change progress and the state of the disk:

- Method 1: Go back to the **Disks** page and view the category of the disk in the **Disk Category** column.
- Method 2: In the left-side navigation pane, choose Maintenance & Monitoring > Tasks. Find the
 change task and view the change progress in the Status column.

Related information

- ModifyDiskSpec
- DescribeDisks
- DescribeTaskAttribute

2.8. Modify the performance levels of ESSDs

You can upgrade the performance level of an enhanced SSD (ESSD) in the Elastic Compute Service (ECS) console. Only pay-as-you-go ESSDs can have their performance level downgraded. You can also call the ModifyDiskSpec operation to modify the performance level of an ESSD. The new performance level immediately takes effect. You do not need to create an ESSD or copy the data.

Context

When you create an ECS instance, you can specify an ESSD as the system disk or a data disk. You can also separately create an ESSD. For more information about ESSDs, see ESSDs.

To modify the performance level of an ESSD, take note of the following items:

- You cannot have overdue payments within your account.
- If your ESSD is attached to a pay-as-you-go ECS instance, make sure that the instance is not in the expired state.
- You can modify the performance level of a new ESSD only after it enters the **Unattached** (*Available*) state.
- After the performance level is modified, you are charged for the ESSD based on the new performance level.
- You can convert PLO ESSDs into PL1, PL2, or PL3 ESSDs. For more information, see Change the category of a disk.

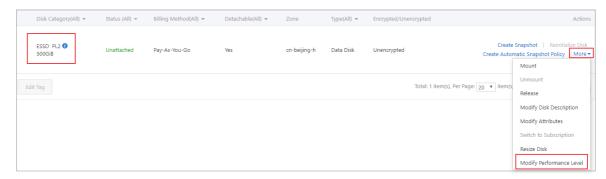
Procedure

1.

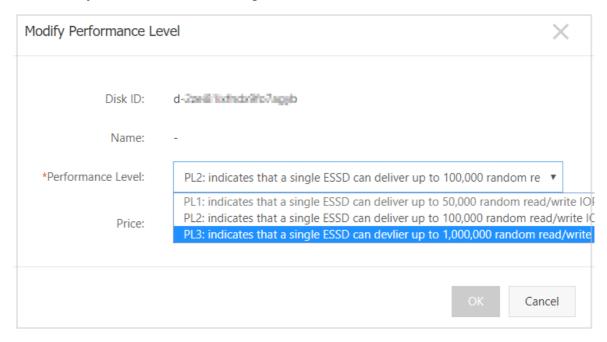
2.

3.

4. Find the ESSD whose performance level you want to modify and choose **More > Modify**Performance Level in the Actions column.



5. In the Modify Performance Level dialog box, select a new Performance Level and click OK.



The performance level that you can select for an ESSD is determined by its storage capacity. If the capacity of your ESSD is small and cannot be upgraded to a higher level, you must resize the ESSD and then upgrade its performance level. For more information, see Overview.

Related information

ModifyDiskSpec

2.9. View the monitoring data of a disk

You can view the monitoring data of disks in the Elastic Compute Service (ECS) or CloudMonitor console, such as IOPS and throughput. You can troubleshoot issues and determine whether to extend or upgrade disks based on their monitoring data. This topic describes how to view the monitoring data of disks in the ECS or CloudMonitor console.

Context

For information about the performance of different disk categories, see EBS performance. Disk performance can be measured by using the following metrics:

- IOPS: measures the number of read and write operations that an Elastic Block Storage (EBS) device can process per second. High IOPS is essential for transaction-intensive applications.
- Throughput: measures the amount of data transferred per second. Unit: MB/s. High throughput is essential for applications that require a large number of sequential read and write operations.

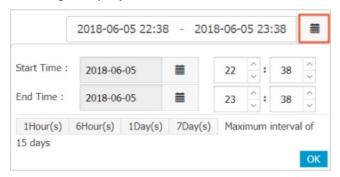
You can view the monitoring data of disks in the ECS console or view the monitoring metrics such as usage, IOPS, and throughput of disks in the CloudMonitor console.

- View the monitoring data of a disk in the ECS console
- View the monitoring data of a disk in the CloudMonitor console

View the monitoring data of a disk in the ECS console

You can view the monitoring data such as IOPS and throughput of disks in the ECS console.

- 1.
- 2.
- 3.
- 4. Find the disk whose monitoring data you want to view and click the disk ID to go to the **Disk Det ails** tab.
- 5. Click the **Disk Monitoring Information** tab.
- 6. In the **Monitoring Information** section, click the icon to specify the beginning and end of the time range to query.



7. View the IOPS and throughput data of the disk. You can click an individual chart to view the data of each metric.



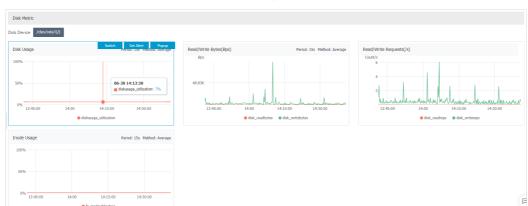
If monitoring data indicates that disk performance is bottlenecked, you can optimize your business applications or upgrade the disk. For information about how to upgrade a disk, see Change the category of a disk.

View the monitoring data of a disk in the CloudMonitor console

You can view the monitoring metrics such as usage, IOPS, and throughput of disks in the CloudMonitor console and configure alert rules based on actual scenarios. The following table describes the metrics of disks. For more information, see Metrics.

The following procedure describes how to view the monitoring data of a disk and configure alert rules for the disk in the CloudMonitor console.

- 1. Log on to the CloudMonitor console.
- 2. Install the CloudMonitor agent. For more information, see Install and uninstall the CloudMonitor agent for C++.
- 3. View the monitoring data of a disk.
 - i. In the left-side navigation pane, click **Host Monitoring**.
 - ii. In the instance list, find the instance whose disk metrics you want to view and click the name of the instance in the Host Name/ID column.
 - iii. On the Host Monitoring page, click the OS Monitoring tab.



iv. In the Disk Metric section, view the monitoring data of a disk that is attached to the instance.

4. Configure alert rules for the disk. For more information, see Alert service.

If monitoring data indicates that disk performance is bottlenecked, you can optimize your business applications or upgrade the disk. For information about how to upgrade a disk, see Change the category of a disk.

If the disk is almost full, delete data that is no longer needed to free up space or extend the disk. For information about how to extend a disk, see Overview.

Related information

- DescribeDiskMonit orDat a
- What is CloudMonitor?

2.10. Query the serial number of a disk

Serial numbers can be used as unique IDs for cloud disks and local disks in an operating system to identify and distinguish different disks. This topic describes how to query the serial numbers of disks in Linux and Windows operating systems.

Limits

When you query serial numbers, take note of the following limits:

- You can guery the serial numbers of disks that are created only on or after June 10, 2020.
- You can query the serial numbers of disks only on I/O optimized instances.
- You can query the serial numbers of cloud disks, including enhanced SSDs (ESSDs), SSDs, ultra disks, and basic disks.
- You can query the serial numbers of local disks, including NVMe SSDs and SATA HDDs.

Query the serial number of a disk in Linux

1.

2. Run the lsblk command to query the device names of the Elastic Block Storage (EBS) devices that are attached to an instance.

```
[root@ecs ~]# lsblk

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT

vdb 254:16 0 20G 0 disk

vda 254:0 0 40G 0 disk
```

3. Run the following command to guery the serial number of an EBS device:

```
udevadm info --query=all --name= Device name | grep ID_SERIAL
```

For example, you can use the following command to query the serial number of the /dev/vda device:

```
[root@ecs ~]# udevadm info --query=all --name=/dev/vda | grep ID_SERIAL
E: ID_SERIAL=bpld4foh3ef8bntl****
```

Query the serial number of a disk in Windows Server 2012 or later

1.

2. Open Windows PowerShell.

Open the Start menu, enter PowerShell in the search field, and then click Windows PowerShell.

3. Run the following command to query the disk corresponding to a logical volume (LV):

```
Get-CimInstance -ClassName Win32_LogicalDiskToPartition |select Antecedent, Dependent | fl
```

Example:

```
PS C:\Users\Administrator> Get-CimInstance -ClassName Win32_LogicalDiskToPartition |sel ect Antecedent, Dependent |fl
Antecedent: Win32_DiskPartition (DeviceID = "Disk #0, Partition #0")
Dependent: Win32_LogicalDisk (DeviceID = "C:")
Antecedent: Win32_DiskPartition (DeviceID = "Disk #1, Partition #0")
Dependent: Win32_LogicalDisk (DeviceID = "D:")
```

4. Run the following command to query the serial number of the disk:

```
Get-Disk |select Number, SerialNumber
```

Example:

```
PS C:\Users\Administrator> Get-Disk |select Number, SerialNumber

Number SerialNumber

-----

1 bp14dzwwr539hzqi****

0 bp16htuqeqnvlee8****
```

Query the serial number of a disk in an operating system earlier than Windows Server 2012

1.

2. Open Windows PowerShell.

Open the Start menu, enter PowerShell in the search field, and then click Windows PowerShell.

3. Run the following command to query the disk corresponding to an LV:

```
Get-WmiObject -Class Win32_LogicalDiskToPartition |select Antecedent, Dependent |fl
```

Example:

```
PS C:\Users\Administrator> Get-WmiObject -Class Win32_LogicalDiskToPartition |select An tecedent, Dependent |fl
Antecedent: \\ecs\root\cimv2:\Win32_DiskPartition.DeviceID="Disk #0, Partition #0"
Dependent: \\ecs\root\cimv2:\Win32_LogicalDisk.DeviceID="C:"
```

4. Run the following command to query the serial number of the disk:

```
Get-WmiObject -Class Win32_PhysicalMedia |select Tag, Serialnumber
```

Example:

2.11. Change the billing methods of a disk

After you create a disk, you can change the billing method of the disk between subscription and payas-you-go. This topic describes the methods to change billing methods and the applicable scope and effective time of each method.

Background information

The billing methods of disks are related to how the disks are created:

• If a disk is created together with a subscription instance or if a subscription disk is separately created on the Disks page, you must pay for the disk before you can use the disk. For more information, see Subscription.

Note Subscription disks can be separately created but must be attached to subscription instances before you use the disks.

• If a disk is created together with a pay-as-you-go instance or if a pay-as-you-go disk is separately created on the Disks page, the billing method of the disk is pay-as-you-go. For more information, see Pay-as-you-go.

Methods

The following table describes different operations that you can perform to change the billing methods of disks.

Direction of change	Operation	Applicable scope	Time to take effect
Subscription to	Change the billing method of a disk	This method can be used for subscription data disks attached to subscription instances. This method cannot be used to change the billing methods of system disks.	From the next billing cycle
pay-as-you-go	Change the billing method of an instance from subscription to pay-as-you-go	This method can be used for system disks and data disks attached to subscription instances. The billing methods of the disks vary based on the billing methods of the instances.	Immediately
Pay-as-you-go to	Change the billing method of a disk	This method can be used for pay-as- you-go data disks attached to subscription instances. This method cannot be used to change the billing methods of system disks.	Immediately
subscription	Change the billing method of an instance from pay-as-you-go to subscription	This method can be used for system disks and data disks attached to pay-as-you-go instances. The billing methods of the disks vary based on the billing methods of the instances.	Immediately

References

ModifyDiskChargeType

2.12. Detach or attach a system disk

You can detach a system disk from an Elastic Compute Service (ECS) instance. If an instance cannot start because its system disk contains corrupted files, you can detach the disk from the instance and attach it to another instance as a data disk for repair. After the disk is repaired, you can attach it back to the original instance.

Prerequisites

Before you detach a system disk from an instance, make sure that the involved resources meet the following requirements:

- The instance is in the **Stopped** state.
- The image used to create the instance is not deleted.
- The image used by the instance is not an Alibaba Cloud Market place image or a custom image derived from an Alibaba Cloud Market place image.

Before you attach a detached system disk to an instance, make sure that the instance meets the following requirements:

- The instance is in the No System Disk state.
- The instance is the instance from which the system disk was detached.

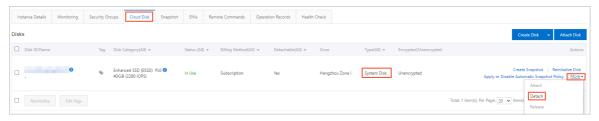
Context

After a system disk is detached from an instance, the instance and the system disk do not support the following operations:

- Start the instance
- Change the billing method of the instance
- Change the instance type
- Change the public bandwidth
- Create a custom image
- Separately release the system disk
- Replace the system disk
- Resize the system disk
- Change the billing method of the disk

Detach a system disk

- 1.
- 2.
- 3.
- 4
- 5. Find the instance from which you want to detach the system disk and make sure that the instance is in the **Stopped** state. Click the instance ID to go to the **Instance Details** page.
- 6. On the Instance Details page, click the Cloud Disk tab.
- 7. In the Actions column corresponding to the system disk, choose More > Detach.



8. In the Detach Disk message, click **OK**.

Note If you do not select Release Disk with Instance in the message, the system disk is retained as a data disk when its attached instance is released. The disk is then billed on a pay-as-you-go basis.

- 9. Check whether the system disk is detached.
 - i. On the Instances page, the instance is in the No System Disk state and has a tag added.



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

iii. The system disk is converted to a data disk, the disk has three tags added, and the billing method remains unchanged.



You can use the tags that are added to the instance and the disk to search for resources. The following table describes the tags.

Tag key	Description	Tag value
acs:ecs:sourceSystem Diskld	The system disk of the instance.	The ID of the disk.
acs:ecs:sourceInstanc eld	The instance to which the system disk was attached.	The ID of the instance.
acs:ecs:diskDeleteProt ection	The release protection feature is enabled for the disk. The disk cannot be manually released.	true
acs:ecs:diskPayType	The billing method of the disk.	The same billing method of the disk before it is detached. Valid values: • Prepaid: subscription • AfterPay: pay-as-you-go

Attach as a system disk

- 1.
- 2.
- 3.
- 4.
- 5. Find the instance that is in the **No System Disk** state and click **Attach Disk** in the **Actions** column.



- 6. In the Attach Disk dialog box, perform the following operations:
 - i. Copy the value of $Source\ System\ Disk\ and\ paste$ it to the $Target\ Disk\ field$.

ii. Configure the parameters described in the following table.

Parameter	Description	Example
Logon Credentials	 Key Pair: Select this option only for Linux instances. Custom Password: You can select this option for both Windows Server and Linux instances. However, this option takes effect only when the default administrator and root usernames are used. 	Key Pair
Release Disk with Instance	Specifies whether to retain the disk when the instance is released. If this option is selected, the disk is released when the instance is released. If this option is not selected, the disk is retained as a data disk and billed on a payas-you-go basis.	Do not select
Delete Automatic Snapshots While Releasing Disk	 Specifies whether to retain the automatic snapshots of the disk when the disk is released. If this option is selected, the automatic snapshots are deleted when the disk is released. If this option is not selected, the automatic snapshots are retained when the disk is released. 	Select

- iii. Click OK.
- iv. Confirm the settings and click Attach.

Attach as a data disk

After a system disk is detached, you can attach it to another instance within the same zone as a data disk. For more information, see Attach a data disk.

Related information

- Det achDisk
- AttachDisk

2.13. Edit the tags of a disk

Tags can be used to identify resources with the same characteristics (such as disks that belong to the same organization or that serve the same purpose) for easy search and management. This topic describes how to edit the tags of an existing disk.

Context

For information about how to use tags, the resources that support tags, and the limits on tags, see Overview and the "Tag limits" section of the Limits topic.

Procedure

- 1.
- 2.
- 3.
- 4. Find the disk whose tags you want to edit, move the pointer over the icon in the Tag column, and then click Edit Tags.
- 5. In the **Edit Tags** dialog box, click Available Tags to select existing tags or click Create to create tags. Then, click **OK**.

What's next

After tags are added to your disks, you can filter the disks by tag to perform different operations. For example, you can create snapshots for disks that have a set of tags and calculate the costs for disks that have a different set of tags.

2.14. Detach a data disk

You can detach a data disk from an Elastic Compute Service (ECS) instance when the disk is no longer needed. You must detach a data disk from an instance before you can attach the disk to another instance within the same zone. This topic describes how to detach a data disk.

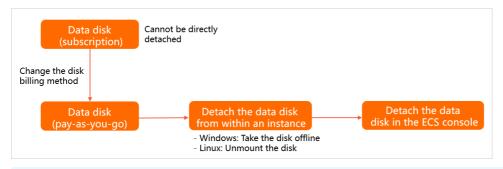
Prerequisites

Before you detach a data disk, make sure that the following conditions are met:

- The disk is attached to an instance and **In Use** is displayed in the **Status** column that corresponds to the disk in the ECS console.
- To avoid data loss and ensure data integrity, we recommend that you stop read and write operations on the disk before you detach it.

Context

The following figure shows the procedure to detach a data disk.



Note Local disks that are used as data disks cannot be detached.

1. If the data disk that you want to detach uses the subscription billing method, change its billing method to pay-as-you-go.

For more information, see Step 1: Change the billing method of the data disk from subscription to pay-as-you-go.

2. If file systems are mounted to the partitions of the data disk, detach the data disk from within the operating system of the instance to which the data disk is attached.

For more information, see Step 2: Detach the data disk from within the operating system of the instance to which the disk is attached.

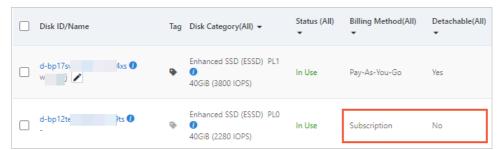
3 Detach the data disk in the ECS console

For more information, see Step 3: Detach the data disk in the ECS console.

Step 1: Change the billing method of the data disk from subscription to pay-as-you-go

If your data disk uses the pay-as-you-go billing method, skip this step.

If your data disk uses the subscription billing method, the data disk cannot be detached, shown as in the following figure. You must change the billing method of your data disk to pay-as-you-go. For more information, see Change the billing method from subscription to pay-as-you-go.



Step 2: Detach the data disk from within the operating system of the instance to which the disk is attached

If the data disk is partitioned and has file systems mounted, perform the following operations to detach the data disk from within the operating system of the instance to which the disk is attached:

For a Linux instance, perform the following steps:

1. Connect to the Linux instance.

For more information, see Connect to a Linux instance by using a password.

2. Run the following command to view the mount information of the data disk:

```
df -h
```

A command output similar to the following one is displayed. In this example, the /dev/vdb1 partition of the data disk is used. In actual scenarios, choose the data disk partitions whose mount information you want to view.

```
[root@ecs ]# df -h
Filesystem Size Used Avail Use% Mounted on
levtmpfs 441M 0 441M 0% /dev
tmpfs 459M 0 459M 0% /dev/shm
tmpfs 459M 468K 459M 1% /run
tmpfs 459M 0 459M 0% /sys/fs/cgroup
/dev/vdal 40G 2.6G 35G 7% /
dev/vdal 40G 49M 38G 1% /mnt
tmpfs 92M 0 92M 0% /run/user/0
```

3. Run the umount command to unmount the file systems from the data disk partitions.

For example, you can run the following command to unmount the file system from the /dev/vdb1 partition of the data disk:

```
umount /dev/vdb1
```

4. Run the following command to view the UUID information of the data disk partitions:

```
blkid
```

The following command output shows the UUID information of the /dev/vdb1 partition of the data disk.

5. Run the following command to check whether the /etc/fstab file contains the automatic mount configuration information of the file systems in the data disk partitions:

```
cat /etc/fstab
```

Find the UUID information obtained in the previous step in the command output. The file system mounted to the /dev/vdb1 partition is configured in /etc/fstab, as shown in the following figure.

- 6. Delete the automatic mount configuration information of the file systems in the data disk partitions from /etc/fstab.
 - **Note** If you do not delete the automatic mount configuration information of the file systems in the data disk partitions from /etc/fstab, the instance cannot be restarted after you detach the data disk from the instance in the ECS console.
 - i. Run the following command to edit /etc/fstab:

```
vim /etc/fstab
```

ii. Press the I key to enter the edit mode.

iii. Delete or comment out the automatic mount configuration information of the file systems in the data disk partitions.

For example, you can add # before the line of the automatic mount configuration information to comment it out.

iv. Press the Esc key, enter :wq, and then press the Enter key to save your changes and exit the edit mode.

For a Windows instance, perform the following steps:

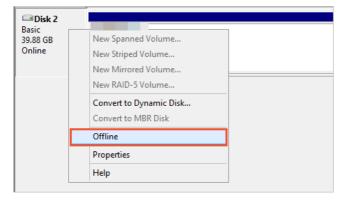
- ? Note In this example, Windows Server 2012 R2 is used.
- 1. Connect to the Windows instance.

For more information, see Connect to a Windows instance by using a password.

2. On the Windows Server desktop, click the Server Manager icon in the lower-left corner.



- 3. In the upper-right corner of the **Server Manager** window, choose **Tools** > **Computer Management**.
- 4. In the left-side navigation pane, choose Computer Management (Local) > Storage > Disk Management.
- 5. Right-click the disk that you want to detach and select Offline.



Step 3: Detach the data disk in the ECS console

You can detach the data disk on the Instances or Disks page. Perform the following operations to detach the data disk on the Instances page in the ECS console:

1.

- 2.
- 3.
- 4. Find the instance from which you want to detach the data disk and click the instance ID.
- 5. On the Instance Details page, click the Cloud Disk tab.
- 6. Find the data disk that you want to detach and choose More > Detach in the Actions column.
 - ? Note If No is displayed in the Detachable(All) column, the data disk cannot be detached. In this case, go to Step 1: Change the billing method of the data disk from subscription to pay-as-you-go.
- 7. In the Detach Disk dialog box, click **OK**.

You can also choose Storage & Snapshots > Disks to go to the Disks page to detach the data disk.

Result

After you perform the preceding steps to detach the data disk, you can choose **Storage & Snapshots > Disks** to find the disk on the Disks page. If **Unattached** is displayed in the **Status(All)** column, the disk is detached.

What's next

- You can attach the data disk to another instance within the same zone. For more information, see Attach a data disk.
- When the data disk is no longer needed, you can back up data stored on the disk and then release the disk. For more information, see Create a snapshot of a disk and Release a disk.

Related information

Det achDisk

2.15. Release a disk

You can manually release disks that you no longer need. When a disk is released, all data stored on the disk is also released and billing for the disk is stopped. By default, disks attached to Elastic Compute Service (ECS) instances are released along with the instances. You can disable the Release Disk with Instance feature for disks to retain them. This topic describes how to release a disk in the ECS console and how to enable or disable the Release Disk with Instance feature.

Prerequisites

- A snapshot is created for the disk that you want to manage to back up disk data. For more information, see Create a snapshot of a disk.
- The disk that you want to manually release or the pay-as-you-go data disk for which you want to enable or disable the Release Disk with Instance feature is in the **Unattached** state. For more information, see Detach a data disk.

Context

• If the Delete Automatic Snapshots While Releasing Disk feature is enabled for a disk, the automatic snapshots of the disk are deleted when the disk is released. You can disable the Delete Automatic Snapshots While Releasing Disk feature for a disk when you attach the disk to an instance. For more

information, see Delete automatic snapshots while releasing a disk.

- Manual snapshots are not released when their associated disks are released.
- You can manually release a disk or enable the Release Disk with Instance feature for the disk. The following table describes how to release different types of disks.

Disk type	Subscription	Pay-as-you-go
System disk	By default, the Release Disk with Instance feature is enabled for subscription system disks. The feature can be disabled for the disks.	By default, the Release Disk with Instance feature is enabled for pay-as-you-go system disks. The feature can be disabled for the disks.
Dat a disk	 By default, the Release Disk with Instance feature is enabled for subscription data disks. The feature can be disabled for the disks. Subscription data disks can be released only after you convert them into payas-you-go disks. 	 Pay-as-you-go data disks can be released along with their associated instances. By default, the Release Disk with Instance feature is enabled for pay-as-you-go data disks that are created together with instances. For separately created pay-as-you-go data disks, you must manually enable the feature. Pay-as-you-go data disks can be manually released.

- When you enable or disable the Release Disk with Instance feature for a system disk or a data disk, take note of the following items:
 - If the Release Disk with Instance feature is enabled for a disk, the disk is automatically released when its associated instance is released.
 - If the Release Disk with Instance feature is disabled for a disk, the disk is retained as a pay-as-you-go data disk 15 days after its associated instance expires, 15 days after a payment becomes overdue for the instance, or when the instance is manually released. If you create disks within regions inside mainland China, you must complete real-name verification for your account to ensure that the disks can be retained.

? Note The retained disks are billed on a pay-as-you-go basis. You can log on to the Billing Management console and view consumption details by disk ID.

Manually release a disk

Perform the following operations to manually release a pay-as-you-go data disk:

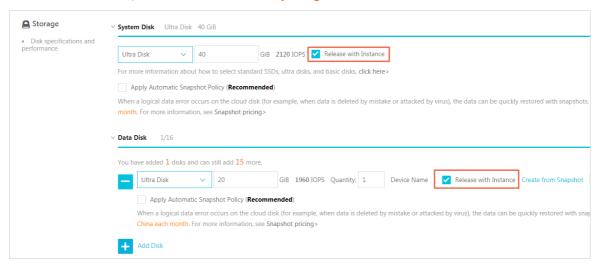
- 1.
- 2.
- 3.
- 4. Find the disk that you want to release and choose More > Release in the Actions column.
- 5. In the Release Disk message, confirm the information and click Confirm Release.

Enable or disable the Release Disk with Instance feature for disks when you create an instance

Perform the following operations to enable or disable the Release Disk with Instance feature for the system disk or data disks when you create an instance in the ECS console:

- 1.
- 2.
- 3. Click Create Instance.
- 4. In the Basic Configurations step, scroll down to the Storage section and optionally clear Release with Instance for the system disk or data disks.

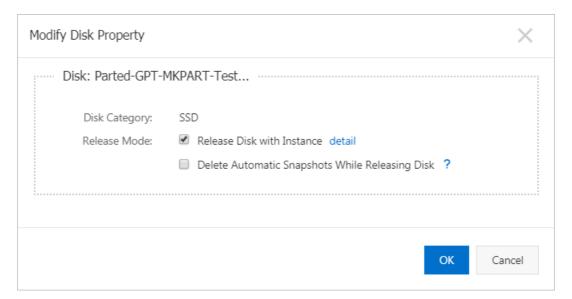
For more information, see Create an instance by using the wizard.



Enable or disable the Release Disk with Instance feature for a disk on the Disks page

Perform the following operations to enable or disable the Release Disk with Instance feature for a payas-you-go data disk on the Disks page in the ECS console:

- 1.
- 2.
- 3.
- 4. Find the instance for whose disk you want to enable or disable the feature and click the instance ID to go to the **Instance Details** page.
- 5. On the Instance Details page, click the Cloud Disk tab.
- 6. Find the disk for which you want to enable or disable the feature and choose **More > Modify Disk**Property in the Actions column.
- 7. In the Modify Disk Property dialog box, select or clear Release Disk with Instance and click OK.
 - ? Note After you clear Release Disk with Instance for a data disk, the disk is retained as a pay-as-you-go data disk 15 days after the associated instance expires, 15 days after a payment becomes overdue for the instance, or when the instance is manually released.



Related information

• DeleteDisk

3.Encrypt disks 3.1. Encrypt a system disk

When you create Elastic Compute Service (ECS) instances or copy custom images by using the ECS console or by calling API operations, you can encrypt the system disks. After the system disks are encrypted, data stored on the disks is also encrypted. Encryption keys can be customer master keys (CMKs) that are created by Key Management Service (KMS) or Bring Your Own Key (BYOK) keys (custom keys imported by using the BYOK feature).

Context

Methods of encrypting system disks

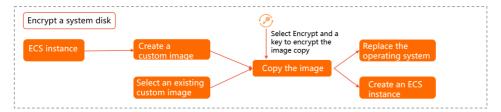
You can use one of the following methods to encrypt a system disk:

When you create an ECS instance, select Disk Encryption and select a key in the Storage section to
encrypt the system disk. For more information, see Encrypt the system disk of an instance when you
create the instance.

Note You cannot encrypt system disks when you create instances in Hong Kong Zone D or Singapore Zone A.

When you copy a custom image, select Encrypt and select a key to encrypt the custom image copy. If
an ECS instance is created from the encrypted custom image copy, the system disk and data disks of
the instance are automatically encrypted. For more information, see Create an encrypted system disk
when you copy a custom image.

The following figure shows how to create an encrypted system disk by using the Copy and Encrypt feature. The Copy and Encrypt feature allows you to encrypt the image copy when you copy a custom image. For more information, see 加密概述.



Scenarios for encrypting a system disk

The following table describes the different scenarios for encrypting a system disk.

Encrypted custom image	System disk encrypted when the instance is created	System disk encrypted
No	No	No
No	Yes (Use Key A) For more information, see Encrypt the system disk of an instance when you create the instance.	Yes (Use Key A)

Encrypted custom image	System disk encrypted when the instance is created	System disk encrypted
Yes (Use Key B) For more information, see Create an encrypted system disk when you copy a custom image.	No	Yes (Use Key B)
Yes (Use Key B) For more information, see Create an encrypted system disk when you copy a custom image.	Yes (Use Key A) For more information, see Encrypt the system disk of an instance when you create the instance.	Yes (Use Key A)

Encrypt the system disk of an instance when you create the instance

You can select Disk Encryption and select a key to encrypt the system disk of an instance when you create the instance.

Requirements

If you want to encrypt the system disk of an instance when you create the instance, the instance must meet the following requirements:

- The instance family of the instance cannot be ecs.ebmg5, ecs.ebmgn5t, ecs.ebmi3, ecs.sccg5, ecs.scch5, ecs.ebmc4, and ecs.ebmhfg5. For more information, see Instance family
- The instance uses only a public or custom image. The instance cannot use an Alibaba Cloud Market place image or a shared image.
- The instance uses only enhanced SSDs (ESSDs).

Create an ECS instance

- 1.
- 2.
- 3.
- 4. On the Instances page, click Create Instance.
- 5. In the Basic Configurations step, perform the following operations.

Note This step describes how to configure the encryption settings of the system disk when you create an instance. For information about other instance configurations, see Create an instance by using the wizard.

- i. Select an appropriate region and zone in the Region section and select an instance type in the Instance Type section.
- ii. Select Public Image or Custom Image in the Image section.
- iii. Select Enhanced SSD (ESSD) and specify a capacity for the system disk in the Storage section.

iv. Select Disk Encryption and select a key from the drop-down list in the Storage section.



By default, Alibaba Cloud uses the **Default Service CMK** managed service key when you select **Disk Encryption** for a system disk. You can also specify the BYOK key that you created in KMS as the encryption key of the system disk. We recommend that you use a BYOK key as the encryption key. For information about how to create a custom CMK, see Create a CMK.

After the system disk is encrypted, a tag is automatically added to the KMS key used to encrypt the disk. The key of the tag is <code>acs:ecs:disk-encryption</code>, and the value of the tag is <code>true</code>. You can log on to the KMS console and click the key ID to view the tag of the key.

Create an encrypted system disk when you copy a custom image

When you copy a custom image, you can select Encrypt to encrypt the custom image copy. This way, system and data disks created from the encrypted custom image copy are automatically encrypted.

Encrypt a custom image

You can encrypt a custom image by using the ECS console or by calling the Copylmage operation.

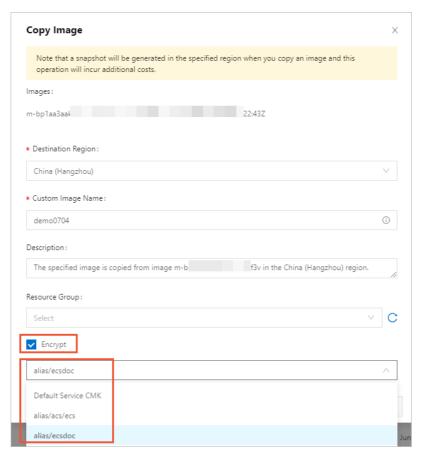
• Encrypt a custom image when you copy an image in the ECS console

This section describes how to copy an existing custom image and encrypt the image copy in the ECS console. Then, you can create an encrypted system disk from the encrypted image copy. If no custom images are available, create a custom image. For more information, see Create a custom image from a snapshot and Create a custom image from an instance.

- i. Log on to the ECS console.
- ii. In the left-side navigation pane, choose Instances & Images > Images.
- iii. In the upper-left corner of the top navigation bar, select a region.
- iv. On the Images page, click the Custom Images tab.
- v. Find the custom image that you want to copy and click **Copy Image** in the **Actions** column.

Note If the size of the custom image is larger than 500 GiB, follow on-screen tips to submit a ticket after you click **Copy Image**.

vi. In the Copy Image dialog box, select Encrypt and select a key from the drop-down list.



By default, Alibaba Cloud uses the **Default Service CMK** managed service key when you select **Disk Encryption** for a system disk. You can also specify the BYOK key that you created in KMS as the encryption key of the system disk. We recommend that you use a BYOK key as the encryption key. For information about how to create a custom CMK, see Create a CMK.

Note The first time that you select Encrypt, click Go to Authorize and follow on-screen tips to attach AliyunECSDiskEncrypt Default Role to allow ECS to access your KMS resources. This step describes only how to configure the encryption settings when you copy a custom image. For more information about other configurations, see Copy a custom image.

vii. Click OK.

You can create an encrypted system disk from the encrypted image copy. After the system disk is encrypted, a tag is automatically added to the KMS key used to encrypt the disk. The key of the tag is acs:ecs:disk-encryption, and the value of the tag is true. You can log on to the KMS console and click the key ID to view the tag of the key.

• Encrypt a custom image by calling the Copylmage operation

In the following example, Alibaba Cloud CLI is used to call the Copylmage operation and specify KMSKeyld to copy a custom image and encrypt the image copy. You can create an encrypted system disk from the encrypted image copy.

```
aliyun ecs CopyImage --RegionId cn-hongkong \
--ImageId m-bp155shrycg3s0****** --DestinationRegionId cn-shenzhen \
--Encrypted true --KMSKeyId e522b26d-abf6-4e0d-b5da-04b7*****3c \
--Tag.N.Key EcsDocumentation
```

Use an encrypted custom image to create an ECS instance

After a custom image is encrypted, the system and data disks of an instance that is created from the image are automatically encrypted. The system and data disks use the same encryption key as the custom image. For information about how to create an ECS instance, see Create an instance by using the wizard.

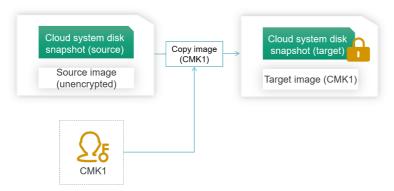
Encryption state change of a system disk

After you copy a custom image and encrypt the custom image copy, the encryption state of the system disks created from the image copy is determined by whether a new CMK is selected during the image copy process. The following section describes how a CMK affects the encryption state of a system disk.

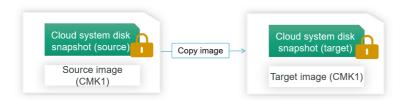
• If you do not select a CMK when you copy an unencrypted custom image, the system disks created from the image copy are unencrypted.



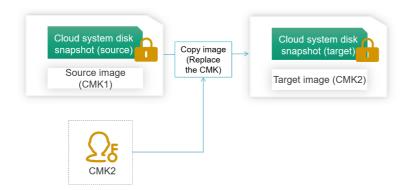
• If you select a CMK when you copy an unencrypted custom image, the image copy is encrypted. To access instances created from the image copy, you must use this CMK.



• If you do not select a CMK when you copy an encrypted image, the image copy is encrypted by using the same encryption key as the original image. To access instances created from the image copy, you must use the encryption key of the original image.



• If you select a new CMK when you copy an encrypted image, the image copy is encrypted by using the new CMK. To access instances created from the image copy, you must use the new CMK.



What's next

You can use the encrypted image copy to create an instance or replace the system disk of an instance as described in the following topics:

- Create an ECS instance by using a custom image
- Replace the operating system of an instance by using a non-public image

Related information

- Copylmage
- CancelCopyImage
- Copy a snapshot

3.2. Encrypt a data disk

This topic describes how to encrypt a data disk. After a data disk is encrypted, both data in transit and data at rest on the disk are encrypted. You can follow the instructions described in this topic to encrypt data disks to meet security compliance requirements.

Context

You can use one of the following methods to encrypt data disks:

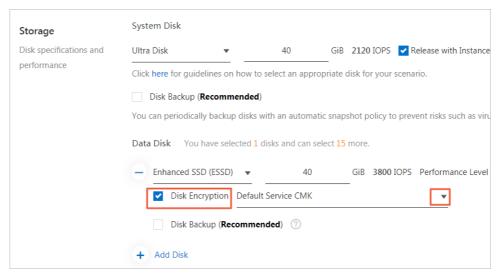
- Create an Elastic Compute Service (ECS) instance from an encrypted custom image that contains data of data disks. The data disks of the instance that were created from the custom image are encrypted. For more information, see Encrypt a system disk.
- When you create an instance, click Add Disk to add data disks and select Disk Encryption and a key for each added data disk. For more information, see Encrypt a data disk when you create an ECS instance.
- When you create an independent disk, select Disk Encryption and a key for each disk. For more information, see Encrypt a data disk when you create the disk.

When you encrypt data disks, you must use the keys in Key Management Service (KMS). For more information, see 加密概述.

Encrypt a data disk when you create an ECS instance

- 1.
- 2.
- 3.
- 4. On the Instances page, click Create Instance.

- 5. In the Basic Configurations step, find the Storage section and perform the following steps.
 - Note This step describes only how to configure the encryption settings when you create an instance. For information about other configurations of the instance, see Create an instance by using the wizard.
 - i. Click Add Disk.
 - ii. Specify the category and capacity of the data disk.
 - iii. Select Disk Encryption and select a key from the drop-down list.



After the data disk is encrypted, a tag is automatically added to the specified KMS key used to encrypt the data disk. The key of the tag is acs:ecs:disk-encryption, and the value of the tag is true. You can log on to the KMS console and click the key ID to view the tag of the key.

Encrypt a data disk when you create the disk

- 1.
- 2.
- 3.
- 4. On the Disks page, click Create Disk.
- 5. Specify the category and capacity of the data disk.
 - **? Note** This step describes only how to configure the encryption settings when you create a disk. For information about other configurations of the disk, see Create a disk.
- 6. In the Storage section, select Disk Encryption and select a key from the drop-down list.



After the data disk is encrypted, a tag is automatically added to the specified KMS key used to encrypt the data disk. The key of the tag is acs:ecs:disk-encryption, and the value of the tag is true. You can log on to the KMS console and click the key ID to view the tag of the key.

Change the encryption state

After a data disk of an instance is created, you cannot change its encryption state. If you want to change the encryption state of its data, perform the procedures described in the following table.

State change	Procedure	Windows Server	Linux
From unencrypted to encrypted	 Log on to the operating system of the instance. For more information, see Connection methodsGuidelines on instance connection. Manually copy data from an unencrypted disk to a new encrypted disk. 	run the ronocony	Run the rsync shell command.
From encrypted to unencrypted	 Log on to the operating system of the instance. For more information, see Connection methodsGuidelines on instance connection. Manually copy data from an encrypted disk to a new unencrypted disk. 	command.	Communic.

Related information

- RunInstances
- CreateDisk

4.Resize cloud disks 4.1. Overview

You can resize disks to meet increasing storage requirements as your business and application data grow. This topic describes the maximum sizes to which disks can be resized and how to resize disks in different scenarios.

Scenarios

You can use the following methods to increase the storage capacity of an instance:

• Resize an existing disk. You must resize the existing partitions of the disk or create new partitions on the disk.

The following table describes the two methods for resizing an existing disk.

Method	Limit	References
Resize a disk online	The instance to which the disk is attached must be in the Running (<i>Running</i>) state. For more information about limits, see the topics in the References column. After you resize the disk, the new size automatically takes effect without the need to restart the instance.	 Resize disks online for Linux instances Resize disks online for Windows instances
Resize a disk offline	The instance to which the disk is attached must be in the Running (<i>Running</i>) or Stopped (<i>Stopped</i>) state. For more information about limits, see the topics in the References column. After you resize the disk, you must restart the instance by using the Elastic Compute Service (ECS) console or by calling the RebootInstance operation for the new size to take effect.	 Resize disks offline for Linux instances Resize disks offline for Windows instances

- Create a disk, attach the disk to an ECS instance, and then partition and format the disk.
- Replace the system disk of an instance and specify a greater size for the new system disk. For more information, see Replace the system disk (public images).

Size ranges of resized system disks

The new size of a resized system disk must be greater than the original size but less than or equal to 500 GiB. The following table describes the size ranges of resized system disks that correspond to different images used by instances.

Image	Size range of a resized system disk (GiB)
CoreOS and FreeBSD	[Max{30, Original size of the system disk}, 500]

lmage	Size range of a resized system disk (GiB)	
Other Linux distributions	[Max{20, Original size of the system disk}, 500]	
Windows Server	[Max{40, Original size of the system disk}, 500]	

For example, the system disk on a CentOS instance is 35 GiB in size. When you resize this system disk, the specified new size must be greater than 35 GiB but less than or equal to 500 GiB.

Maximum sizes of resized data disks

The new size of a resized data disk must be greater than the original size. The following table describes the maximum sizes of resized data disks of different categories.

Disk category	Maximum size of a resized data disk (GiB)
Enhanced SSD (ESSD), standard SSD, or ultra disk	32,768
Basic disk	2,000

Billing method

For information about the billing methods of resized disks, see EBS devices.

4.2. Resize disks online for Linux instances

You can resize system and data disks to extend their capacities. This topic describes how to resize a disk for an Elastic Compute Service (ECS) Linux instance when the instance is running.

Prerequisites

The requirements described in the following table are met before you can resize a disk online for a Linux instance.

Resource	Requirement			
Instance	 The instance is I/O optimized. The public image used by the instance supports online resizing of disks. For information about images that support online resizing of disks, see Operating systems (images) that support online resizing of disks. Disks on instances of the ecs.ebmc4.8xlarge, ecs.ebmhfg5.2xlarge, and ecs.ebmg5.24xlarge instance types cannot be resized online. The instance is in the Running (Running) state. The version of the Linux kernel of the instance is 3.6.0 or later. You can run the uname -a command to check the kernel version. If the version of the Linux kernel is earlier than 3.6.0, you must resize the partitions of disks after you resize the disks. For more information, see Procedure for instances with kernels earlier than 3.6.0 and Resize partitions and file systems of Linux data disks. Note If your instance does not meet the preceding requirements on resizing disks online, you can resize its disks offline. For more information, see Resize disks offline for Linux instances. 			
Disk	 The disk is in the In Use (In Use) state. The disk is an enhanced SSD (ESSD), a standard SSD, or an ultra disk. After you renew a subscription instance and downgrade its configurations, the subscription disks of the instance cannot be resized for the remainder of the current billing cycle. When you resize a disk of a specific category, the specified new disk capacity cannot exceed the maximum single-disk capacity allowed for the disk category. For more information, see the "Elastic Block Storage (EBS) limits" section in Limits. Notice A disk in the master boot record (MBR) partition format cannot be resized to 2 TiB or larger. To resize an MBR disk to larger than 2 TiB, we recommend that you create a disk with the desired capacity, partition and format the new disk to the GUID Partition Table (GPT) format, and then copy data from the original MBR disk to the new GPT disk. For more information about how to partition and format disks to GPT, see Partition and format a data disk larger than 2 TiB in size. 			

Context

The following table describes the resources used in the examples of this topic.

Resource	Description
Image used by the instance	Alibaba Cloud Linux 2.1903 LTS 64-bit public image

Resource	Description
System disk	/dev/vda: uses the MBR partition format and the ext4 file system, and is resized from 40 GiB to 60 GiB.
Data disks	 /dev/vdb: uses the MBR partition format and the ext4 file system, and is resized from 40 GiB to 60 GiB. /dev/vdc: uses the GPT partition format and the XFS file system, and is resized from 40 GiB to 60 GiB.

Step 1: Create a snapshot

Before you resize a disk online, you must create a snapshot for the disk to back up the data stored on the disk.

- 1.
- 2.
- 3.
- 4.
- 5. Find the instance whose disk you want to resize and click the instance ID.
- 6. On the Instance Details page, click the Cloud Disk tab.
- 7. Find the disk that you want to resize and click Create Snapshot in the Actions column.
- 8. In the Create Disk Snapshot dialog box, enter a snapshot name, configure tags, and then click OK.
- 9. Click the **Snapshot** tab to view the snapshot.

After the snapshot is created, **100%** appears in the **Progress** column that corresponds to the snapshot. Then, you can proceed to perform subsequent operations.

Step 2: Resize the disk in the ECS console

- 1. On the Instance Details page, click the Cloud Disk tab.
- 2. Find the disk that you want to resize and choose More > Resize Disk in the Actions column.

To batch resize disks, log on to the ECS console by using your Alibaba Cloud account and choose **Storage & Snapshots > Disks** in the left-side navigation pane. On the Disks page, select the disks that you want to resize and click **Resize Disk** in the lower part of the page. Disks that are attached to the same instance cannot be resized together.

Note Alibaba Cloud accounts are permitted to batch resize disks but RAM users are not.
To batch resize disks, log on to the ECS console by using your Alibaba Cloud account.

- On the Resize Disks page, select Online Resizing and set the Size after Resize parameter.
 The specified Size after Resize value must be greater than the current capacity.
- 4. Verify the price. Read and select ECS Service Terms, and then click Confirm.
- 5. Read the notes, click I have read the notes. Resize, and then complete the payment.

Notice ■ Notice

- After you resize the disk in the ECS console, you must resize the partitions and file systems of the disk within the instance before you can use the new capacity of the disk.
- If you use Logical Volume Manager (LVM) to manage your disk partitions, you must use LVM
 to resize the partitions and file systems after you resize your disk in the ECS console. For
 more information, see Resize an LV by using LVM.

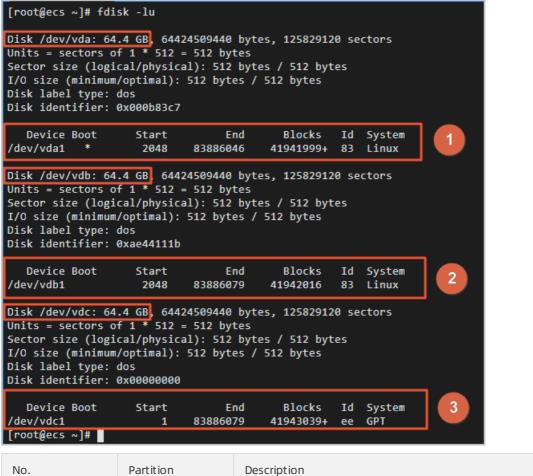
Step 3: View the disk partitions

You can log on to the instance to view the partition types (such as MBR and GPT) and file system types (such as ext4 and XFS) of the system disk and data disks. Subsequent resize operations vary based on the partition and file system types of the disks.

- 1. Log on to the instance. For more information, see Connect to a Linux instance by using a password.
- 2. Run the following command to view the disks attached to the instance:

```
fdisk -lu
```

The following command output shows the /dev/vda1 partition of the system disk and the /dev/vdb1 and /dev/vdc1 partitions of data disks.



No.	Partition	Description
①	/dev/vda1	A partition of the system disk. A value of Linux for System indicates that the partition is in the MBR format.
2	/dev/vdb1	A partition of a data disk. A value of Linux for System indicates that the partition is in the MBR format.
3	/dev/vdc1	A partition of a data disk. A value of GPT for System indicates that the partition is in the GPT format.

Note If the capacity of the system disk (/dev/vda) in the command output is still 40 GiB (Disk /dev/vda: 42.9 GB), the disk was not resized. We recommend that you restart the instance in the ECS console.

3. Run the following command to check the file system types of existing partitions:

```
df -Th
```

A command output similar to the following one is displayed. In the command output, the Mounted on column lists the mount points that correspond to the file systems.

```
[root@ecs ~]# df -Th
Filesystem
              Type
                        Size Used Avail Use% Mounted on
                                           0% /dev
devtmpfs
              devtmpfs
                        1.8G
                                 0 1.8G
tmpfs
              tmpfs
                        1.8G
                                 0 1.8G
                                           0% /dev/shm
                                           1% /run
tmpfs
               tmpfs
                              612K 1.8G
                        1.8G
              tmpfs
                        1.8G
                                           0% /sys/fs/cgroup
tmpfs
                                0 1.8G
/dev/vda1
                         40G 1.8G 36G
                                           5% /
              ext4
tmpfs
              tmpfs
                        361M
                                0 361M
                                           0% /run/user/0
/dev/vdb1
                                     38G
                                           1% /media/vdb
                         40G
                               49M
              ext4
/dev/vdc1
              xfs
                         40G
                               74M
                                     40G
                                           1% /media/vdc
[root@ecs ~]#
```

Step 4: Resize partitions

When you view the partitions of a resized disk, you may find that the partitions and file systems within the instance were not resized. This step describes how to resize the partitions of the resized disk within the instance.

1. Install the gdisk tool on the instance.

If the partitions are in the GPT format, you must perform this step. If the partitions are in the MBR format, skip this step.

```
yum install gdisk -y
```

- 2. Install the growpart tool.
 - If the instance runs an Alibaba Cloud Linux 2 or 3 operating system or a CentOS 7 or later operating system, run the following command:

```
yum install -y cloud-utils-growpart
```

- Note Cent OS 8 has reached its end of life (EOL). If your instance runs a Cent OS 8 operating system, change the Cent OS 8 repository address. For more information, see Change Cent OS 8 repository addresses.
- If the instance runs a Debian 8 or later operating system or an Ubuntu 14 or later operating system, run the following commands.

Run the following command to update the software repository:

```
apt-get update
```

Run the following command to install cloud-guest-utils:

```
apt-get install -y cloud-guest-utils
```

3. Run the following command to resize partitions:

```
growpart /dev/vda 1
```

In this example, the first partition of the system disk is resized. /dev/vda is the system disk and is the partition number. Use a space to separate /dev/vda from 1.

```
[root@ecs ~]# growpart /dev/vda 1
CHANGED: partition=1 start=2048 old: size=83883999 end=83886047 new: size=125827039 end=125829087
```

Modify the command accordingly to resize other partitions.

• Run the following command to resize the /dev/vdb data disk:

```
growpart /dev/vdb 1
```

• Run the following command to resize the /dev/vdc data disk:

```
growpart /dev/vdc 1
```

? Note

- If a single disk has successive partitions, you need to only resize the last partition to resize the disk. For example, the /dev/vdb data disk has three partitions named /dev/vdb1 , /dev/vdb2 , and /dev/vdb3 , you can run only the growpart /dev/vdb 3 command to resize the /dev/vdb data disk.
- When you run the growpart /dev/vda 1 command, the unexpected output in sfdisk
 --version [sfdisk, from util-linux 2.23.2] error message may appear. For information about how to troubleshoot this issue, see the "FAQ" section of this topic.

Step 5: Resize file systems

This step describes how to resize the file systems of partitions within the instance.

- 1. Resize the file systems within the instance based on the file system types that you obtained.
 - To resize ext* (such as ext4) file systems, run the following commands.

To resize the file system of the /dev/vda1 partition of the system disk, run the following command:

resize2fs /dev/vda1

To resize the file system of the /dev/vdb1 partition of the data disk, run the following command:

• To resize the xfs file system of the /dev/vdc1 partition of the data disk, run the following command:

xfs_growfs /media/vdc is the mount point of the /dev/vdcl partition. Replace it with the mount point of your partition.

• To resize the btrfs file system of a data disk, run the following command:

One /mountpoint is the mount point of the data disk. Replace it with the mount point of your data disk.

2. Run the following command to check the resize results:

```
df -Th
```

A command output similar to the following one is displayed.

[root@ecs ~]#	df -Th					
Filesystem	Type	Size	Used	Avail	Use%	Mounted on
devtmpfs	devtmpfs	1.8G	0	1.8G	0%	/dev
tmpfs	tmpfs	1.8G	0	1.8G	0%	/dev/shm
tmpfs	tmpfs	1.8G	492K	1.8G	1%	/run
tmpfs	tmpfs	1.8G	0	1.8G	0%	/sys/fs/cgroup
/dev/vda1	ext4	59G	1.9G	55G	4%	/
tmpfs	tmpfs	361M	0	361M	0%	/run/user/0
/dev/vdb1	ext4	59G	52M	56G	1%	/media/vdb
/dev/vdc1	xfs	60G	94M	60G	1%	/media/vdc

After you resize the file systems, check whether their actual sizes are consistent with the new specified values.

- If the file system sizes are increased to the new specified values and the business programs on the instance can run normally, the resize operation is complete.
- If the file system sizes are not increased to the new specified values, use the snapshot that you created in Step 1 to roll back the disk.

Operating systems (images) that support online resizing of disks

The following Linux public images and custom images derived from these public images support online resizing of disks:

- Alibaba Cloud Linux: Alibaba Cloud Linux 2 and Alibaba Cloud Linux 3
- CentOS:
 - o CentOS 6: CentOS 6.8 and later
 - o Cent OS 7: Cent OS 7.2 and later
 - o Cent OS 8 and later
- Red Hat Enterprise Linux (RHEL):
 - o RHEL 6: RHEL 6.9 and later
 - o RHEL 7: RHEL 7.4 and later
 - o RHEL 8 and later
- Ubuntu: Ubuntu 16 and later
- Debian: Debian 8 and later
- SUSE: SUSE 12 SP2 and later
- openSUSE: openSUSE 42.3 and later

FAO

• Question: When the growpart /dev/vda 1 command is run on the Linux instance, the unexpected output in sfdisk --version [sfdisk, from util-linux 2.23.2] error message appears. What do I do?

Answer:

- i. Run the locale command to check the character encoding type of the instance. If the character encoding type is not *en_US.UTF-8*, switch it to *en_US.UTF-8*.
 - a. Run the following command to switch the character encoding type:

```
LANG=en_US.UTF-8
```

b. If the error persists, run the following command to switch the character encoding type:

```
export LC ALL=en US.UTF-8
```

c. If the error still persists, run the following command to switch the character encoding type:

```
localectl set-locale LANG=en US.UTF-8
```

d. If the preceding solutions do not resolve the error, run the following command to switch the character encoding type:

```
export LANGUAGE=en_US.UTF-8
```

ii. If the error persists, run the reboot command to restart the instance.

Notice

If the partition is resized after you switch the character encoding type, we recommend that you switch it back to the original character encoding type.

• Question: When the growpart /dev/vda 1 command is run on the Linux instance, the -bash: growpart: command not found error message appears. What do I do?

Answer:

i. Run the uname -a command to check the version of the Linux kernel. The procedure described in this topic applies to Linux kernel 3.6.0 and later.

If the version of the Linux kernel is earlier than 3.6.0, you must resize the partitions of disks after you resize the disks. For more information, see Procedure for instances with kernels earlier than 3.6.0 and Resize partitions and file systems of Linux data disks.

- ii. Install the growpart tool.
 - If the instance runs a CentOS 7 or later operating system, run the following command:

```
yum install -y cloud-utils-growpart
```

Note CentOS 8 has reached its end of life (EOL). If your instance runs a CentOS 8 operating system, change the CentOS 8 repository address. For more information, see Change CentOS 8 repository addresses.

• If the instance runs a Debian 8 or later operating system or an Ubuntu 14 or later operating system, run the following command:

```
apt install -y cloud-guest-utils
```

 Question: The growpart tool used to resize partitions cannot be installed on an instance that runs CentOS 6.5. Why?

Answer: The Linux kernel version of CentOS 6 is earlier than 3.6.0. To install the growpart tool on CentOS 6, perform the following steps:

- i. Change the YUM repository address for Cent OS 6. For more information, see Change the Cent OS 6 source address.
 - ? Note CentOS 6 has reached its EOL. To use YUM to install CentOS software packages, you must change the YUM repository address for CentOS 6.
- ii. Install the dracut-modules-growroot tool on the instance that runs CentOS 6 to resize the partitions of disks. For more information, see Procedure for instances with kernels earlier than 3.6.0.

Other scenarios for disk resizing

- If you want to create partitions on a resized data disk, see Option 2: Create and format MBR partitions or Option 4: Create and format GPT partitions.
- If a raw data disk contains a file system but no partitions, see Option 5: Resize the file system of a raw data disk.

Related information

- •
- RebootInstance
- ResizeDisk
- AttachDisk

4.3. Resize disks offline for Linux instances

You can resize a system disk or data disk to extend its capacity. If your Elastic Compute Service (ECS) Linux instance does not support online resizing of disks, you can resize the disks on the instance offline. After you resize a disk on an instance offline, you must restart the instance for the change to take effect. The restart operation interrupts services that run on the instance for a short period of time. We recommend that you resize disks offline during off-peak hours.

Prerequisites

The requirements described in the following table must be met before you can resize a disk offline for a Linux instance.

Resource	Requirement
Instance	 The instance is in the Running (Running) or Stopped (Stopped) state.
Disk	•
	? Note

Context

Step 2: Resize the disk and start or restart the instance in the ECS console

- 1. On the Instance Details page, click the Cloud Disk tab.
- 2. Find the disk that you want to resize and choose More > Resize Disk in the Actions column.

If you want to batch resize disks, log on to the ECS console by using your Alibaba Cloud account and choose **Storage & Snapshots > Disks** in the left-side navigation pane. On the Disks page, select the disks that you want to resize and click **Resize Disk** in the lower part of the page. Disks that are attached to the same instance cannot be resized together.

Note Alibaba Cloud accounts are permitted to batch resize disks but Resource Access Management (RAM) users are not. Log on to the ECS console by using your Alibaba Cloud account.

- On the Resize Disks page, set the Size after Resize parameter.
 The specified Size after Resize value must be greater than the current disk capacity.
- 4. Verify the price, read and select ECS Service Terms, and then click Confirm.
- 5. Read the notes, click I have read the notes. Resize, and then complete the payment.

6. Start or restart the instance in the ECS console.



- If your instance is in the Running (Running) state, you must restart the instance in the ECS console. If your instance is in the Stopped (Stopped) state, you must start the instance in the ECS console.
- The disk resizing operation takes effect only after you start or restart the instance by using the ECS console or by calling the RebootInstance operation. If you restart the instance from within the operating system, the disk resizing operation does not take effect.
- The offline disk resizing operation takes effect only after you start or restart the instance. If you have resized a disk offline but cannot start or restart the instance due to business requirements, you can resize the disk online again to make the offline resizing operation take effect when the instance is in the Running state. You are charged for resizing the disk online. You must resize the disk based on your business requirements. For more information, see What do I do if I have resized a disk offline but do not want to restart the instance?.

The following procedure describes how to restart an instance.

- i. In the left-side navigation pane, choose Instances & Images > Instances.
- ii. On the Instances page, find the instance that you want to restart and choose **More** > **Instance Status** > **Restart** in the **Actions** column.
- iii. In the **Restart Instance** dialog box, click **OK**.

Related information

•

- RebootInstance
- ResizeDisk
- AttachDisk

4.4. Resize disks online for Windows instances

You can resize a system disk or a data disk to extend its capacity. This topic describes how to resize a disk for a Windows Elastic Compute Service (ECS) instance when the instance is running.

Prerequisites

The requirements described in the following table are met before you can resize a disk online for a Windows instance.

Resource	Requirement
Instance	 The instance is I/O optimized. The instance does not run Windows Server 2003. The instance is not an ecs.ebmc4.8xlarge, ecs.ebmhfg5.2xlarge, or ecs.ebmg5.24xlarge instance. The instance is in the Running (Running) state. The version of the Red Hat virtio-scsi driver is later than 58011. For more information about how to check the version of the Red Hat virtio-scsi driver and update the driver, see Update Red Hat virtio drivers of Windows instances
	Note If your instance does not meet the preceding requirements on resizing disks online, you can resize its disks offline. For more information, see Resize disks offline for Windows instances.
Disk	 The disk is in the In Use (In Use) state. The disk is an enhanced SSD (ESSD), a standard SSD, or an ultra disk. The disk is formatted with New Technology File System (NTFS) file systems. After you renew a subscription instance and downgrade its configurations, you cannot resize the subscription disks on the instance for the remainder of the current billing cycle. When you resize a disk of a specific category, the specified new disk capacity cannot exceed the maximum single-disk capacity allowed for the disk category. For more information, see the "Elastic Block Storage (EBS) limits" section in Limits.
	 Notice A disk in the master boot record (MBR) partition format cannot be resized to 2 TiB or larger. To resize an MBR disk to larger than 2 TiB, you can use one of the following methods: (Recommended) Create a disk larger than 2 TiB, partition and format the new disk to GUID Partition Table (GPT), and then copy data from the original MBR disk to the new GPT disk. For more information about how to partition and format a disk to GPT, see Partition and format a data disk larger than 2 TiB in size.
	 Convert the MBR partition to a GPT partition within the instance. When you convert the partition format, data may be lost. Proceed with caution. For more information, see Convert the partition format of a data disk on a Windows instance.

Context

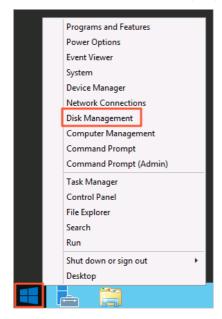
The following table describes the resources used in the examples of this topic.

Resource	Description
Image used by the instance	Windows Server 2012 R2 64-bit
System disk	The system disk is resized from 40 GiB to 60 GiB and is formatted with NTFS file systems.
Data disk	The data disk is resized from 40 GiB to 60 GiB and is formatted with NTFS file systems.

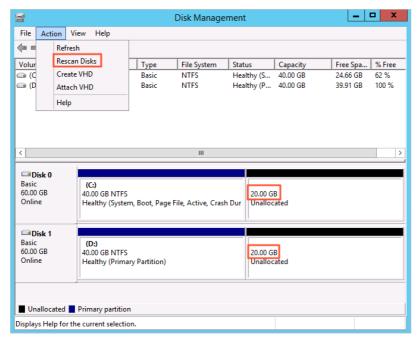
Step 3: Resize the file system of the system disk or data disk partition

After you resize a disk of an instance in the ECS console, the file systems of the disk partitions are not resized. You must connect to the instance to extend the file systems. The system and data disks are resized from 40 GiB to 60 GiB and are formatted with NTFS file systems. The following steps describes how to resize the file system of a system disk partition.

- 1. Log on to the instance. For more information, see Connect to a Windows instance by using a password or key.
- 2. On the Windows Server desktop, right-click the Start icon and select Disk Management.

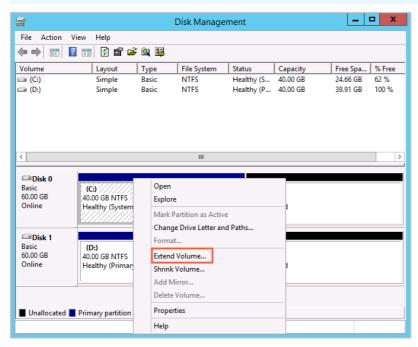


3. In the **Disk Management** dialog box, choose **Action > Rescan Disks** to view the unallocated disk capacity.



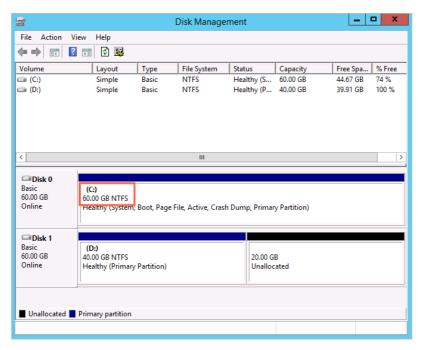
Disk 0 is the system disk and Disk 1 is a data disk.

- 4. Right-click the blank space in the Disk 0 section and select Extend Volume...
 - **? Note** If you resize the file system of a data disk partition, right-click the blank space in the data disk (such as **Disk 1**) section. If your disk has multiple partitions, we recommend that you resize the partition that is adjacent to the left of the unallocated partition. Otherwise, Windows changes the selected basic disk into a dynamic disk.



5. Extend the volume as instructed in **Extend Volume Wizard**.

The new space is automatically added to the original volume, as shown in the following figure.



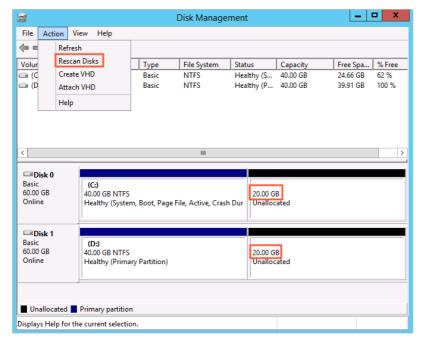
After the disk is resized, check whether the new capacity of the disk is consistent with the specified capacity.

- If the file system size is increased to the specified value and business programs on the instance can run normally, the file system is resized.
- If the file system size is not increased to the specified value, use the snapshot that you created in Step 1 to roll back the disk.

Step 4: (Optional) Create a partition on the data disk

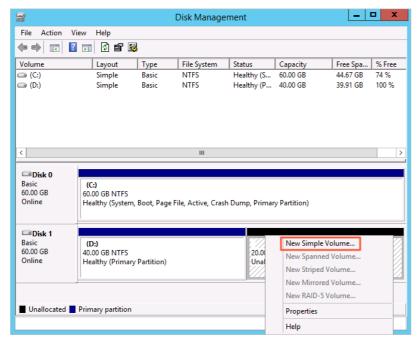
If you want to create a partition in the new disk space, perform the following steps:

- 1. Log on to the instance. For more information, see Connect to a Windows instance by using a password or key.
- 2. On the Windows Server desktop, right-click the Start icon and select Disk Management.
- 3. In the **Disk Management** dialog box, choose **Action > Rescan Disks** to view the unallocated disk capacity.



Disk 0 is the system disk and Disk 1 is a data disk.

4. Right-click the blank space in the Disk 1 section and select New Simple Volume...



5. Extend the volume as instructed in **New Simple Volume Wizard**. A partition is created in the new disk space.

Related information

- RebootInstance
- ResizeDisk
- AttachDisk

4.5. Resize disks offline for Windows instances

You can resize a system disk or a data disk to extend its capacity. If your Windows Elastic Compute Service (ECS) instance does not support online resizing of disks, you can resize the disks on the instance offline. After you resize a disk on an instance offline, you must restart the instance for the change to take effect. The restart operation interrupts services that run on the instance for a short period of time. We recommend that you resize disks offline during off-peak hours.

Prerequisites

The requirements described in the following table are met before you can resize a disk offline for a Windows instance.

Resource	Requirement
Instance	 The instance is in the Running (Running) or Stopped (Stopped) state. If the instance runs a Windows Server 2003 operating system, only data disks of the instance can be resized and the system disk cannot.

Resource	Requirement
	 The disk is in the In Use (In Use) state. The disk is an enhanced SSD (ESSD), a standard SSD, or an ultra disk. The disk is formatted with a New Technology File System (NTFS) file system. After you renew a subscription instance and downgrade its configurations, you cannot resize the subscription disks on the instance for the remainder of the current billing cycle. When you resize a disk of a specific category, the specified new disk capacity cannot exceed the maximum single-disk capacity allowed for the disk category. For more information, see the "Elastic Block Storage (EBS) limits" section in Limits.
Disk	 A disk in the master boot record (MBR) partition format cannot be resized to 2 TiB or larger. To resize an MBR disk to larger than 2 TiB, you can use one of the following methods: (Recommended) Create a disk larger than 2 TiB, partition and format the new disk to GUID Partition Table (GPT), and then copy data from the original MBR disk to the new GPT disk. For more information about how to partition and format a disk to GPT, see Partition and format a data disk larger than 2 TiB in size. Convert the MBR partition to a GPT partition within the instance. When you convert the partition format, data may be lost. Proceed with caution. For more information, see Convert the partition format of a data disk on a Windows instance.

Step 2: Resize the disk and start or restart the instance in the ECS console

- 1. On the Instance Details page, click the Cloud Disk tab.
- 2. Find the disk that you want to resize and choose More > Resize Disk in the Actions column.

If you want to batch resize disks, log on to the ECS console by using your Alibaba Cloud account and choose **Storage & Snapshots > Disks** in the left-side navigation pane. On the Disks page, select the disks that you want to resize and click **Resize Disk** in the lower part of the page. Disks that are attached to the same instance cannot be resized together.

Note Alibaba Cloud accounts are permitted to batch resize disks but Resource Access Management (RAM) users are not. Log on to the ECS console by using your Alibaba Cloud account.

- 3. On the **Resize Disks** page, set the **Size after Resize** parameter.
 - The specified Size after Resize value must be greater than the current disk capacity.
- 4. Verify the price, read and select ECS Service Terms, and then click Confirm.
- 5. Read the notes, click I have read the notes. Resize, and then complete the payment.
- 6. Start or restart the instance in the ECS console.

? Note

- If your instance is in the Running (Running) state, you must restart the instance in the ECS console. If your instance is in the Stopped (Stopped) state, you must start the instance in the ECS console.
- The disk resizing operation takes effect only after you start or restart the instance by using the ECS console or by calling the RebootInstance operation. If you restart the instance from within the operating system, the disk resizing operation does not take effect.
- The offline disk resizing operation takes effect only after you start or restart the instance. If you have resized a disk offline but cannot start or restart the instance due to business requirements, you can resize the disk online again to make the offline resizing operation take effect when the instance is in the Running state. You are charged for resizing the disk online. You must resize the disk based on your business requirements. For more information, see What do I do if I have resized a disk offline but do not want to restart the instance?.

The following procedure describes how to restart an instance.

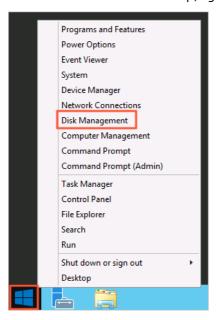
- i. In the left-side navigation pane, choose Instances & Images > Instances.
- ii. On the Instances page, find the instance that you want to restart and choose **More** > **Instance Status** > **Restart** in the **Actions** column.
- iii. In the Restart Instance dialog box, click OK.

Step 3: Resize the file system of the system disk or data disk partition

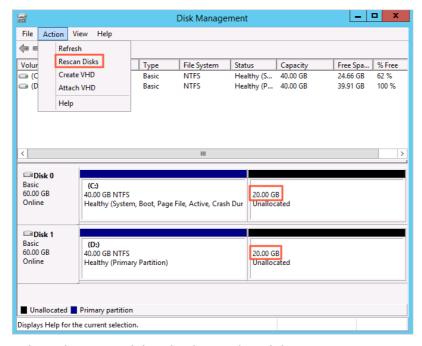
After you resize a disk of an instance in the ECS console, the file systems of the disk partitions are not resized. You must connect to the instance to extend the file systems. The system and data disks are resized from 40 GiB to 60 GiB and are formatted with NTFS file systems. The following steps describes how to resize the file system of a system disk partition.

1. Log on to the instance. For more information, see Connect to a Windows instance by using a password or key.

2. On the Windows Server desktop, right-click the Start icon and select Disk Management.

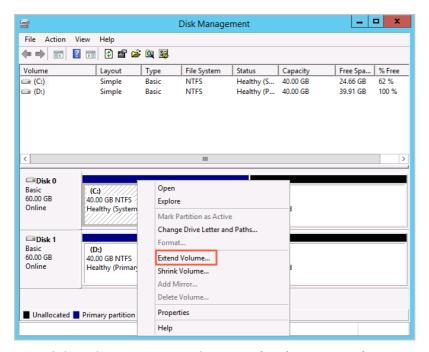


3. In the **Disk Management** dialog box, choose **Action > Rescan Disks** to view the unallocated disk capacity.



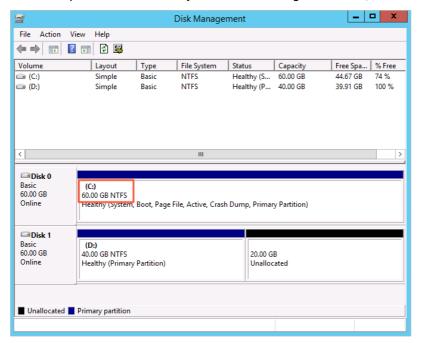
Disk 0 is the system disk and Disk 1 is a data disk.

- 4. Right-click the blank space in the **Disk 0** section and select **Extend Volume...**
 - Note If you resize the file system of a data disk partition, right-click the blank space in the data disk (such as Disk 1) section. If your disk has multiple partitions, we recommend that you resize the partition that is adjacent to the left of the unallocated partition. Otherwise, Windows changes the selected basic disk into a dynamic disk.



5. Extend the volume as instructed in **Extend Volume Wizard**.

The new space is automatically added to the original volume, as shown in the following figure.



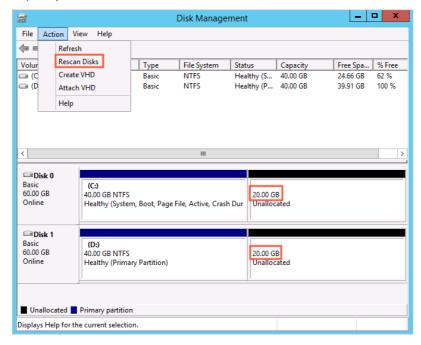
After the disk is resized, check whether the new capacity of the disk is consistent with the specified capacity.

- If the file system size is increased to the specified value and business programs on the instance can run normally, the file system is resized.
- If the file system size is not increased to the specified value, use the snapshot that you created in Step 1 to roll back the disk.

Step 4: (Optional) Create a partition on the data disk

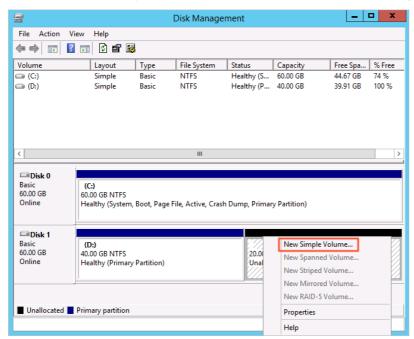
If you want to create a partition in the new disk space, perform the following steps:

- 1. Log on to the instance. For more information, see Connect to a Windows instance by using a password or key.
- 2. On the Windows Server desktop, right-click the Start icon and select Disk Management.
- 3. In the **Disk Management** dialog box, choose **Action > Rescan Disks** to view the unallocated disk capacity.



Disk 0 is the system disk and Disk 1 is a data disk.

4. Right-click the blank space in the Disk 1 section and select New Simple Volume...



5. Extend the volume as instructed in **New Simple Volume Wizard**. A partition is created in the new disk space.

Related information

- •
- RebootInstance
- ResizeDisk
- AttachDisk

4.6. Update Red Hat virtio drivers of Windows instances

You can resize the disks of a Windows Elastic Compute Service (ECS) instance online without the need to restart the instance. You can also query the serial numbers of disks from the operating system of the instance. If you want to resize the disks of a Windows instance created before March 30, 2019 or query the serial numbers of the disks attached to an instance created before May 1, 2020, perform the operations described in this topic to check whether the Red Hat virtio driver of the instance must be updated.

Context

- Red Hat virtio drivers are supported only in Windows Server 2008 and later.
- If an ECS instance has multiple data disks attached, the driver update process may take 1 to 2 minutes to complete.

Procedure

Perform the following operations to update the Red Hat virtio driver of a Windows instance:

- 1. Step 1: Check the version of the driver
- 2. Step 2: Update the driver

Step 1: Check the version of the driver

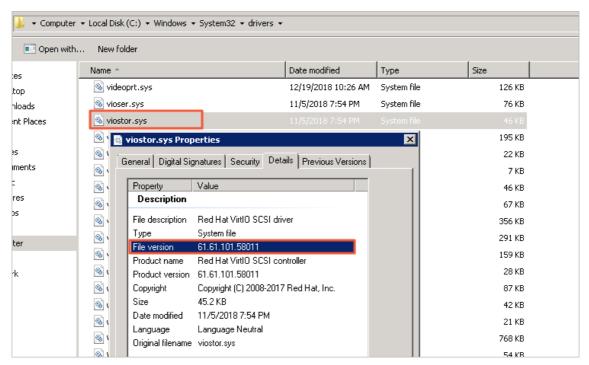
You can use one of the following methods to check the version of the driver:

- Method 1: Use the PowerShell script to check the version of the driver
 - i. Connect to the Windows instance. For more information, see Connect to a Windows instance by using a username and password.
 - ii. Open Command Prompt.
 - iii. Enter powershell to access the PowerShell interactive interface.
 - iv. Enter and run the following command to check the version and determine whether the ECS instance supports resizing disks online based on the command output:

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

- Method 2: Manually check the version of the driver
 - i. Connect to the Windows instance.

- ii. Go to the *C:\Windows\System32\drivers* system directory.
- iii. Right-click the *viostor.sys* file, select **Properties**, and then view the file version on the **Details** tab.



Determine whether the requirements of the file version are met based on the query results, and perform corresponding operations.

Scenario	Version of the Red Hat virtio-scsi driver	Operation		
Resize a disk	58011 or later	Resize the disk online. For more information, see Resize disks online for Windows instances.		
online	Earlier than 58011	Proceed to the next step.		
Query the serial number	58015 or later	Query the serial number of the disk. For more information, see Query the serial number of a disk.		
of a disk	Earlier than 58015	Proceed to the next step.		

Step 2: Update the driver

If your Windows instance can access the Internet, we recommend that you perform the operations described in this section to update the virtio driver of your instance in a quick manner. If the instance cannot access the Internet or if you want to batch update the virtio drivers, perform the operations described in Manually update virtio drivers of Windows instances.

- Connect to your Windows instance.
 For more information, see Connection methods.
- 2. In the Windows instance, download the script used to update the virtio driver. Click InstallVirtio.ps1 to download the InstallVirtio.ps1 script.

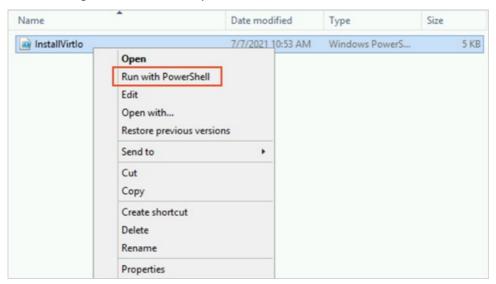
3. Run the *InstallVirtIo.ps1* script to update the virtio driver.

In this example, *InstallVirtIo.ps1* is downloaded to the *C:\test* directory.

i. Access the *C:\test* directory.

Access the directory to which *InstallVirtIo.ps1* is downloaded.

ii. Select and right-click InstallVirtIo.ps1 and then select Run with PowerShell.



Alternatively, you can press the *Shift* key and right-click a blank area in the folder window at the same time. Then, select **Open PowerShell window here**. Run the *InstallVirtIo.ps1* script in Windows PowerShell.

? Note

- If you are logged on to the Windows instance as a regular user, you must run this script as an administrator. If you are logged on as an administrator, you can directly run this script.
- If you are prompted that you are not allowed to run the script when you manually run the script, you must first modify the script execution policy of PowerShell. For more information, see FAQ.
- 4. After the script is run, restart the instance.

For more information, see Restart an instance. The updated virtio driver takes effect after your instance is restarted.

What's next

- For information about how to resize a disk online, see Resize a disk online for Linux instances.
- For information about how to query the serial number of a disk, see Query the serial number of a disk.
- For more information about how to manually update the Red Hat virtio drivers on Windows instances, see Manually update virtio drivers of Windows instances.

5. Support for the NVMe protocol and multi-attach feature

5.1. Overview of disks that support NVMe

Non-Volatile Memory Express (NVMe) is a host controller interface protocol used to accelerate the transfer of data from non-volatile memory. Alibaba Cloud enhanced SSDs (ESSDs) support NVMe. Each ESSD can be simultaneously attached to multiple Elastic Compute Service (ECS) instances that support NVMe for data sharing. This topic describes ESSDs that support NVMe and the limits for and operations of attaching the ESSDs.

ESSDs and NVMe

Note ESSDs that support NVMe and features related to these ESSDs are in public review in some zones including Hangzhou Zone I, Shanghai Zone B, Beijing Zone K, and Shenzhen Zone F.

ESSDs that support NVMe are called shared NVMe disks. Shared NVMe disks support concurrent read and write access from multiple ECS instances and provide high reliability, concurrency, and performance. Shared NVMe disks provide the multi-attach and I/O blocking features.

- A shared NVMe disk can be attached to a maximum of 16 ECS instances at the same time.
- You can run NVMe commands to control the permissions of ECS instances on shared NVMe disks. For more information about NVMe commands, see NVM Express Base Specification.

The preceding features improve service availability without compromising data reliability. Meanwhile, when a single point of failure (SPOF) occurs, you can use shared NVMe disks to quickly schedule and restore data. Data sharing among multiple ECS instances greatly reduces storage costs and improves service flexibility. Shared NVMe disks are suitable for high-availability databases and distributed database clusters that each consist of one write node and multiple read-only nodes.

Shared NVMe disks can be attached to ECS instances that support NVMe. For example, after shared NVMe disks are attached to Linux instances based on NVMe, you can run the lsblk command to check the device names and partition names of the shared NVMe disks, as shown in the following figure.

```
[root@ecs ~]# lsblk

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT

nvmeOn1 259:0 0 40G 0 disk

_nvmeOn1p1 259:1 0 40G 0 part /

nvmeIn1 259:2 0 100G 0 disk

[root@ecs ~]#
```

Description of the command output:

- The device names of the shared NVMe disks are displayed in the /dev/nvmeXn1 format. Examples: / dev/nvme0n1, /dev/nvme1n1, and /dev/nvme2n1.
- The partition names of the shared NVMe disks are displayed in the <Device name of the ESSD>p<Partition number> format. Examples: /dev/nvme0n1p1, /dev/nvme1n1p1, and /dev/nvme1n1p2.

Shared NVMe disks support the multi-attach feature and help enterprises to migrate their high-availability services to the cloud. For more information, see Enable multi-attach.

Limits

Before you can attach shared NVMe disks to an ECS instance based on NVMe, the resources of the instance must meet the conditions described in the following table.

Resource	Description				
	The instance family must support NVMe by default. Currently, the g7se, c7se, and r7se instance families support NVMe.				
Instance family	Note You can call the DescribeInstanceTypes operation to query instance families and check the NvmeSupport response parameter to determine whether the instance families support NVMe.				
	The image must contain the NVMe driver. The following public images are installed with the NVMe driver.				
	Note Only some public Linux and Windows images support the NVMe driver.				
lmage	 Alibaba Cloud Linux 2.1903 and later Ubuntu 16.04 and later Debian GNU/Linux 10 and later CentOS: CentOS 7: CentOS 7.6 and later CentOS 8: CentOS 8.0 and later openSUSE 15.2 and later Fedora 33 and later Fedora CoreOS Windows Server: Windows Server 2019 Windows Server 2016 				
Disk	 Disk category: ESSD. Creation method: Create ESSDs when you create instances that support NVMe. When you create disks, select ESSD as the disk category and enable the multi-attach feature. 				

Billing

NVMe is free of charge. You are charged for resources that support NVMe by using their individual billing methods. For more information about the billing of ECS resources, see Overview.

Operations

This section describes the operations that you can perform on shared NVMe disks.

- When you create an ECS instance that supports NVMe, you can create ESSDs. These ESSDs are shared NVMe disks. For more information, see Create an instance by using the wizard.
- You can enable the multi-attach feature when you separately create ESSDs. For more information, see Enable multi-attach.
- Format ESSDs and create file systems. The operations vary based on the partition formats of ESSDs.

Notice The following operations are applicable only to ESSDs that are created along with instances but not to ESSDs for which the multi-attach feature is enabled.

- Partition ESSDs that are larger than 2 TiB in size to GUID Partition Table (GPT). For more information, see Partition and format a data disk larger than 2 TiB in size.
- o Partition ESSDs that are smaller than or equal to 2 TiB in size to Master Boot Record (MBR). For more information, see Partition and format a data disk on a Linux instance or Partition and format a data disk on a Windows instance.

5.2. Enable multi-attach

When you create an enhanced SSD (ESSD), you can enable the multi-attach feature for the ESSD. After multi-attach is enabled for an ESSD, the ESSD can be attached to up to 16 Elastic Compute Service (ECS) instances that support the Non-Volatile Memory Express (NVMe) protocol within the same zone to allow concurrent read and write access from the instances.

Benefits

Note The multi-attach feature is in public review in the Hangzhou Zone I, Shanghai Zone B, Beijing Zone K, and Shenzhen Zone F.

This feature is suitable for high-availability databases and distributed database clusters that each consist of one write node and multiple read-only nodes and provides the following benefits:

- Use of NVMe commands: NVMe commands can be used to control the permissions of ECS instances on ESSDs. This helps improve service availability without compromising data durability. For more information about NVMe commands, see NVM Express Base Specification.
- Cross-instance data sharing: This feature enables data sharing by multiple ECS instances to reduce storage costs and improve service flexibility.
- Disaster recovery: This feature allows fast scheduling of services to normal ECS instances to ensure service continuity in single-point-of-failure (SPOF) scenarios.

Limits

The following limits apply to the multi-attach feature:

- Instance families that support the multi-attach feature include the ecs.g7se, ecs.c7se, and ecs.r7se.
- Only pay-as-you-go ESSDs are supported.
- Data disks are supported but system disks are not.
- The multi-attach feature can be enabled only when ESSDs are created. This feature cannot be enabled or disabled for created ESSDs.

After the multi-attach feature is enabled for an ESSD, we recommend that you use cluster file
systems such as Oracle Cluster File System version 2 (OCFS2), Global File System 2 (GFS2), Veritas
Cluster File System (Veritas CFS), Oracle Automatic Storage Management Cluster File System (Oracle
ACFS), and Databricks File System (DBFS) on the ESSD.

Warning When an ESSD for which the multi-attach feature is enabled is attached to multiple ECS instances, data cannot be synchronized among the instances and data inconsistency may occur if file systems such as EXT2, EXT3, EXT4, XFS, and New Technology File System (NTFS) are used.

• The performance of ESSDs is capped. When an ESSD is attached to multiple ECS instances, the total performance of this ESSD on all these instances cannot exceed the maximum performance that the ESSD can deliver.

ESSDs for which multi-attach is enabled are subject to the functionality limits described in the following table

Feature	Limit
Disk attaching	A single ESSD can be attached to up to 16 instances that support the NVMe protocol.
Disk billing method change	Not supported.
Release of disks with instances	Not supported.
Disk re-initialization	Not supported.
Disk category change	Not supported.
ESSD performance level change	Not supported.
Disk resizing	Only offline resizing of disks is supported. For more information, see Resize disks offline for Linux instances.
Snapshot-consistent group	Not supported.
Application-consistent snapshot	Not supported.

Use multi-attach

Perform the following steps to use the multi-attach feature:

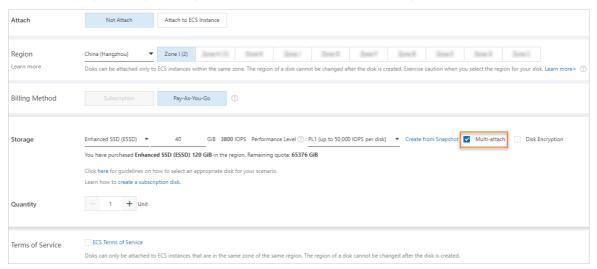
- 1. Step 1: Create an ESSD for which the multi-attach feature is enabled
- 2. Step 2: Attach the ESSD to multiple instances that support the NVMe protocol

Step 1: Create an ESSD for which the multi-attach feature is enabled

To use the multi-attach feature, you must enable this feature when you create ESSDs.

- 1.
- 2.
- 3. In the upper-right corner of the Disks page, click Create Disk.

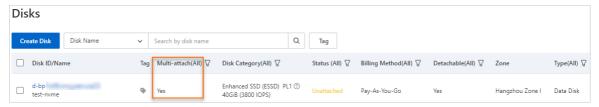
4. On the Disk page, configure the parameters described in the following table.



Parameter	Description					
Attach	Select Not Attach to specify not to attach the ESSD to be created to ECS instances.					
Billing Method	Select Pay-As-You-Go.					
Storage	 Configure storage settings and enable the multi-attach feature. Select Enhanced SSD (ESSD) from the drop-down list and set the Capacity and Performance Level parameters. If the number of instances within a same zone is two or more, select Multi-attach to attach a single ESSD to multiple instances in the zone. To encrypt an ESSD, select Disk Encrypt and then select an encryption key. 					
Quantity	Specify the number of disks that you want to create. Note Pay-as-you-go disks have capacity quotas. The total capacity that you purchase cannot exceed your capacity quota balance. You can use the following formula to calculate the total disk capacity: Total capacity = Capacity of a single disk × Number of disks. After you select a disk category, the purchased capacity and the capacity quota balance are displayed on the disk buy page.					
Terms of Service	Read and select <i>ECS Terms of Service</i> .					

Note For more information about how to create disks, see Create a disk.

- 5. Confirm the configurations and fees and click **Preview**.
- 6. In the Preview message, confirm the purchase information and click **Create**. After the ESSD is created, you can view it on the **Disks** page.

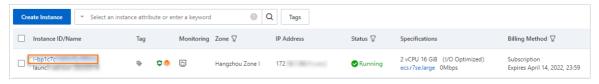


Step 2: Attach the ESSD to multiple instances that support the NVMe protocol

Before you attach the ESSD to instances, make sure that the following requirements are met:

- The ESSD and the instances reside in the same zone.
- The instance families and images of the instances comply with the NVMe protocol. For more information, see Limits.
 - 1. Obtain the ID of the instance to which you want to attach the ESSD.

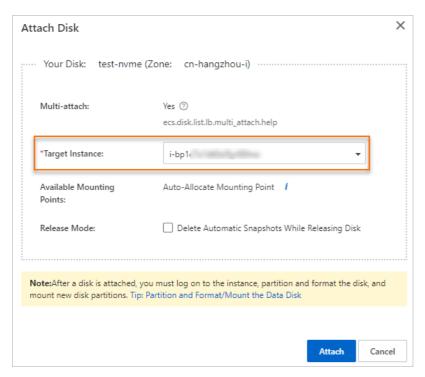
In the left-side navigation pane, choose **Instances & Images > Instances**. On the Instances page, view and copy the instance ID in the instance list.



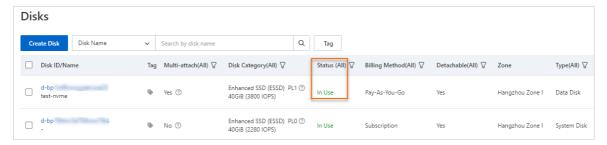
- 2.
- 3. Find the ESSD that you want to attach and choose More > Attach in the Actions column.
- 4. In the Attach Disk dialog box, select an instance ID from the Target Instance drop-down list and click Attach.

You must select the instance ID that you obtained in Step 1 from the **Target Instance** drop-down list.

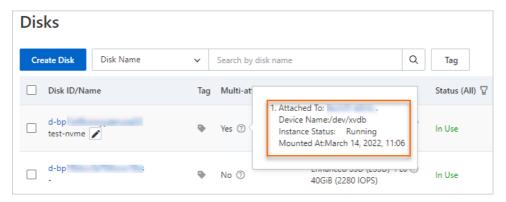
? Note You can attach the ESSD to only a single instance each time you perform the preceding steps. To attach the ESSD to multiple instances, repeat the preceding steps.



After the ESSD is attached, you can find that the state of the ESSD on the **Disks** page changes to **In Use**.



You can click the o icon in the Multi-attach (All) column to view the attach information of the ESSD on one or more instances.



What to do next

After the ESSD is attached to multiple instances, we recommend that you create cluster file systems based on your business requirements. Common cluster file systems include OCFS2, GFS2, Veritas CFS, Oracle ACFS, and DBFS.

☐ Warning When an ESSD for which the multi-attach feature is enabled is attached to multiple ECS instances, data cannot be synchronized among the instances and data inconsistency may occur if file systems such as EXT2, EXT3, EXT4, XFS, and New Technology File System (NTFS) are used.

6.Performance6.1. EBS performance

This topic describes the performance metrics and specifications of various Elastic Block Storage (EBS) devices, including cloud disks and local disks.

Performance metrics

The key metrics used to measure the performance of EBS devices include IOPS, throughput, and latency. Some EBS devices also have capacity requirements. For example, enhanced SSDs (ESSDs) at different performance levels (PLs) have different capacity ranges.

IOPS

IOPS measures the number of read/write operations that can be performed per second. High IOPS is critical for transaction-intensive applications such as databases. A standard SSD can deliver the committed IOPS only when it is attached to an I/O optimized instance. For more information about I/O optimized instances, see Instance family. The following table describes the common IOPS metrics.

Metric	Description	Data access method		
Total IOPS	The total number of I/O operations per second.	Access locations on storage devices in a continuous or non-continuous manner.		
Random read IOPS	The average number of random read I/O operations per second.	Access locations on storage devices in a		
Random write IOPS	The average number of random write I/O operations per second.	non-continuous manner.		
Sequential read IOPS	The average number of sequential read I/O operations per second.	Access locations on storage devices in a		
Sequential write IOPS	The average number of sequential write I/O operations per second.	continuous manner.		

Throughput

Throughput measures the amount of data transferred per second. Unit: MB/s. High throughput is critical for applications such as Hadoop offline computing applications that require a large number of sequential read/write operations.

Latency

132

Latency measures the amount of time required for an EBS device to process an I/O request. Unit: seconds, milliseconds, or microseconds. High latency may cause performance to degrade or lead to errors in applications that require low latency.

- For latency-sensitive applications such as databases, we recommend that you use ESSDs, standard SSDs, or local SSDs.
- For applications such as Hadoop offline computing applications that require high throughput but not low latency, we recommend that you use Elastic Compute Service (ECS) instances of the d1 or d1ne instance family that have local SATA HDDs attached.

Capacity

Capacity is the amount of storage space. Unit: TiB, GiB, MiB, or KiB. EBS capacity is measured in binary units. For example, 1 GiB equals 1,024 MiB.

You cannot use capacity as a metric to measure the performance of EBS devices. However, the performance of EBS devices varies based on their capacity. The larger the capacity of an EBS device is, the stronger its processing capabilities are. EBS devices of the same category have the same I/O performance per unit capacity. However, the performance of a cloud disk linearly increases with its capacity up to the single-disk maximum performance of the disk category. ESSDs of different capacity ranges have different PLs.

For more information about how to test the performance of different categories of EBS devices, see Test the performance of EBS devices and Test the IOPS performance of an ESSD.

Performance of cloud disks

The following table describes the performance and typical use scenarios of different categories of cloud disks.

ltem	ESSD				Standard SSD	Ultra disk	Basic disk ③
PL	PL3	PL2	PL1	PL0	N/A	N/A	N/A
Single- disk capacity range (GiB)	1,261 to 32,768	461 to 32,768	20 to 32,768	40 to 32,768	20 to 32,768	20 to 32,768	5 to 2,000
Maximum IOPS	1,000,000	100,000	50,000	10,000	25,000 ^①	5,000	Several hundreds
Maximum throughp ut (MB/s)	4,000	750	350	180	300 ^①	140	30 to 40
Formula used to calculate single- disk IOPS ②	min{1,800 + 50 × Capacity, 1,000,000}	min{1,800 + 50 × Capacity, 100,000}	min{1800 + 50 × Capacity, 50,000}	min{1,800 +12 × Capacity, 10,000}	min{1,800 + 30 × Capacity, 25,000}	min{1,800 + 8 × Capacity, 5,000}	None
Formula used to calculate singledisk throughp ut (MB/s)	min{120 + 0.5 × Capacity, 4,000}	min{120 + 0.5 × Capacity, 750}	min{120 + 0.5 × Capacity, 350}	min{100 + 0.25 × Capacity, 180}	min{120 + 0.5 × Capacity, 300}	min{100 + 0.15 × Capacity, 140}	None
Data durability	99.999999 9%	99.999999 9%	99.999999 9%	99.99999 99%	99.99999 99%	99.99999 99%	99.99999 99%

ltem	ESSD			Standard SSD	Ultra disk	Basic disk ③	
Average single- channel random write latency in millisecon ds (block size = 4 K)	0.2	0.2	0.2	0.3 to 0.5	0.5 to 2	1 to 3	5 to 10

• ^① The performance of standard SSDs varies based on the sizes of data blocks. Smaller data blocks result in lower throughput and higher IOPS, as described in the following table.

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

- ^② In the following examples, a standard SSD is used to describe how to calculate single-disk performance:
 - Single-disk maximum IOPS: The baseline IOPS is 1,800. The IOPS increases by 30 per additional GiB until it reaches the maximum IOPS 25,000.
 - Single-disk maximum throughput: The baseline throughput is 120 MB/s. The throughput increases by 0.5 MB/s per additional GiB until it reaches the maximum throughput 300 MB/s.
- ullet Basic disks are cloud disks of a previous generation and are no longer available for purchase.

Performance of local disks

For information about the performance of local Non-Volatile Memory Express (NVMe) SSDs and SATA HDDs, see Local disks.

6.2. Storage I/O performance

Storage I/O performance, also known as storage read/write performance, is the performance that can be delivered when disks are attached to instances of different ECS instance types. Metrics of storage I/O performance include IOPS and throughput.

I/O size

I/O (Input/Output or Read/Write) is random or sequential data requests initiated by an application. The volume of I/O requests is also known as the I/O size in KiB, such as 4 KiB, 256 KiB, and 1,024 KiB.

When you design the underlying storage architecture or select an instance type, you must consider metrics such as IOPS, I/O size, and throughput. You can use the following formula to calculate throughput based on the I/O size and IOPS: IOPS × I/O size = Throughput. You can select different Elastic Block Storage devices and instance types based on the I/O request features of applications to achieve the best results.

- For the applications such as offline analysis and data warehousing applications that require a large I/O size, we recommend that you select the big data instance families that deliver high throughput.
- For the applications such as OLTP databases and enterprise-level applications like SAP that require low-latency, random, and small-sized I/O, we recommend that you select enhanced SSDs (ESSDs) and standard SSDs that deliver high IOPS.

For more information about IOPS and throughput, see EBS performance.

Storage I/O performance of instances

? Note This section and subsequent sections apply only to the new generation of enterprise-level instance families, including hfg7, hfc7, hfr7, g6e, c6e, and r6e. For more information, see Storage I/O performance of the new generation of enterprise-level instance families. These sections are not applicable to local disks.

The new generation of enterprise-level instance families of Alibaba Cloud ECS provides isolation of storage I/O performance. Dedicated storage bandwidths are assigned to ECS instances and disks to avoid storage I/O preemption among ECS instances. The new generation of enterprise-level instance families ensures consistent storage I/O performance of applications even during peak hours.

If your business application is one of the following I/O-sensitive applications that require consistent storage I/O performance, we recommend that you select a new-generation instance family that provides isolation of storage I/O performance:

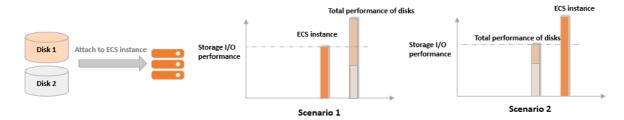
- Large and medium-sized databases, such as Oracle, MySQL, SQL Server, PostgreSQL, Cassandra, and MongoDB databases
- Enterprise-level applications such as ERP and CRM

Relationship between instance types and storage I/O performance

The storage I/O performance of ECS instances varies based on their instance families, instance types, and attached disks. The storage I/O performance of an instance is subject to the instance type. The higher specifications an instance type has, the higher storage I/O performance (IOPS and throughput) it provides.

After you create an ECS instance and attach disks to it, the final storage I/O performance of the ECS instance is determined as described in the following section:

- Scenario 1: If the total storage performance of the attached disks exceeds the maximum storage I/O performance that the instance type can deliver, the final storage I/O performance of the instance is limited to the maximum storage I/O performance of that instance type.
- Scenario 2: If the total storage performance of the attached disks does not exceed the maximum storage I/O performance that the instance type can deliver, the final storage I/O performance of the instance is limited to the total storage I/O performance of the disks.



For example, the ecs.g6.8xlarge instance type can deliver up to 60,000 IOPS. If a 1,600 GiB PL2 ESSD is attached to an instance of this instance type and the IOPS of the PL2 ESSD is 81,800, the maximum storage IOPS of the instance is still 60,000. For information about the performance levels of ESSDs, see ESSDs.

When you understand the relationship between instance storage performance and disk storage performance, you can choose instance types and Elastic Block Storage devices based on their performance specifications. This prevents performance bottlenecks caused by improper configurations.

Storage I/O performance of the new generation of enterprise-level instance families

The following tables list the storage I/O performance of the new generation of enterprise-level instance families. For information about other specifications of the instance families, see Instance family.

Note In these tables, the maximum IOPS is measured at an I/O size of 4 KiB and the maximum throughput is measured at an I/O size of 1,024 KiB. For information about test methods, see Test the performance of EBS devices.

Storage I/O performance of g7se

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.g7se.large	3	3	375
ecs.g7se.xlarge	6	4	500
ecs.g7se.2xlarge	10	6	750
ecs.g7se.3xlarge	12	8	1,000
ecs.g7se.4xlarge	15	10	1,250
ecs.g7se.6xlarge	20	12	1,500
ecs.g7se.8xlarge	30	16	2,000
ecs.g7se.16xlarge	50	32	4,000
ecs.g7se.32xlarge	100	64	8,000

Storage I/O performance of hfg7

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfg7.large	2.0	1.0	125
ecs.hfg7.xlarge	3.0	1.5	187.5
ecs.hfg7.2xlarge	4.5	2.0	250
ecs.hfg7.3xlarge	6.0	2.5	312.5
ecs.hfg7.4xlarge	7.5	3.0	375
ecs.hfg7.6xlarge	9.0	4.0	500
ecs.hfg7.8xlarge	10.5	5.0	625
ecs.hfg7.12xlarge	15.0	8.0	1,000
ecs.hfg7.24xlarge	30.0	16.0	2,000

Storage I/O performance of hfc7

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfc7.large	2.0	1.0	125
ecs.hfc7.xlarge	3.0	1.5	187.5
ecs.hfc7.2xlarge	4.5	2.0	250
ecs.hfc7.3xlarge	6.0	2.5	312.5
ecs.hfc7.4xlarge	7.5	3.0	375
ecs.hfc7.6xlarge	9.0	4.0	500
ecs.hfc7.8xlarge	10.5	5.0	625
ecs.hfc7.12xlarge	15.0	8.0	1,000
ecs.hfc7.24xlarge	30.0	16.0	2,000

Storage I/O performance of hfr7

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfr7.large	2.0	1.0	125
ecs.hfr7.xlarge	3.0	1.5	187.5

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfr7.2xlarge	4.5	2.0	250
ecs.hfr7.3xlarge	6.0	2.5	312.5
ecs.hfr7.4xlarge	7.5	3.0	375
ecs.hfr7.6xlarge	9.0	4.0	500
ecs.hfr7.8xlarge	10.5	5.0	625
ecs.hfr7.12xlarge	15.0	8.0	1,000
ecs.hfr7.24xlarge	30.0	16.0	2,000

Storage I/O performance of g6t

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.g6t.large	2.0	1.0	125
ecs.g6t.xlarge	4.0	1.5	187.5
ecs.g6t.2xlarge	5.0	2.0	250
ecs.g6t.4xlarge	8.0	3.0	375
ecs.g6t.8xlarge	15.0	5.0	625
ecs.g6t.13xlarge	24.0	8.0	1,000
ecs.g6t.26xlarge	48.0	16.0	2,000

Storage I/O performance of c6t

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.c6t.large	2.0	1.0	125
ecs.c6t.xlarge	4.0	1.5	187.5
ecs.c6t.2xlarge	5.0	2.0	250
ecs.c6t.4xlarge	8.0	3.0	375
ecs.c6t.8xlarge	15.0	5.0	625
ecs.c6t.13xlarge	24.0	8.0	1,000

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.c6t.26xlarge	48.0	16.0	2,000

Storage I/O performance of g6e

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.g6e.large	2.0	1.0	125
ecs.g6e.xlarge	4.0	1.5	187.5
ecs.g6e.2xlarge	5.0	2.0	250
ecs.g6e.4xlarge	8.0	3.0	375
ecs.g6e.8xlarge	15.0	5.0	625
ecs.g6e.13xlarge	24.0	8.0	1,000
ecs.g6e.26xlarge	48.0	16.0	2,000

Storage I/O performance of c6e

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.c6e.large	2.0	1.0	125
ecs.c6e.xlarge	4.0	1.5	187.5
ecs.c6e.2xlarge	5.0	2.0	250
ecs.c6e.4xlarge	8.0	3.0	375
ecs.c6e.8xlarge	15.0	5.0	625
ecs.c6e.13xlarge	24.0	8.0	1,000
ecs.c6e.26xlarge	48.0	16.0	2,000

Storage I/O performance of r6e

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.r6e.large	2.0	1.0	125
ecs.r6e.xlarge	4.0	1.5	187.5
ecs.r6e.2xlarge	5.0	2.0	250

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.r6e.4xlarge	8.0	3.0	375
ecs.r6e.8xlarge	15.0	5.0	625
ecs.r6e.13xlarge	24.0	8.0	1,000
ecs.r6e.26xlarge	48.0	16.0	2,000

Storage I/O performance of g6

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.g6.large	1.0	1.0	125
ecs.g6.xlarge	2.0	1.5	187.5
ecs.g6.2xlarge	2.5	2.0	250
ecs.g6.3xlarge	3.0	2.5	312.5
ecs.g6.4xlarge	4.0	3.0	375
ecs.g6.6xlarge	5.0	4.0	500
ecs.g6.8xlarge	6.0	5.0	625
ecs.g6.13xlarge	10.0	8.0	1,000
ecs.g6.26xlarge	20.0	16.0	2,000

Storage I/O performance of c6

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.c6.large	1.0	1.0	125
ecs.c6.xlarge	2.0	1.5	187.5
ecs.c6.2xlarge	2.5	2.0	250
ecs.c6.3xlarge	3.0	2.5	312.5
ecs.c6.4xlarge	4.0	3.0	375
ecs.c6.6xlarge	5.0	4.0	500
ecs.c6.8xlarge	6.0	5.0	625

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.c6.13xlarge	10.0	8.0	1,000
ecs.c6.26xlarge	20.0	16.0	2,000

Storage I/O performance of r6

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.r6.large	1.0	1.0	125
ecs.r6.xlarge	2.0	1.5	187.5
ecs.r6.2xlarge	2.5	2.0	250
ecs.r6.3xlarge	3.0	2.5	312.5
ecs.r6.4xlarge	4.0	3.0	375
ecs.r6.6xlarge	5.0	4.0	500
ecs.r6.8xlarge	6.0	5.0	625
ecs.r6.13xlarge	10.0	8.0	1,000
ecs.r6.26xlarge	20.0	16.0	2,000

Storage I/O performance of hfg6

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfg6.large	1.0	1.0	125
ecs.hfg6.xlarge	2.0	1.5	187.5
ecs.hfg6.2xlarge	2.5	2.0	250
ecs.hfg6.3xlarge	3.0	2.5	312.5
ecs.hfg6.4xlarge	4.0	3.0	375
ecs.hfg6.6xlarge	5.0	4.0	500
ecs.hfg6.8xlarge	6.0	5.0	625
ecs.hfg6.10xlarge	10.0	8.0	1,000
ecs.hfg6.16xlarge	12.0	10.0	1,250

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfg6.20xlarge	20.0	16.0	2,000

Storage I/O performance of hfc6

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfc6.large	1.0	1.0	125
ecs.hfc6.xlarge	2.0	1.5	187.5
ecs.hfc6.2xlarge	2.5	2.0	250
ecs.hfc6.3xlarge	3.0	2.5	312.5
ecs.hfc6.4xlarge	4.0	3.0	375
ecs.hfc6.6xlarge	5.0	4.0	500
ecs.hfc6.8xlarge	6.0	5.0	625
ecs.hfc6.10xlarge	10.0	8.0	1,000
ecs.hfc6.16xlarge	12.0	10.0	1,250
ecs.hfc6.20xlarge	20.0	16.0	2,000

Storage I/O performance of hfr6

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfr6.large	1.0	1.0	125
ecs.hfr6.xlarge	2.0	1.5	187.5
ecs.hfr6.2xlarge	2.5	2.0	250
ecs.hfr6.3xlarge	3.0	2.5	312.5
ecs.hfr6.4xlarge	4.0	3.0	375
ecs.hfr6.6xlarge	5.0	4.0	500
ecs.hfr6.8xlarge	6.0	5.0	625
ecs.hfr6.10xlarge	10.0	8.0	1,000
ecs.hfr6.16xlarge	12.0	10.0	1,250

Instance type	Maximum IOPS (K, 4 KiB I/O)	Maximum storage bandwidth (Gbit/s)	Maximum throughput (MB/s, 1,024 KiB I/O)
ecs.hfr6.20xlarge	20.0	16.0	2,000

6.3. Test the performance of EBS devices

This topic describes how to use the fio tool on a Linux Elastic Compute Service (ECS) instance to test the performance of Elastic Block Storage (EBS) devices. EBS devices include cloud disks and local disks, and the performance metrics of these disks include IOPS, throughput, and latency.

Prerequisites

An EBS device is created and attached to a Linux ECS instance.

Note If you want to test only the performance of a specific category of EBS devices, we recommend that you use a new pay-as-you-go data disk. You can release the disk after the test is complete.

Context

You can use other tools to test the performance of EBS devices, but you may obtain different baseline performance. For example, tools such as **dd**, **sysbench**, and **iometer** may be affected by test parameters and file systems and return inaccurate results. The performance results in this topic are obtained from a **fio** test on a Linux instance. These results are used as performance references for Alibaba Cloud EBS devices. We recommend that you use the **fio** tool to test the performance of EBS devices for both Linux and Windows instances.

Warning

- You can obtain accurate test results by testing raw disk partitions. However, you may destroy the file system structure in a raw disk partition if you directly test the partition. Before you test a raw disk, we recommend that you back up your data by taking a snapshot of the disk. For more information, see Create a snapshot of a disk.
- We recommend that you do not test a disk where the operating system is located or a disk
 that stores important data. To prevent data loss, we recommend that you use a new Elastic
 Compute Service (ECS) instance that contains no data for the test.

Procedure

1.

2. Before you test an EBS device, make sure that it is 4 KiB aligned.

```
sudo fdisk -lu
```

If the value of Start in the command output is divisible by 8, the device is 4 KiB aligned. Otherwise, perform 4 KiB alignment before you proceed with the test.

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

3. Run the following commands in sequence to install libaio and FIO:

```
sudo yum install libaio -y
sudo yum install libaio-devel -y
sudo yum install fio -y
```

4. Run the following command to switch the path:

```
cd /tmp
```

- 5. Run the test commands. For more information about the commands, see the following sections:
 - For information about commands used to test the performance of cloud disks, see Commands used to test the performance of cloud disks.
 - For information about commands used to test the performance of local disks, see Commands used to test the performance of local disks.

Commands used to test the performance of cloud disks

For information about how to test the IOPS of an enhanced SSD (ESSD), see Test the IOPS performance of an ESSD.

Note In this example, the device name of the cloud disk is /dev/your_device. Replace it with your actual device name. For example, if the device name of the cloud disk that you want to test is /dev/vdb, replace /dev/your device with /dev/vdb in the following commands.

• Command used to test the random write IOPS of a cloud disk:

```
fio -direct=1 -iodepth=128 -rw=randwrite -ioengine=libaio -bs=4k -size=1G -numjobs=1 -run time=1000 -group_reporting -filename=/dev/your_device -name=Rand_Write_Testing
```

• Command used to test the random read IOPS of a cloud disk:

```
\label{limit}  \mbox{fio -direct=1 -iodepth=128 -rw=randread -ioengine=libaio -bs=4k -size=1G -numjobs=1 -runtime=1000 -group\_reporting -filename=/dev/your\_device -name=Rand\_Read\_Testing}
```

Command used to test the sequential write throughput of a cloud disk:

```
fio -direct=1 -iodepth=64 -rw=write -ioengine=libaio -bs=1024k -size=1G -numjobs=1 -runti me=1000 -group_reporting -filename=/dev/your_device -name=Write_PPS_Testing
```

• Command used to test the sequential read throughput of a cloud disk:

• Command used to test the random write latency of a cloud disk:

```
fio -direct=1 -iodepth=1 -rw=randwrite -ioengine=libaio -bs=4k -size=1G -numjobs=1 -group _reporting -filename=/dev/your_device -name=Rand_Write_Latency_Testing
```

• Command used to test the random read latency of a cloud disk:

fio -direct=1 -iodepth=1 -rw=randread -ioengine=libaio -bs=4k -size=1G -numjobs=1 -group_reporting -filename=/dev/your_device -name=Rand_Read_Latency_Testing

Commands used to test the performance of local disks

The following test commands are applicable to local Non-Volatile Memory Express (NVMe) SSDs and local Serial Advanced Technology Attachment (SATA) HDDs.

Note In this example, the device name of the local disk is /dev/your_device. Replace it with your actual device name. For example, if the device name of the local disk that you want to test is /dev/vdb, replace /dev/your device with /dev/vdb in the following commands.

• Command used to test the random write IOPS of a local disk:

```
fio -direct=1 -iodepth=32 -rw=randwrite -ioengine=libaio -bs=4k -numjobs=4 -time_based=1 -runtime=1000 -group_reporting -filename=/dev/your_device -name=test
```

Command used to test the random read IOPS of a local disk:

```
fio -direct=1 -iodepth=32 -rw=randread -ioengine=libaio -bs=4k -numjobs=4 -time_based=1 -runtime=1000 -group reporting -filename=/dev/your device -name=test
```

• Command used to test the sequential write throughput of a local disk:

```
fio -direct=1 -iodepth=128 -rw=write -ioengine=libaio -bs=128k -numjobs=1 -time_based=1 -runtime=1000 -group_reporting -filename=/dev/your_device -name=test
```

Command used to test the sequential read throughput of a local disk:

```
fio -direct=1 -iodepth=128 -rw=read -ioengine=libaio -bs=128k -numjobs=1 -time_based=1 -r untime=1000 -group reporting -filename=/dev/your device -name=test
```

• Command used to test the random write latency of a local disk:

Command used to test the random read latency of a local disk:

```
fio -direct=1 -iodepth=1 -rw=randread -ioengine=libaio -bs=4k -numjobs=1 -time_based=1 -r untime=1000 -group reporting -filename=/dev/your device -name=test
```

• Command used to test the sequential write latency of a local disk:

```
\label{limited_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_final_f
```

• Command used to test the sequential read latency of a local disk:

```
fio -direct=1 -iodepth=1 -rw=read -ioengine=libaio -bs=4k -numjobs=1 -time_based=1 -runti me=1000 -group reporting -filename=/dev/your device -name=test
```

fio parameter settings

fio commands use specific parameters. The following table describes the parameter settings in the fio command used to test the random write IOPS (*randwrite*) of a cloud disk.

Parameter setting	Description	
-direct=1	Indicates that the I/O buffer is ignored during the test and data is directly written.	
-iodepth=128	Indicates that when asynchronous I/O is used, up to 128 I/O requests can be concurrently made.	
-rw=randwrite	Indicates that the read/write policy is random writes. Other valid values for rw: • randread: random reads • read: sequential reads • write: sequential writes • randrw: random reads and writes	
-ioengine=libaio	 Indicates that libaio (the Linux-native asynchronous I/O facility) is used for the test. An application can use I/O in the following ways: Synchronous In synchronous I/O mode, a thread sends a single I/O request at a time and waits for the request to complete. In this case, the iodepth value is always less than 1 for a single thread. You can increase the iodepth value by using multiple concurrent threads. The iodepth value reaches its upper limit when 16 to 32 threads are running concurrently. Asynchronous In asynchronous I/O mode, a thread uses libaio to send multiple I/O requests at a time and waits for all these requests to complete. Asynchronous I/O helps reduce the number of interactions and make interactions more efficient. 	
-bs=4k	 Indicates that the size of each block for one I/O is 4 KiB. The default value is also 4 KiB. When IOPS is tested, we recommend that you set bs to a small value such as 4k. When throughput is tested, we recommend that you set bs to a large value such as 1024k. 	
-size=1G	Indicates that the size of the test file is 1 GiB.	
-numjobs=1	Indicates that the number of test threads is 1.	
-runtime=1000	Indicates that the test duration is 1,000 seconds. If this parameter is not specified, the file of the size specified by -size is written in blocks of the size specified by -bs.	
-group_reporting	Indicates that in the test results, statistics are displayed for a group of threads, not for each individual thread.	

Parameter setting	Description
- filename=/dev/your_devic e	Indicate that the device name of the cloud disk is /dev/your_device. Replace the value with the actual name of your cloud disk.
- name=Rand_Write_Testing	Indicates that the name of the test task is <i>Rand_Write_Testing</i> . You can specify a name.

6.4. Test the IOPS performance of an ESSD

This topic describes how to test the IOPS performance of an enhanced SSD (ESSD). The specifications of the disk and the test conditions affect the test results. If you configure the test conditions as described in the following example to make full use of the performance of the multi-core and high-concurrency system, you can obtain an IOPS value of one million when you perform a stress test on the ESSD.

Test conditions

- Sample operation: random write (randwrite).
- Image: We recommend that you use a later version of a Linux public image provided by Alibaba Cloud, such as Cent OS 7.4 64-bit, Cent OS 7.3 64-bit, Cent OS 7.2 64-bit, or Alibaba Cloud Linux 2.1903 64-bit.
- Tool: We recommend that you use FIO.
- Instance type: We recommend that you use ecs.g7se.32xlarge. For more information, see g7se, storage-enhanced general-purpose instance family.
- ESSD: We recommend that you use an ESSD at the performance level 3 (PL3 ESSD). In this example, the device name of the ESSD is /dev/your_device. Replace it with your actual device name. For more information, see ESSDs.

Warning

- You can obtain accurate test results by testing raw disk partitions. However, you may
 destroy the file system structure in a raw disk partition if you directly test the partition.
 Before you test a raw disk, we recommend that you back up your data by taking a
 snapshot of the disk. For more information, see Create a snapshot of a disk.
- We recommend that you do not test a disk where the operating system is located or a
 disk that stores important data. To prevent data loss, we recommend that you use a new
 Elastic Compute Service (ECS) instance that contains no data for the test.

Procedure

1.

2. Run the following commands in sequence to install libaio and FIO:

```
sudo yum install libaio -y
sudo yum install libaio-devel -y
sudo yum install fio -y
```

3. Run the following command to switch the path:

```
cd /tmp
```

4. Run the following command to create the *test 100w.sh* script:

```
vim test100w.sh
```

5. Paste the following content to the test 100w.sh script.

For more information, see Details about the test 100w.sh script.

```
function RunFio
numjobs=$1 # The number of test threads. In this example, the value is 10.
iodepth=$2  # The maximum number of concurrent I/O requests. In this example, the val
             # The data block size per I/O. In this example, the value is 4k.
bs=$3
              # The read and write policy. In this example, the value is randwrite.
filename=$5  # The name of the test file. In this example, the value is /dev/your devi
 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++))
     list=`cat /sys/block/your device/mq/*/cpu list | awk '{if(i<=NF) print $i;}' i="$i</pre>
" | tr -d ',' | tr '\n' ','`
    if [ -z $list ]; then
        break
     cpulist=${cpulist}${list}
spincpu=`echo $cpulist | cut -d ',' -f 2-${nu}`
echo $spincpu
fio --ioengine=libaio --runtime=30s --numjobs=${numjobs} --iodepth=${iodepth} --bs=${b
s} --rw=${rw} --filename=${filename} --time based=1 --direct=1 --name=test --group repo
rting --cpus allowed=$spincpu --cpus_allowed_policy=split
echo 2 > /sys/block/your device/queue/rq affinity
RunFio 10 64 4k randwrite /dev/your device
```

- 6. Modify the parameter settings in the *test100w.sh* script based on the actual conditions.
 - Replace your device with the actual device name of the ESSD. Example: nvme1n1.
 - Modify 10, 64, 4k, randwrite, and /dev/your_device in RunFio 10 64 4k randwrite /dev/your_device based on the actual conditions.
 - To proceed with this operation, set filename to a device name. Example: [/dev/vdb]. If you do not want to risk data loss, set filename to a file path. Example: [/mnt/test.image].

Warning

- You can obtain accurate test results by testing raw disk partitions. However, you may
 destroy the file system structure in a raw disk partition if you directly test the partition.
 Before you test a raw disk, we recommend that you back up your data by taking a
 snapshot of the disk. For more information, see Create a snapshot of a disk.
- We recommend that you do not test a disk where the operating system is located or a disk that stores important data. To prevent data loss, we recommend that you use a new Elastic Compute Service (ECS) instance that contains no data for the test.
- 7. Run the following command to test the performance of the ESSD:

```
sh test100w.sh
```

If the IOPS=*** result is displayed in the command output, the performance stress test for the ESSD is complete.

Details about the test100w.sh script

• In the script, the following command sets the rg affinity parameter to 2:

```
echo 2 > /sys/block/your_device/queue/rq_affinity

Value of rq_affinity

Description
```

Value of rq_affinity	Description
1	Indicates that the block device delivers received I/O completion events to the group of the vCPUs that submit the corresponding I/O requests. In scenarios where multiple threads run concurrently, I/O completion events may be delivered only to one vCPU and cause a performance bottleneck.
2	Indicates that the block device delivers received I/O completion events to the vCPUs that submit the corresponding I/O requests. In scenarios where multiple threads run concurrently, each vCPU can deliver its maximum performance.

• The following command runs jobs to bind queues to different CPU cores:

Note In normal mode, a device has a single request queue. This request queue becomes a performance bottleneck when multiple threads concurrently process I/O requests. In multi-queue mode, a device can have multiple request queues to process I/O requests and deliver the maximum backend storage performance. For example, assume that you have four I/O threads. To make full use of multi-queue mode and improve storage performance, you must bind these threads to the CPU cores that correspond to different request queues.

Parameter	Description	Example
numjobs	The number of I/O threads.	10
/dev/your_device	The device name of the ESSD.	/dev/nvme1n1
cpus_allowed_policy	The parameter provided by FIO to bind vCPUs. FIO provides the cpus_allowed_policy and cpus_allowed parameters to bind vCPUs.	split

The preceding command runs jobs to bind queues to different CPU cores. These queues have different queue IDs. To view the ID of the CPU core to which a queue is bound, perform the following steps:

- o Run the ls /sys/block/your_device/mq/ command. In this command, replace your_device with your actual device name. Example: nvme1n1. This command returns the ID of the queue for an ESSD whose device name is in the /dev/vd* format.
- Run the __cat /sys/block/your_device/mq//cpu_list command. In this command, replace __your_device __with your actual device name. Example: nvme1n1. This command returns the ID of the CPU core to which the queue for an ESSD is bound. The device name of the ESSD is in the /dev/vd* format.

7.Dedicated block storage clusters

7.1. Overview

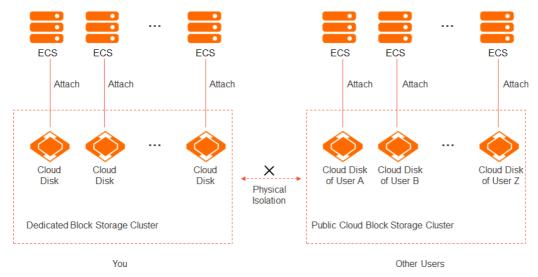
Dedicated Block Storage Cluster is a service that provides block storage resources in clusters. Each dedicated block storage cluster is physically isolated and allows its owner exclusive access to its resources. You can use dedicated block storage clusters to improve the security and O&M efficiency of your business data storage.

Features

Note Dedicated Block Storage Cluster is currently available in the China (Heyuan), Indonesia (Jakarta), and China South 1 Finance regions.

Dedicated block storage clusters are physically isolated from public block storage clusters. The owner of each dedicated block storage cluster has exclusive access to all resources in the cluster. Dedicated block storage clusters are physically isolated from public block storage clusters. The owner of each dedicated block storage cluster has exclusive access to all resources in the cluster. Dedicated Block Storage Cluster provides a monitoring feature that allows you to view cluster information such as cluster states and capacities in real time.

Unlike public block storage clusters, dedicated block storage clusters offer physical isolation and exclusive resources. After you purchase a dedicated block storage cluster, you can create dedicated disks in the cluster. Dedicated disks provide the same features as their public counterparts and grant you exclusive access to their physical resources. You can use the Elastic Compute Service (ECS) console or call API operations to manage disks in dedicated block storage clusters.



Disks in dedicated block storage clusters can be used in the following ways:

- The disks can be used as system disks on ECS instances. You must enable Dedicated Block Storage Cluster when you create ECS instances before you can attach disks in dedicated block storage clusters to the instances.
- The disks can be used as data disks on ECS instances. You can enable Dedicated Block Storage Cluster

and specify a dedicated block storage cluster when you add data disks during the creation of an ECS instance. The added data disks are all created in the specified dedicated block storage cluster. Alternatively, you can manually create disks in a dedicated block storage cluster. The created disks can be attached only as data disks to ECS instances within the same zone as the cluster, regardless of whether the instances reside on dedicated hosts.

Use scenarios

Resources in dedicated block storage clusters are physically isolated from those in public block storage clusters. Each dedicated block storage cluster is dedicated to its owner. Dedicated block storage clusters are suitable for industries or sectors that require high data security, such as finance, healthcare, and public service sector.

Billing

For information about the billing of dedicated block storage clusters, see Billing.

Limits

The following table describes the limits of Dedicated Block Storage Cluster.

ltem	Limits	
	 Minimum capacity: 61,440 GiB Maximum capacity: 2,359,296 GiB The capacity of each dedicated block storage cluster can be changed by multiples of 12,288 GiB. 	
Cluster capacity	Note If the capacity of a dedicated block storage cluster is less than 576 TiB, the maximum throughput supported per TiB does not exceed 52 MB/s. If the capacity of a dedicated block storage cluster is greater than 576 TiB, the maximum throughput supported per TiB does not exceed 26 MB/s.	
Disk category	Dedicated block storage clusters of different performance types support different disk categories. Basic dedicated block storage clusters support enhanced SSDs at performance level 0 (PL0 ESSDs). Performance dedicated block storage clusters support enhanced SSDs at performance level 1 (PL1 ESSDs).	
	For more information about ESSDs, see ESSDs. Note The configurations of ESSDs in dedicated block storage clusters cannot be changed.	

Lifecycle management

You

Create a dedicated block storage cluster

Performance cluster

Performance cluster

Create PL0
ESSDs

Monitor the cluster

Monitor the cluster

The following figure shows the lifecycle of a dedicated block storage cluster from creation to O&M.

You can perform the following operations:

1. Create a dedicated block storage cluster.

Dedicated block storage clusters are classified into the basic and performance types. Select a cluster type that is appropriate to your business requirements. For more information about how to create a dedicated block storage cluster, see Create a dedicated block storage cluster.

2. Monitor a dedicated block storage cluster.

After a dedicated block storage cluster is created, you can view its information such as state, capacity, and dedicated disks. For more information, see View the information of a dedicated block storage cluster.

3. Create disks in a dedicated block storage cluster.

Dedicated block storage clusters of different performance types support different disk categories. System disks and data disks are created in different ways in dedicated block storage clusters. For more information, see Create disks in a dedicated block storage cluster.

7.2. Billing

This topic describes the billing methods and billable items of the Dedicated Block Storage Cluster service.

? Note The Dedicated Block Storage Cluster service is now available in the China (Heyuan), Indonesia (Jakarta), and China South 1 Finance regions.

Billable items

Fee for a dedicated block storage cluster = Cluster capacity × Subscription duration × Unit price of the disk category. The following table describes the billable items.

Description

Billable item	Description
Cluster capacity	Cluster capacity is billed by TiB. If you create a dedicated block storage cluster in the Elastic Compute Service (ECS) console, the system converts GiB into TiB based on the following formula: 1 TiB = 1024 GiB.
Subscription duration	The subscription duration can range from six months to three years. Specific prices are displayed on the buy page in the ECS console.
	The categories of disks that can be created vary based on the types of clusters. Specific prices are displayed on the buy page in the ECS console.
Unit price of the disk category	• Enhanced SSDs (ESSDs) at performance level PL0 (PL0 ESSDs) can be created in basic dedicated block storage clusters.
	 ESSDs at performance level PL1 (PL1 ESSDs) can be created in performance dedicated block storage clusters.

Billing method

Dedicated block storage clusters support only the subscription billing method.

7.3. Create a dedicated block storage cluster

After you create a dedicated block storage cluster, you can exclusively use its storage resources. This topic describes how to create a dedicated block storage cluster.

Context

Note Dedicated Block Storage Cluster is supported in the China (Heyuan), Indonesia (Jakarta), and China South 1 Finance regions.

Dedicated block storage clusters are classified into basic and performance types. When you create a dedicated block storage cluster, choose a cluster type that best suits your business needs.

You are charged for creating dedicated block storage clusters. For more information, see Billing.

Procedure

- 1. Go to the Dedicated block storage cluster page.
 - i.
 - ii. In the left-side navigation pane, choose **Storage & Snapshots > Dedicated Block Storage**Cluster

iii.

- 2. On the Dedicated Block Storage Cluster page, click **Create Dedicated Block Storage Cluster**.
- 3. On the **Create Dedicated Block Storage Cluster** page, configure the parameters described in the following table for the dedicated block storage cluster.

Parameter	Description	
Billing Method	Only the Subscription billing method is supported.	
Region and Zone	Select the region and zone in which to create the dedicated block storage cluster. Note After a dedicated block storage cluster is created, you cannot change its region or zone. You must have an understanding of the region and zone in which cloud resources are located before you select a region and zone for a dedicated block storage cluster. For example, you can select the region and zone in which Elastic Compute Service (ECS) instances are located to create a dedicated block storage cluster.	
Cluster Capacity	Set the capacity of the dedicated block storage cluster. The following limits apply: Minimum capacity: 61,440 GiB Maximum capacity: 2,359,296 GiB The capacity of each dedicated block storage cluster can be changed by multiples of 12,288 GiB. Mote If the capacity of a dedicated block storage cluster is less than 576 TiB, the maximum throughput supported per TiB does not exceed 52 MB/s. If the capacity of a dedicated block storage cluster is greater than 576 TiB, the maximum throughput supported per TiB does not exceed 26 MB/s.	
Dedicated Block Storage Cluster Name	Set the name of the dedicated block storage cluster.	
Cluster Performance Type	 Set the type of the dedicated block storage cluster. You can select Performance or Basic. Take note of the following items: Basic dedicated block storage clusters support enhanced SSDs at performance level 0 (PL0 ESSDs). Performance dedicated block storage clusters support enhanced SSDs at performance level 1 (PL1 ESSDs). For more information about ESSDs, see ESSDs. 	
Term	Set the subscription duration of the dedicated block storage cluster.	

4. Check the configuration fees, read and select ECS Service Terms, and then click Confirm Order.

The dedicated block storage cluster is delivered within 21 business days after the payment is made.

Result

After the dedicated block storage cluster is created, you can choose **Storage & Snapshots** > **Dedicated Block Storage Cluster** in the left-side navigation pane to view the information about the cluster.

What's next

After a dedicated block storage cluster is created, you can perform the following operations:

- View the basic monitoring information of the dedicated block storage cluster. For more information, see View the information of a dedicated block storage cluster.
- Create ESSDs in the dedicated block storage cluster. For more information, see Create disks in a dedicated block storage cluster.

7.4. View the information of a dedicated block storage cluster

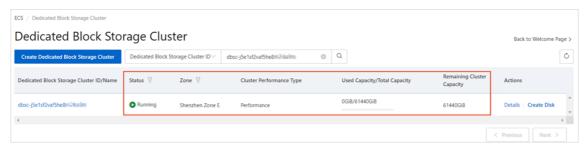
After you create a dedicated block storage cluster, you can view the basic information of the cluster and the information of disks created in the cluster.

Context

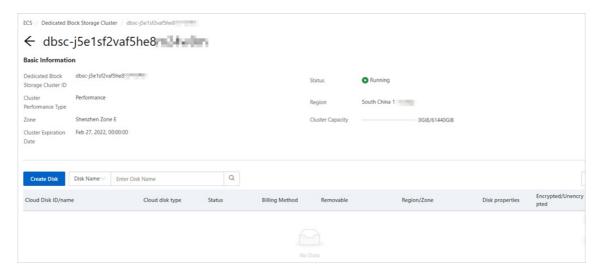
Dedicated Block Storage Cluster is available in the China (Heyuan), Indonesia (Jakarta), and China South 1 Finance regions.

Procedure

- 1.
- 2. In the left-side navigation pane, choose **Storage & Snapshots > Dedicated Block Storage Cluster**.
- 3.
- 4. View the information of a dedicated block storage cluster.
 - In the dedicated block storage cluster list, view the states, zones, and capacities of dedicated block storage clusters.



 In the Actions column corresponding to a dedicated block storage cluster, click Details to go to the cluster details page and view the basic information of the cluster. If disks are created in the dedicated block storage cluster, the information of the disks is displayed.



7.5. Create disks in a dedicated block storage cluster

After a dedicated block storage cluster is created, you can create disks in the dedicated block storage cluster. The method used to create system disks in a dedicated block storage cluster is different from that used to create data disks. This topic describes how to create disks in a dedicated block storage cluster in different scenarios.

Context

Note Dedicated Block Storage Cluster is available in the China (Heyuan), Indonesia (Jakarta), and China South 1 Finance regions.

Disks in a dedicated block storage cluster are exclusively reserved for your use and are physically isolated from other disks in the public cloud. Take note of the following items:

- The category of disks that are available for a dedicated block storage cluster depends on the cluster type. Only enhanced SSDs at performance level 0 (PL0 ESSDs) can be created in basic dedicated block storage clusters. Only enhanced SSDs at performance level 1 (PL1 ESSDs) can be created in performance dedicated block storage clusters.
- Both system disks and data disks can be created in dedicated block storage clusters, but in different ways. For more information, see the following sections:
 - o Create a system disk in a dedicated block storage cluster
 - o Create a data disk in a dedicated block storage cluster

Create a system disk in a dedicated block storage cluster

If you want to create the system disk for an Elastic Compute Service (ECS) instance in a dedicated block storage cluster, you must enable Dedicated Block Storage Cluster when you create the instance.

Note Only pay-as-you-go instances support Dedicated Block Storage Cluster. Make sure to set Billing Method to Pay-As-You-Go when you create the instance.

1.

2.

3. In the top navigation bar, select a region.

The instance must be created within the same region as the dedicated block storage cluster that you want to use.

- 4. Click Create Instance.
- 5. Enable Dedicated Block Storage Cluster.

For more information about how to create an instance, see Create an instance by using the wizard. Take note of the following parameters:

- Billing Method: Select Pay-As-You-Go.
- Region: Select the region and zone to which the dedicated block storage cluster belongs.
- Dedicated Block Storage Cluster: Select Enable Dedicated Block Storage Cluster and select the dedicated block storage cluster in which to create the system disk.

After the instance is created, the system disk of the instance belongs to the selected dedicated block storage cluster.

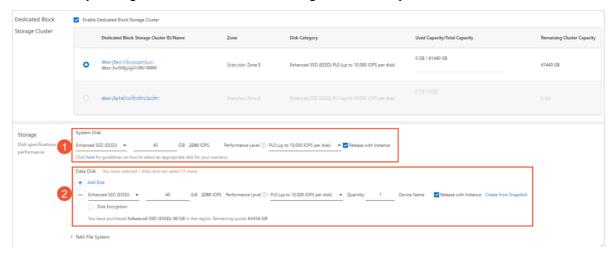
Create a data disk in a dedicated block storage cluster

You can use one of the following methods to create a data disk in a dedicated block storage cluster:

• Method 1: Enable Dedicated Block Storage Cluster when you create an instance

If you enable Dedicated Block Storage Cluster when you create an instance, the system disk and data disks of the instance automatically belong to the dedicated block storage cluster that you selected. For more information, see Create a system disk in a dedicated block storage cluster.

For example, when you create an instance, you enable Dedicated Block Storage Cluster and specify a system disk (as shown in Section ① of the following figure) and a data disk (as shown in Section ② of the following figure). After the instance is created, the system disk and data disk of the instance automatically belong to the dedicated block storage cluster that you selected.



- Method 2: Create a data disk in a dedicated block storage cluster
 - 1. Go to the Dedicated Block Storage Cluster page.

i.

ii. In the left-side navigation pane, choose **Storage & Snapshots > Dedicated Block Storage Cluster.**

iii

- 2. Find the dedicated block storage cluster in which you want to create a data disk and click **Create Disk** in the **Actions** column.
- 3. On the Create Disk page, configure the parameters described in the following table:

Parameter	Description	
Region	Select the region and zone to which the dedicated block storage cluster that you want to use belongs.	
Dedicated Block Storage Cluster	By default, Dedicated Block Storage Cluster is enabled. If Dedicated Block Storage Cluster is not enabled, select Enable Dedicated Block Storage Cluster . If Dedicated Block Storage Cluster is enabled and multiple dedicated block storage clusters belong to the selected region and zone, select the dedicated block storage cluster that you want to use from the dedicated block storage cluster list.	
Storage	Select PL0 ESSDs if the type of your dedicated block storage cluster is Basic, or select PL1 ESSDs if the type of your dedicated block storage cluster is Performance.	
Quantity	Specify the number of disks that you want to create. Note When you create disks, take note of the number of data disks on the instance to which you want to attach the disks. The number of data disks that can be attached to an ECS instance is limited. For more information, see the "Elastic Block Storage (EBS) limits" section in Limits. Pay-as-you-go disks have capacity quotas. The total capacity that you purchase cannot exceed your remaining capacity quota. You can use the following formula to calculate the total disk capacity: Total capacity = Capacity of a single disk × Number of disks. After you configure the disk category, the purchased capacity and the remaining capacity quota are displayed on the disk buy page.	
Terms of Service	Read and select ECS Terms of Service.	
Others (Optional)	Configure the Disk Name, Description, Tags, and Resource Group parameters.	

4. Click Preview. After you verify disk information in the Preview message, click Create. After the disk is created, you can view the disk information in the dedicated block storage cluster that you selected. For more information, see View the information of a dedicated block storage cluster.

What to do next

If you perform the operations described in Method 2 to create a disk in a dedicated block storage cluster, you must attach the disk to an instance within the same zone and format the disk before the disk can be used.

- 1. Attach the disk to a specified instance.
 - For more information, see Attach a data disk.
- 2. Log on to the instance and format the disk.

For more information, see Partition and format a data disk on a Linux instance or Partition and format a data disk on a Windows instance.

8. Async replication

8.1. Overview

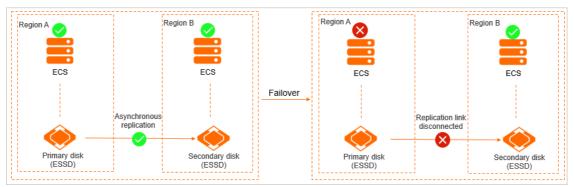
Async replication is a feature that protects data across regions or across zones within the same region based on the data replication capability of Elastic Block Storage (EBS). This feature allows data to be asynchronously replicated between disks across regions or across zones within the same region for disaster recovery. You can use this feature to implement disaster recovery for critical business to protect data in your databases and improve business continuity.

Use scenarios

The async replication feature is suitable for users that require high data security and need to implement disaster recovery and cross-region migration of business data.

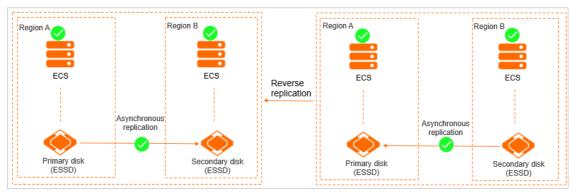
• Disaster recovery scenario

If the primary disk in a replication pair fails, you can use the failover feature to switch the primary and secondary disks over. During this switchover, the failover feature disconnects the original replication link and fails services over to the disaster recovery system by attaching the new primary disk (original secondary disk) to an Elastic Compute Service (ECS) instance that is used for disaster recovery.



• Cross-region migration scenario

If you want to migrate your business data across regions, you can use the reverse replication subfeature, instead of the image or snapshot replication feature.



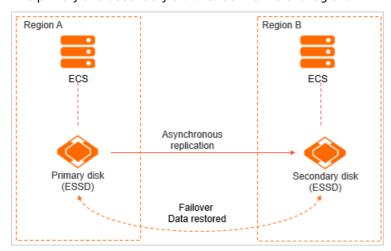
Feature description

Note The async replication feature is supported in the China (Hangzhou), China (Shanghai), China (Beijing), China (Shenzhen), China (Heyuan), China (Chengdu), China (Hong Kong), Singapore (Singapore), US (Silicon Valley), and US (Virginia) regions.

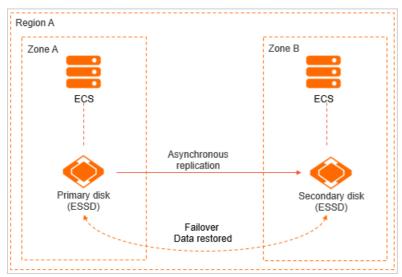
How the async replication feature works

The async replication feature allows data to be asynchronously replicated from an enhanced SSD (ESSD) (primary disk) to an ESSD (secondary disk) that has the same specifications. The primary and secondary disks can reside in different zones within the same region or reside in different regions. If the primary disk fails, you can fail over to the secondary disk. After the primary disk recovers, you can restore data from the secondary disk. The following figures show how the async replication feature works when the primary and secondary disks reside in different zones within the same region or reside in different regions.

• The primary and secondary disks reside in different regions



• The primary and secondary disks reside in different zones within the same region



Operations

The following table describes the operations supported by the async replication feature.

Operation	Description	References
Create a replication pair	Before you can use the async replication feature to implement disaster recovery across regions or across zones within the same region, you must create a replication pair.	Create a replication pair
Enable the async replication feature	After you create a replication pair, you must enable the async replication feature to replicate data from the primary disk to the secondary disk cross regions or across zones within the same region on a periodic basis.	Enable async replication for replication pairs
Implement disaster recovery	After you create and activate a replication pair, if the primary disk fails, you can use the failover and reverse replication sub-features of async replication to implement disaster recovery.	Implement disaster recovery
Delete a replication pair	The functionality of disks in a replication pair is limited. If a replication pair is no longer needed, you can delete it.	Delete a replication pair

Terms

The following table describes the terms related to the async replication feature.

Term	Description
asynchronous replication	The async replication feature supports only asynchronous replication. Unlike synchronous replication, asynchronous replication replicates data from a disk to another disk across regions or across zones within the same region on a periodic basis. In asynchronous replication, the data on the source disk may be inconsistent with that on the destination disk.
primary disk	The disk from which to replicate data for disaster recovery. The primary disk is also called the source disk.
secondary disk	The disk to which to replicate data. The secondary disk is also called the destination disk.
recovery point objective (RPO)	The amount of data that may be lost due to a disk exception. RPO is measured by time. It is used as a data metric and holds a default value of 15 minutes in async replication. This value indicates that data written to a primary disk from up to 15 minutes prior may be lost when an exception occurs on the disk.
recovery time objective (RTO)	The duration of time that it takes for a primary disk to recover after an exception occurs on the disk. RTO is used as a data metric in async replication. For example, if the value of RTO is 1 hour, it indicates that a primary disk can have its data restored and be back in service within 1 hour after an exception occurs on the disk.

Term	Description
replication pair	The relationship that is established between a primary disk and a secondary disk, and the configurations for asynchronous replication.
failover	A sub-feature of async replication that allows you to enable read and write permissions on the secondary disk and fail over to the disk.
reverse replication	A sub-feature of async replication that can reverse replication relationships to replicate data from secondary disks to primary disks.

Billing

You are charged for the use of async replication on a subscription basis and must pay bandwidth fees.

Precautions

The following table describes the limits that apply to the async replication feature.

ltem	Limits
Replication pairs that can be created per zone in an Alibaba Cloud account	16
Replication pairs that can be created for a single disk	1
Replication cycle	15 minutes (Data is asynchronously replicated from a primary disk to a secondary disk every 15 minutes.)
Primary disk category	A primary disk must be an ESSD used as a data disk and uses the pay-as-you-go billing method.
Secondary disk category	A secondary disk must be of the same disk category and have the same performance level and capacity as the associated primary disk.

The limits described in the following table apply to primary and secondary disks when you use the async replication feature.

ltem	Support by the primary disk	Support by the secondary disk
Read and write operations	√	x ^①
Disk deletion	×	×
Disk initialization	×	×
Disk resizing	x	×
Disk attaching	√ .	×
Snapshot creation	√	√ ^②

ltem	Support by the primary disk	Support by the secondary disk
Rollback based on snapshots	1	×
Disk category change	×	×
Performance level change	×	×
Disk encryption	×	×
Multi-attach	×	×
Disk migration along with instances	×	×

Note:

- 1: After a replication pair is activated, the secondary disk enters the read-only state and no users have write permissions on the disk.
- ②: Due to RPO, data of a snapshot created for a primary disk may not be consistent with that of a snapshot created at the same time for the associated secondary disk.

8.2. Create a replication pair

This topic describes how to create a replication pair. Before you can use the async replication feature to implement disaster recovery for disks across regions or across zones within the same region, you must create replication pairs.

Prerequisites

The source disk (primary disk) from which to replicate data and the destination disk (secondary disk) to which to replicate data are created. Note that the primary and secondary disks must be of the same category and have the same performance level and capacity. For more information, see Create a disk.

Context

Note The async replication feature is supported in the China (Hangzhou), China (Shanghai), China (Beijing), China (Shenzhen), China (Heyuan), China (Chengdu), China (Hong Kong), Singapore (Singapore), US (Silicon Valley), and US (Virginia) regions.

When you create a replication pair, take note of the following items:

- The first time you log on to the Elastic Block Storage (EBS) console, you must create a service-linked role for EBS as instructed. For information about the service-linked role for EBS, see Service-linked role for EBS.
- The primary disk in a replication pairs must be an unencrypted enhanced SSD (ESSD) that is used as a data disk and is not being resized, upgraded, or downgraded.
- The async replication feature replicates data from primary disks to secondary disks and the replicated data overwrites the original data on the secondary disks. We recommend that you use an empty disk as the secondary disk.

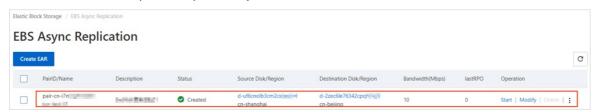
For more information, see Precautions.

Procedure

- 1. Log on to the EBS console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > EBS Async Replication**.
- 3. On the EBS Async Replication page, click Create EAR.
- 4. Configure parameters, confirm the price, and then click **Buy Now**.

Configure the following parameters:

- Region and Zone: Select the region and zone where the primary disk is located.
- o Disk: Select the ID of the primary disk.
- **Disaster Recovery Region and Zone**: Select the region and zone where the secondary disk is located.
- o Destination Disk: Select the ID of the secondary disk.
- o Relationship Pair Name: Enter a name for the replication pair.
- **Description**: Enter a description for the replication pair.
- **Bandwidth**: Select a bandwidth value based on your business needs. The greater the bandwidth value, the higher the rate at which data is asynchronously replicated.
- **Subscription Period**: Select a subscription period for the replication pair.
- 5. Confirm that the replication pair configurations and the price are correct. Then, click **Purchase**. After you complete the payment, you can go back to the EBS Async Replication page on the EBS console to view the replication pair that you created.



What's next

After the replication pair is created, you must enable the async replication feature to implement disaster recovery for disks. For more information, see Enable async replication for replication pairs.

8.3. Enable async replication for replication pairs

After you create a replication pairs, you must enable the async replication feature for the pair so that data can be asynchronously replicated from the primary disk to the secondary disk on a periodic basis. This topic describes how to enable the async replication feature.

Prerequisites

- If you want to use automatic synchronization by clicking **OK** when you enable async replication for a replication pair, make sure that the pair is in the **Created**, **Syncing**, **Normal**, or **Stopped** state.
- If you want to use manual synchronization by clicking **Copy Data** when you enable async replication for a replication pair, make sure that the pair is in the **Created**, **One-time Syncing**, or **Stopped** state.

Context

If a replication pair is added to a replication pair-consistent group, async replication cannot be separately enabled for the pair. In this case, you can only batch manage all replication pairs in the group together. For more information, see Activate async replication for replication pair-consistent groups.

Procedure

- 1. Log on to the EBS console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > EBS Async Replication**.
- 3. In the top navigation bar, select a region. Example: China (Shanghai).

 You can select the region of a primary or secondary disk to view the list of associated replication pair.
- 4. In the replication pair list of the EBS Async Replication page, find the created replication pair that you want to activate and click **Activate** in the **Operation** column.
 - ? Note If you want to batch enable async replication for multiple replication pairs, select the replication pairs and click Batch Activate in the lower part of the page.
- 5. In the message that appears, you can click Copy Data or OK to activate the replication pairs.
 - Notice After async replication is enabled for a replication pair, data stored on the secondary disk is overwritten. We recommend that you click Create Snapshot to back up data stored on the secondary disk to prevent data loss.
 - Copy Data: The system immediately replicates data once, and then stops the replication pair
 after data is replicated. Subsequently, you can manually synchronize data or have data
 automatically synchronized from the primary disk to the secondary disk based on your business
 requirements.
 - **OK**: The system periodically replicates data based on the RPO of the replication pair to synchronize data between the primary and secondary disks.

The first time you activate a replication pair, the replication pair enters the **Initial Syncing** state. Wait until data is synchronized from the primary disk to the secondary disk.



After the replication pair is activated and data between the primary disk and the secondary disk is initially synchronized, the state of the replication pair changes to **Normal**. Subsequently, the system asynchronously replicates data from the primary disk to the secondary disk at an interval of 15 minutes for disaster recovery. You can check when data was last replicated from the primary disk to the secondary disk in the **Recent recovery point** column corresponding to the replication pair.



8.4. Implement disaster recovery

After you create and activate a replication pair, if the primary disk fails, you can use the async replication feature to implement disaster recovery. This topic describes how to use the async replication feature to implement disaster recovery for a primary disk.

Prerequisites

Before you use the reverse replication sub-feature to perform a reverse replication, make sure that the primary disk has been detached from its associated Elastic Compute Service (ECS) instance and is in the Unattached state. For more information, see Detach a data disk.

Context

When you use the async replication feature to implement disaster recovery, take note of the following items:

- The async replication feature provides the failover and reverse replication sub-features. If the primary disk in a replication pair fails, you can use the failover sub-feature to enable read and write permissions on the secondary disk, attach the secondary disk to a temporary ECS instance, and then fail over to the disk. After the primary disk recovers, you can use the reverse replication sub-feature to reverse the replication relationship and replicate the latest data from the secondary disk to the primary disk for disaster recovery.
- If a replication pair is added to a replication pair-consistent group, a failover or reverse replication cannot be separately performed on the replication pair. In this case, you can only batch manage all replication pairs within the replication pair-consistent group together. For more information, see Use replication pair-consistent groups to implement disaster recovery.

Step 1: Perform a failover

The failover sub-feature allows you to enable read and write permissions on a secondary disk and perform a failover to the disk. We recommend that you create a temporary ECS instance in advance within the region and zone where the secondary disk is located. This way, if the primary disk fails, you can use the failover sub-feature to enable read and write permissions on the secondary disk, attach the secondary disk to the temporary ECS instance, and then fail over to the secondary disk. You can continue your business on the secondary disk until the primary disk recovers. For information about how to create an ECS instance, see Creation method overview.

- 1. Log on to the EBS console.
- 2. In the left-side navigation pane, choose Enterprise Feature > EBS Async Replication
- 3. In the top navigation bar, select a region. Example: China (Beijing).
 When you use failover and reverse replication to implement disaster recovery for a primary disk, you must switch to the region where the secondary disk is located.
- 4. Find the replication pair to which the faulty primary disk belongs, move the pointer over the icon in the **Operation** column, and then click **Failover**.
- 5. In the Failover message, read the notes and click OK.
 - The failover sub-feature suspends the async replication feature. Only use this sub-feature when your primary disk fails to prevent data loss.
 - After failover is enabled, the state of the replication pair changes to Failovered. At this point, you can attach the secondary disk to a temporary ECS instance and fail over to the disk.

Step 2: Perform a reverse replication

After the primary disk recovers, you can use the reverse replication sub-feature to replicate the latest data from the secondary disk to the primary disk for disaster recovery.

- In the top navigation bar, select a region. Example: China (Beijing).
 When you use failover and reverse replication to implement disaster recovery for a primary disk, you must switch to the region where the secondary disk is located.
- 2. Find the replication pair on which you have performed a failover, move the pointer over the icon in the **Operation** column, and then click **Reprotect**.
- 3. In the **Reprotect** dialog box, read the notes and click **Create Snapshot** to create a snapshot for the primary disk.
 - In reverse replication, the original data on the primary disk is overwritten by the data replicated from the secondary disk. We recommend that you create snapshots for the primary disk to back up disk data and prevent data loss. If you have manually created a snapshot for the primary disk after the disk recovers, you do not need to create a snapshot for the disk in this dialog box. You are charged for the snapshots. For information about billing details about snapshots, see Snapshots.
- After a snapshot is created for the primary disk, click OK.
 At this point, the state of the replication pair changes to Stopped.
 - Notice After the reverse replication sub-feature is used, the replication relationship is reversed. The original primary disk is automatically changed into the secondary disk and the original secondary disk is changed into the primary disk. For example, assume that before a reverse replication is performed, Enhanced SSD (ESSD) A in the China (Beijing) region is the primary disk and ESSD B in the China (Shanghai) region is the secondary disk. After a reverse replication is performed, the replication relationship is reversed and ESSD B is changed into the primary disk and ESSD A is changed into the secondary disk.
- 5. In the **Operation** column, click **Activate**.

You must activate the replication pair in this step to asynchronously replicate data from the original secondary disk to the original primary disk.

- After data is asynchronously replicated from the original secondary disk to the original primary disk, the state of the replication pair changes to **Normal** and disaster recovery is implemented.
- 6. (Optional)Revert the relationship between the disks in the replication pair.

After you perform the preceding steps to perform a reverse replication, the original relationship in the replication pair is reversed. If you want to revert the relationship, perform the following steps:

- i. View the region in the Secondary Disk/Region/Zone column corresponding to the replication pair and select the region in the top navigation bar to switch to that region.
 For example, assume that __cn-beijing __is displayed in the Secondary Disk/Region/Zone column corresponding to the replication pair, which indicates the China (Beijing) region. Select China (Beijing) in the top navigation bar. For more information about regions, see Regions and zones.
- ii. Find the replication pair on which you have performed a reverse replication. Move the pointer over the icon in the Operation column to perform a failover and reverse replication again.

iii. After the relationship in the replication pair is reverted, click **Activate** in the **Operation** column to activate the replication pair again.

8.5. Delete a replication pair

The functionality of disks in a replication pair is limited. If an existing replication pair is no longer needed to implement disaster recovery or needs to be replaced, you can delete it. This topic describes how to delete a replication pair.

Prerequisites

Before you delete a replication pair, make sure that the replication pair is stopped.

Context

When you delete a replication pair, take note of the following items:

- When you delete a replication pair, data on the primary and secondary disks is not changed. After the
 replication pair is deleted, the secondary disk becomes writable and stores data that was replicated
 up until the last asynchronous replication. When you create a replication pair again for the primary
 disk and secondary disk, the previously replicated data on the secondary disk is cleared and a full
 replication must be performed to replicate all data from the primary disk to the secondary disk.
 Proceed with caution.
- If a replication pair is added to a replication pair-consistent group, you cannot separately delete the replication pair. The replication pair can only be deleted after it is removed from the replication pair-consistent group. For more information, see Remove replication pairs.

Procedure

- 1. Log on to the EBS console.
- 2. In the left-side navigation pane, choose Enterprise Feature > EBS Async Replication.
- 3.
- 4. (Optional)On the EBS Async Replication page, find the replication pair that you want to stop and click **Stop** in the **Operation** column.
 - Note If the replication pair is already in the Stopped state, skip this step.

If you want to batch stop multiple replication pairs, select the replication pairs and click **BatchStop** in the lower part of the page.

- 5. When the replication pair is in the **Stopped** state, click **Delete** in the **Operation** column.
- 6. In the message that appears, click **OK**.

To delete the replication pair, you must complete the Short Message Service (SMS) verification as instructed.

When the replication pair is deleted, the secondary disk rolls back to the point in time when the last asynchronous replication was complete and drops all the data that is being replicated from the primary disk. The functionality limits are lifted from the secondary disk. You can attach the disk and read data from or write data to the disk.

? Note If you delete a replication pair before your subscription bandwidth has expired, to apply for a refund.

9. Replication pair-consistent groups

9.1. Overview

The replication pair-consistent group feature allows you to batch manage multiple cloud disks in disaster recovery scenarios. This feature helps ensure that the data of all disks within the same replication pair-consistent group can be restored to the same point in time to allow for the recovery of one or more instances. This topic describes the use scenarios, limits, supported operations, and relevant terms of the replication pair-consistent group feature.

Use scenarios

The replication pair-consistent group feature is applicable to the following scenarios:

• Management of virtual groups

When a business system is deployed in a cluster file system that spans multiple Elastic Compute Service (ECS) instances, all disks on the instances need to be managed together as a virtual group to ensure that each of these disks is write order-consistent in asynchronous replication. Examples of virtual group scenarios: a self-managed MySQL cluster is built on multiple ECS instances, a single Logical Volume Manager (LVM) logical volume is created across multiple disks, and MySQL or SAP HANA clusters are migrated to the cloud.

Multi-disk protection and recovery

When a business system spans multiple disks, a policy needs to be applied to protect and restore all these disks together.

• Multi-disk dat a backup

Disks on multiple ECS instances that reside within the same region need to be batch backed up with high point-in-time consistency.

• Disaster recovery of distributed application systems

Disaster recovery needs to be implemented for distributed application systems, such as super computing systems, large websites, and multi-application collaborative systems.

Feature introduction

(?) Note The replication pair-consistent group feature is supported in the China (Hangzhou), China (Shanghai), China (Beijing), China (Shenzhen), China (Heyuan), China (Chengdu), China (Hong Kong), Singapore (Singapore), US (Silicon Valley), and US (Virginia) regions.

Feature principles

The replication pair-consistent group feature allows data to be asynchronously replicated between disks across zones within the same region or across regions. When a primary site fails, you can fail over to the corresponding secondary site and restore data from the secondary site.

Supported operations

The following table describes the operations supported by the replication pair-consistent group feature.

Operation	Description	References
Create a replication pair-consistent group	Before you can use the replication pair-consistent group feature to implement multi-disk geo-disaster recovery, you must create a replication pair-consistent group. After a replication pair-consistent group is created, you can modify its name and description based on your business requirements.	Create a replication pair-consistent group
Add replication pairs	After a replication pair-consistent group is created, you can add replication pairs that replicate in the same direction as the group to the group.	Add replication pairs
Remove replication pairs	You can remove replication pairs that you no longer need from replication pair-consistent groups. When a replication pair is removed from a replication pair-consistent group, the replication pair is disassociated from the group but is not deleted.	Remove replication pairs
Activate async replication for one or more replication pair-consistent groups	After replication pairs are added to a replication pair-consistent group, you must activate these replication pairs by activating async replication for the group to asynchronously replicate data from the disks in primary site to the disks in the secondary site on a periodic basis.	Activate async replication for replication pair-consistent groups
Stop async replication for one or more replication pair- consistent groups	If you no longer need to replicate data by using a replication pair-consistent group or if you want to perform a failover on a replication pair-consistent group, you can stop the member replication pairs by stopping async replication for the group.	Disable the async replication feature for replication pair-consistent groups
Implement disaster recovery	After you create and activate a replication pair- consistent group, if disks in the primary site fail, you can use this group to batch restore the disks.	Use replication pair- consistent groups to implement disaster recovery
Delete a replication pair-consistent group	You can delete a replication pair-consistent group that you no longer need.	Delete a replication pair-consistent group

Terms

The following table describes the terms related to the replication pair-consistent group feature.

Term	Description
asynchronous replication	Asynchronous replication is an operation that replicates data between disks across zones or regions on a periodic basis. In asynchronous replication, data on the source disk (primary disk) and destination disk (secondary disk) is crashconsistent but is inconsistent in time. For more information, see Overview.

Term	Description
replication pair	A pair of disks that have an asynchronous replication relationship.
primary site	A primary data center that can run independently to support the normal operation of business in normal cases.
secondary site	The data center that serves as a disaster recovery site for a primary site. If the primary site fails, the secondary site takes over business to ensure business continuity.
recovery point objective (RPO)	The amount of data that may be lost due to a disk exception. RPO is measured in time. It is used as a data metric and holds a default value of 15 minutes in replication pair-consistent groups. This value indicates that if an exception occurs on disks in the primary site, data written to the disks from up to 15 minutes prior to the exception may be lost when the disks are restored.
failover	A feature that allows you to enable read and write permissions on the disks in the secondary site and fail over to the secondary site.
reverse replication	A feature that can reverse replication relationships to replicate the latest data from the disks in the secondary site to the disks in the primary site.

Limits

The following table describes the quotas used for the replication pair-consistent group feature.

ltem	Limit
Replication pairs that can be added to a single replication pair-consistent group	17
RPO	15 minutes (Data is restored to up to 15 minutes prior to when disks in a primary site fail.)

The limits described in the following table apply to primary and secondary disks when you use the replication pair-consistent group feature.

Item	Support by the primary disk	Support by the secondary disk
Read and write operations	√	x ^①
Disk deletion	×	×
Disk initialization	×	×
Disk resizing	×	×
Disk attaching	√ .	×
Snapshot creation	√	$\sqrt{2}$
Rollback based on snapshots	1	×

ltem	Support by the primary disk	Support by the secondary disk
Disk category change	×	×
Performance level change	×	×
Disk encryption	x	×
Multi-attach	×	×
Disk migration along with instances	×	×

Note:

- ①: After a replication pair-consistent group is activated, the secondary disks enter the read-only state and no users have write permissions on the disks.
- ②: Due to RPOs, data of a snapshot created for a primary disk may not be consistent with that of a snapshot created at the same time for the associated secondary disk.

9.2. Create a replication pairconsistent group

After you create a replication pair-consistent group, you can add multiple replication pairs to the group so that you can batch manage these replication pairs between the primary and secondary sites.

Procedure

- 1. Log on to the Elastic Block Storage (EBS) console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > Replication Pair-consistent Group**.

3.

- 4. On the Replication Pair-Consistent Group page, create Create replication group.
- 5. On the **Create Replication Pair-consistent Group** page, configure parameters and click **Confirm.**

Configure the following parameters on the Create Replication Pair-consistent Group page:

- o Group Name: Enter a name for the replication pair-consistent group.
- **Production region**: Select a region for the primary site (primary or production region).
- **Production Zone**: Select a zone for the primary site (primary or production zone).
- **Disaster Recovery Region**: Select a region for the secondary site (secondary or disaster recovery region).
- **Disaster Recovery Zone**: Select a zone for the secondary site (secondary or disaster recovery zone).
- **RPO**: the recovery point objective (RPO) of the replication pair-consistent group. By default, this parameter is set to 15 minutes.

After a replication pair-consistent group is created, you can modify its name and description based on your business requirements. Before you modify the name or description of a replication pair-

consistent group, make sure that the group is in the **Created** or **Stopped** state.

What's next

After a replication pair-consistent group is created, you can add replication pairs that replicate in the same direction as the group to the group. For more information, see Add replication pairs.

9.3. Add replication pairs

You can add replication pairs that replicate in the same direction as a replication pair-consistent group to the group so that you can manage these replication pairs as a batch by using the group.

Prerequisites

- A replication pair-consistent group is created and in the Created or Stopped state. For information about how to create a replication pair-consistent group, see Create a replication pair-consistent group.
- Replication pairs are created and in the Created or Stopped state. For information about how to create replication pairs, see Create a replication pair.

Context

When you add replication pairs to a replication pair-consistent group, take note of the following items:

• Direction of data replication

A replication pair and a replication pair-consistent group replicate in the same direction if they have the same primary region (production region), primary zone (production zone), secondary region (disaster recovery region), and secondary zone (disaster recovery zone). A replication pair can be added only to a replication pair-consistent group that replicates in the same direction as it.

• Number of replication pairs

Up to 17 replication pairs can be added to a single replication pair-consistent group.

- Impact of adding replication pairs to a replication pair-consistent group
 - The recovery point objective (RPO) of the replication pair-consistent group takes effect on the replication pairs in place of their original RPOs.
 - You can only modify the name and description of the replication pairs. You cannot start, stop, delete, or perform a failover or reverse replication on individual replication pairs. These replication pairs can only be managed as a batch by using the replication pair-consistent group.

Procedure

- 1. Log on to the Elastic Block Storage (EBS) console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > Replication Pair-consistent Group**.

3.

- 4. On the **Replication Pair-Consistent Group** page, find the replication pair-consistent group to which you want to add replication pairs and click the ID of the group.
- 5. On the details page of the replication pair-consistent group, click Add replication pair.

Note After replication pairs are added to the replication pair-consistent group, you can click Activate in the upper-right corner of the group details page to activate the replication pairs.

6. In the Add a disaster recovery pair to the consistency replication group dialog box, select replication pairs and click Confirm.

What to do next

After replication pairs are added to the replication pair-consistent group, you must activate the replication pairs by activating async replication for the group. For more information, see Activate async replication for replication pair-consistent groups.

9.4. Remove replication pairs

You can remove replication pairs from replication pair-consistent groups based on your business requirements. When a replication pair is removed from a replication pair-consistent group, the replication pair is disassociated from the group but is not deleted.

Prerequisites

Replication pairs are added to a replication pair-consistent group which is in the **Created** or **Stopped** state.

Procedure

- 1. Log on to the Elastic Block Storage (EBS) console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > Replication Pair-consistent Group**.
- 3.
- 4. On the **Replication Pair-Consistent Group** page, find the replication pair-consistent group to which you want to add replication pairs and click the ID of the group.
- 5. In the replication pair list, select the replication pair that you want to remove and click **Delete** in the **Operation** column.

Note If you want to remove multiple replication pairs at the same time, select the replication pairs and click Batch remove in the lower part of the page.

9.5. Activate async replication for replication pair-consistent groups

After you add replication pairs to a replication pair-consistent group, you must activate the replication pairs by activating async for the group to asynchronously replicate data from disks in the primary site to disks in the secondary site on a periodic basis. After replication pairs are activated, the system first performs a full synchronization to synchronize all data from disks in the primary site to disks in the secondary site and then periodically synchronizes incremental data based on the recovery point objective (RPO) of the replication pair-consistent group.

Prerequisites

- If you want to use manual synchronization by clicking **Confirm** when you activate async replication for replication pair-consistent groups, make sure that the groups are in the **Created**, **Syncing**, **Normal**, **Stopped** state.
- If you want to use manual synchronization by clicking **Copy Data** when you activate async replication for replication pair-consistent groups, make sure that the groups are in the **Created**, **One-time Syncing**, or **Stopped** state.

Procedure

- 1. Log on to the Elastic Block Storage (EBS) console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > Replication Pair-consistent Group**.

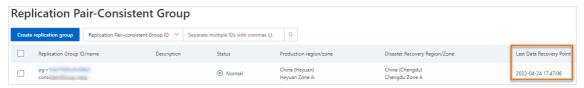
3.

- 4. You can use one of the following methods to activate async replication for replication pair-consistent groups:
 - Activate a single replication pair-consistent group.
 - In the list of replication pair-consistent groups, select the replication pair-consistent group that you want to activate and click **Activate** in the **Operation** column.
 - Activate multiple replication pair-consistent groups at the same time.
 - In the list of replication pair-consistent groups, select the replication pair-consistent groups that you want to activate and click **Batch Activate** in the lower part of the Replication Pair-Consistent Group page.
 - Activate all replication pairs in a single replication pair-consistent group.
 - Click the ID of a replication pair-consistent group. On the details page of the replication pair-consistent group, click **Activate** in the upper-right corner to activate all the replication pairs in the group.
- 5. In the message that appears, you can click Copy Data or Confirm to activate the replication pairs.
 - Copy Data: The system immediately replicates data once, and then stops the replication pairconsistent group after data is replicated. Subsequently, you can manually synchronize data or have data automatically synchronized from disks in the primary site to disks in the secondary site based on your business requirement.
 - **Confirm**: The system periodically replicates data based on the RPO of the replication pair-consistent group to synchronize data between disks in the primary and secondary sites.

The first time you activate async replication for a replication pair-consistent group, the state of the group changes to **Syncing**. You must wait until all data is synchronized.



When the state of a replication pair-consistent group changes to **Normal**, async replication is activated for group and all data on disks in the primary site has been synchronized to disks in the secondary site. Subsequently, the system asynchronously replicates incremental data from disks in the primary site to disk in the secondary site based on the RPO of the replication pair-consistent group. In the **Last Data Recovery Point** column of the replication pair-consistent group list, you can find the time when data was last replicated from disks in the primary site to disks in the secondary site.



9.6. Disable the async replication feature for replication pair-consistent groups

After replication pairs in replication pair-consistent groups are activated, if data does not need to be replicated from the primary disks to the secondary disks or if a failover needs to be performed, you can disable the async replication feature for the replication pair-consistent groups. After the async replication feature is disabled, the primary and secondary disks roll back to the point in time at which the last asynchronous replication was completed. The secondary disks drop all the data that is being replicated from the primary disks.

Prerequisites

The replication pair-consistent groups are in the **Synchronizing**, **Normal**, **Stopping**, **Stop Failed**, or **Stopped** state.

Procedure

- 1. Log on to the Elastic Block Storage (EBS) console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > Replication Pair-consistent Group**.

3.

- 4. Use one of the following methods to disable the async replication feature for one or more replication pair-consistent groups:
 - Disable the async replication feature for a single replication pair-consistent group.
 In the replication pair-consistent group list of the Replication Pair-Consistent Group page, select the desired group and click Stop in the Operation column.
 - Batch disable the async replication feature for multiple replication pair-consistent groups.
 In the replication pair-consistent group list of the Replication Pair-Consistent Group page, select the desired groups and click Batch stop in the lower part of the page.
 - o Disable the async replication feature for all replication pairs in a replication pair-consistent group.

In the replication pair-consistent group list of the Replication Pair-Consistent Group page, click the ID of the desired group. On the group details page, click **Stop** in the upper-right corner of the page.

After the async replication feature is disabled, the primary and secondary disks roll back to the point in time at which the last asynchronous replication was completed. The secondary disks drop all the data that is being replicated from the primary disks. After the preceding steps are performed, **Stopped** is displayed in the Status column corresponding to the replication pair-consistent groups.

9.7. Use replication pair-consistent groups to implement disaster recovery

After you create and activate a replication pair-consistent group, if the disks at the primary site (primary disks) fail, you can use the replication pair-consistent group to batch restore data stored on the primary disks. This topic describes how to use replication pair-consistent groups to implement disaster recovery.

Context

- Replication pair-consistent groups support the failover and reverse replication features. When the
 primary disks fail, you can use the failover feature to enable read and write permissions on data
 stored on the disks at the secondary site (secondary disks) and replicate the data from the primary
 disks to the secondary disks. Then, you can attach the secondary disks to temporary Elastic Compute
 Service (ECS) instances that are created for the failover to ensure service availability. After the issues
 of the primary disks are resolved, you can use the reverse replication feature to replicate the latest
 data stored on the secondary disks back to the primary disks.
- If you use a replication pair-consistent group to implement failover and reverse replication, the async replication feature is enabled for all replication pairs in the replication pair-consistent group. This indicates that all the replication pairs implement failover and reverse replication at the same time.
- After disaster recovery is implemented in a replication pair-consistent group, the data stored on all disks in the group is restored to the same point in time.

Step 1: Perform a failover

You can use the failover feature to enable read and write permissions on data stored on the secondary disks. We recommend that you create temporary ECS instances in the zone where the secondary disks reside based on your business requirements. If the primary disks fail, the secondary disks whose data can be read and written are attached to the temporary instances to continue to provide services until the issues of the primary disks are resolved.

Notice The failover feature suspends the async replication feature. Make sure that you only use the failover feature when your primary disks fail to prevent data loss.

- 1. Log on to the Elastic Block Storage (EBS) console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > Replication Pair-consistent Group**.

3.

- 4. On the Replication Pair-Consistent Group page, find the replication pair-consistent group to which the failed primary disks belong, move the pointer over the icon in the **Operation** column, and then click **Failover**.
 - **Note** Alternatively, click the ID of the replication pair-consistent group. On the group details page, click **Failover** in the upper-right corner of the page.
- 5. In the Failover message, read the notes and click OK.
 - After the failover is performed, **Failed Over** is displayed in the Status column corresponding to the group.
 - The failover is performed for all replication pairs in the group. At this point, you can attach the secondary disks to temporary ECS instances to continue to provide services.

Step 2: Perform a reverse replication

After the issues of the primary disks are resolved, you can use the reverse replication feature to replicate the latest data stored on the secondary disks back to the primary disks to implement disaster recovery.

Notice After a reverse replication is performed, the original data stored on the primary disks is overwritten by the data that is replicated from the secondary disks. We recommend that you create snapshots for the primary disks to prevent data loss. For more information, see Create a snapshot of a disk.

- 1. Log on to the Elastic Block Storage (EBS) console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > Replication Pair-consistent Group**.

3.

- 4. On the Replication Pair-Consistent Group page, find the replication pair-consistent group on which you have performed a failover, move the pointer over the icon in the **Operation** column, and then click **Reverse replication**.
 - **Note** Alternatively, click the ID of the replication pair-consistent group. On the group details page, click **Reverse replication** in the upper-right corner of the page.
- 5. In the Reverse replication message, read the notes and click OK.

After the reverse replication is performed, **Stopped** is displayed on the Status column corresponding to the group.

After the reverse replication is performed, the relationship of the original primary and secondary sites in the replication pair is reversed. The original primary site becomes the new secondary site, and the original secondary site becomes the new primary site. For example, assume that before a reverse replication is performed, China (Heyuan) is the primary site and China (Chengdu) is the secondary site. After the reverse replication is performed, China (Chengdu) becomes the primary site and China (Heyuan) becomes the secondary site



6. Find the replication pair-consistent group on which the reverse replication is performed and click **Activate** in the **Operation** column.

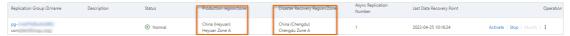
After this step is performed, data stored on the original secondary disks is asynchronously replicated back to the original primary disks.

If the data is replicated back to the original primary disks, **Normal** is displayed in the Status column corresponding to the replication pair-consistent group and disaster recovery is completed.

7. (Optional)Restore the relationship of the primary and secondary sites in the replication pairconsistent group to its original status.

After the reverse replication is performed, the relationship of the original primary and secondary sites is reversed. To restore the relationship, perform the following steps:

- i. Find the replication pair-consistent group on which the reverse replication is performed, move the pointer over the icon in the Operation column, and then subsequently click Failover to perform a failover and click Reverse replication to perform a reverse replication.
- ii. After the relationship of the primary and secondary sites in the replication pair-consistent group is restored to its original status, for example, the primary site is restored to China (Heyuan) and the secondary site is restored to China (Chengdu), click Activate in the Operation column to enable the async replication feature for the replication pair-consistent group.



9.8. Delete a replication pairconsistent group

When a replication pair-consistent group is no longer needed, you can delete it to release resources. After the replication pair-consistent group is deleted, the resources of disks at the primary and secondary sites are not deleted.

Prerequisites

- The replication pair-consistent group is in the Created, Creation Failed, Stopped, Failed Over, Deleting, Deletion Failed, or Invalid state.
- Before you delete a replication pair-consistent group, make sure that all replication pairs in the replication pair-consistent group are deleted. For more information, see Remove replication pairs.

Procedure

- 1. Log on to the Elastic Block Storage (EBS) console.
- 2. In the left-side navigation pane, choose **Enterprise Feature > Replication Pair-consistent Group**.

3.

4. On the Replication Pair-Consistent Group page, find the replication pair-consistent group that you want to delete, move the pointer over the icon in the **Operation** column, and then click

Delet e.

5. In the message that appears, click **OK**.

After the replication pair-consistent group is deleted, the secondary disks roll back to the point in time at which the last asynchronous replication was completed and drop all the data that is being replicated from the primary disks.

10.Storage capacity units 10.1. Overview

Storage capacity units (SCUs) are subscription storage resource plans that can be used to offset the pay-as-you-go bills of various storage resources. Compared with disks purchased together with subscription Elastic Compute Service (ECS) instances or resource plans of a specific resource, SCUs offer a better combination of cost-effectiveness and flexibility together with pay-as-you-go resources.

Scenarios

SCUs are applicable to scenarios that require low storage costs and high flexibility in disk creation. Examples:

- Scenarios where frequent interactions and adjustments are required in different runtime environments including development, test, and production environments, such as DevOps and microservice.
- Scenarios where multiple projects are incubated at the same time and applications are delivered or released on a frequent basis, such as container cloud-native and mobile game.

You can use SCUs in one of the following patterns:

- Pattern 1: You have created multiple pay-as-you-go disks. The historical bills of these disks show an approximate monthly average of disk space consumption. You can purchase SCUs to reduce monthly storage costs.
- Pattern 2: The cost budget for a quarter or fiscal year has been planned in advance and requires a centralized purchase or advanced payment. After you estimate the required storage capacity based on historical data and your budget, you can purchase SCUs and create pay-as-you-go disks to fulfill your storage requirements.

Comparison between different purchase methods of disks

The following table compares the purchase methods of disks.

Purchase method	Purpose	Form	Scenario
SCU	Offsets the hourly bills of pay-as-you-go disks.	A subscription storage resource plan that requires an all-upfront payment. It is used to offset the bills of pay-as-you-go disks but does not provide storage capacity.	For more information, see the Scenarios section in this topic.
Pay-as-you- go disk	Can be resized, encrypted, attached to or detached from an ECS instance, and have snapshots.	A disk that provides storage capacity and uses the payas-you-go billing method.	Applications or services with burst traffic such as temporary scaling, temporary testing, and scientific computing.

184 > Document Version: 20220712

Purchase method	Purpose	Form	Scenario
Subscription disk	Can be resized, encrypted, attached to a subscription ECS instance, and have snapshots.	A disk that provides storage capacity and is purchased together with a subscription ECS instance.	Web services that have 24/7 support, or data disks that store data for subscription ECS instances for an extended period of time.

Specifications

SCUs are measured by capacity. An SCU can have a capacity of 20 GiB, 40 GiB, 100 GiB, 500 GiB, 1 TiB, 2 TiB, 5 TiB, 10 TiB, 20 TiB, or 50 TiB. Different capacities are suited for different scenarios from daily use by developers to enterprise-level applications.

- Individual developers can purchase SCUs of 20 GiB to 500 GiB for daily use to have long-time cost-effectiveness and flexibility in the use of disks.
- SCUs of 1 TiB to 50 TiB are suitable for enterprise-level applications, which can meet the requirements that cloud-native scenarios (such as DevOps, microservice, and containerization) have for flexibility in the use of Elastic Block Storage (EBS) devices.

The following section describes the offset rules of SCUs:

- SCUs can be used to offset the bills of enhanced SSDs (ESSDs), standard SSDs, ultra disks, and basic disks. SCUs cannot be used to offset the bills of local disks.
- SCUs can be used to offset the bills of Capacity Apsara File Storage NAS file systems and Performance NAS file systems. SCUs cannot be used to offset the bills of Extreme NAS file systems or Infrequent Access (IA) storage media.
- SCUs can be used to offset the bills of snapshots.
- SCUs can be used to offset the bills of Object Storage Service (OSS) Standard, Infrequent Access, and Archive storage classes.

Billing methods

SCUs use the subscription billing method. For more information, see Storage capacity units.

Limits

The following limits apply to SCUs:

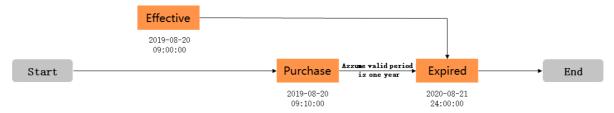
- SCUs can be used to offset the bills of only pay-as-you-go resources. SCUs cannot be used to offset the bills of pay-as-you-go disks attached to preemptible instances.
- You can configure an effective time for each SCU. The configured time must fall within six months of the SCU creation.
- You cannot call API operations to create or manage SCUs.
- The capacity limits of SCUs are determined by the capacity of EBS devices in your account. For more information about the capacity limits of SCUs that you can purchase in a region, see 使用限制.

Lifecycle management

• If you set Effective Time to Now when you purchase an SCU in the ECS console, the following results occur:

The SCU takes effect at the top of the hour that you purchase the SCU and expires at 00:00:00 on the day following the expiration date. It can offset the bills of pay-as-you-go disks starting from the time it takes effect until it expires.

Assume that you purchase a 10 TiB SCU with a validity period of one year at 09:10:00 on August 20, 2019. The SCU takes effect starting at 09:00:00 on August 20, 2019 and expires at 24:00:00 on August 21, 2020.



• If you set Effective Time to Specify Effective Time when you purchase an SCU in the ECS console, the following results occur:

The SCU takes effect at the specified time and expires at 00:00:00 on the day following the expiration date. It can offset the bills of pay-as-you-go disks starting from the time it takes effect until it expires.

Assume that you purchase a 10 TiB SCU with a validity period of one year at 09:15:00 on August 20, 2019 and set the effective time of the SCU to 01:00:00 of November 19, 2019. The SCU takes effect at 01:00:00 on November 19, 2019 and expires at 24:00:00 on November 20, 2020.



Related information

- Usage rules
- Storage capacity units
- Create an SCU

10.2. Usage rules

After you purchase a storage capacity unit (SCU) in a region, the SCU can be automatically applied to offset the bills of eligible pay-as-you-go storage resources within that region until the SCU expires. You are charged for the storage resource usage beyond the capacity of the SCU on a pay-as-you-go basis.

Deduction factors

The amount of storage capacity that an SCU can offset varies based on the types of resources to which the SCU is applied. For information about the deduction factors of different resource types, visit the SCU buy page.

Order of application

Service-specific storage resource plans take precedence over SCUs to offset the bills of storage resources within the same account. For example, if you have both Object Storage Service (OSS) resource plans and SCUs, the OSS resource plans are applied first.

Examples

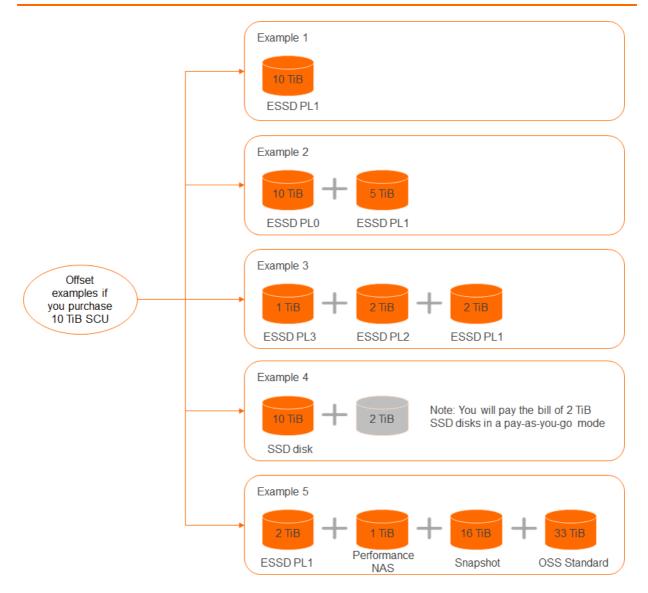
If you purchase a 10 TiB SCU in a region, the SCU can be applied to one of the following resource types to offset pay-as-you-go bills up to a specific amount of capacity:

- Disks: 2.5 TiB of enhanced SSDs (ESSDs) at performance level 3 (PL3 ESSDs), 5 TiB of ESSDs at performance level 2 (PL2 ESSDs), 10 TiB of ESSDs at performance level 1 (PL1 ESSDs), 20 TiB of ESSDs at performance level 0 (PL0 ESSDs), 10 TiB of standard SSDs, 28 TiB of ultra disks, or 33 TiB of basic disks
- Snapshots: 83 TiB of snapshots
- Apsara File Storage NAS file systems: 28 TiB of Capacity NAS file systems or 5.4 TiB of Performance NAS file systems
- OSS: 83 TiB of OSS Standard locally redundant storage (LRS) capacity, 66 TiB of OSS Standard zoneredundant storage (ZRS) capacity, 125 TiB of OSS Infrequent Access (IA) LRS capacity, 100 TiB of OSS IA ZRS capacity, or 333 TiB of OSS Archive LRS capacity
- Hybrid Backup Recovery (HBR): 33.89 TiB of HBR backup storage capacity

Note The preceding examples are for reference only. The amount of capacity for which an SCU can offset bills is calculated based on the deduction factors of the region to which the SCU belongs.

If you purchase a 10 TiB SCU when you use multiple types of storage resources, the SCU can be applied to a combination of different storage resource types.

- Example 1: The SCU is applied to offset the bills for 10 TiB of PL1 ESSDs.
- Example 2: The SCU is applied to offset the bills for 10 TiB of PL0 ESSDs and 5 TiB of PL1 ESSDs.
- Example 3: The SCU is applied to offset the bills for 1 TiB of PL3 ESSDs, 2 TiB of PL2 ESSDs, and 2 TiB of PL1 ESSDs.
- Example 4: The SCU is applied to offset the bills for 10 TiB out of 12 TiB of standard SSDs, and you are charged for the remaining 2 TiB of standard SSDs on a pay-as-you-go basis.
- Example 5: The SCU is applied to offset the bills for 2 TiB of PL1 ESSDs, 1 TiB of Performance NAS file systems, 16 TiB of snapshots, 20 TiB of OSS Standard LRS capacity, and 5 TiB of HBR backup storage capacity.



10.3. Create an SCU

This topic describes how to create a storage capacity unit (SCU) in the ECS console.

Prerequisites

You do not have overdue payments in your account.

Procedure

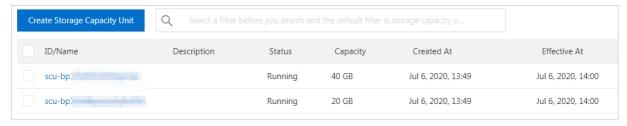
- 1.
- 2.
- 3.
- 4. Click Create Storage Capacity Unit.
- 5. Configure the following parameters on the SCU buy page.

Parameter	Description	Example
Region	The region of the SCU.	China (Hangzhou)
SCU Specification	SCU Specification The capacity of the SCU.	
Deductible Product The deductible products and the corresponding deduction factors.		None
SCU Name	The name of the SCU.	Example for SCU deduction
Quantity	The quantity of SCUs that you want to purchase.	1
Validity Period	The validity period of the SCU.	1 Month
Effective Time	 The effective time of the SCU. Valid values: Now: The SCU takes effect immediately after purchase. Specify Effective Time: The SCU takes effect at the specified time after purchase. 	Now
Terms of Service	Read and select ECS Terms of Service and Service Terms for Trial Use.	None

- 6. Verify the total price specified by **Total** and click **Purchase**.
- 7. Click Create Order and complete the payment.

Result

After you purchase an SCU, you can find it in the **Creating** state on the **Storage Capacity Units** page. When the status of the SCU changes to **Running**, the SCU begins to offset bills of eligible pay-as-you-go resources in the current region.



Block Storage·FAQ Elastic Compute Service

11.FAQ

This topic provides answers to frequently asked questions about Elastic Block Storage (EBS) devices.

- FAQ about storage capacity units (SCUs)
 - What is an SCU?
 - Which EBS devices can use SCUs?
 - o Can an SCU be used alone?
 - What rules apply when you use SCUs?
 - o How are SCUs billed?

• FAQ about enhanced SSDs (ESSDs)

- o What is an ESSD?
- What specifications do ESSDs have?
- What are the similarities and differences between ESSDs, standard SSDs, and ultra disks?
- How is the performance level of an ESSD measured?
- How do I test the performance of an ESSD?
- What is the relationship between the storage performance of an ESSD and the storage performance of the instance to which the ESSD is attached?
- o How are ESSDs billed?
- Which instance families do ESSDs support?

Common FAQ

- What must I consider when I select zones to create disks and then attach the disks to ECS instances?
- What are the common operations that can be performed on a disk?
- How do I query the usage and free space of EBS devices?
- Can I shrink a disk or a Shared Block Storage device?
- How do I release a subscription disk that has not expired?
- What is I/O optimization? Can I upgrade an existing ECS instance to an I/O optimized instance?
- How does Alibaba Cloud deal with EBS resource contention?

• FAQ about performance testing

- What tools can I use to test the performance of EBS devices?
- Why does my instance go down when I use fio to test the I/O performance of the instance?
- How do I test the performance of an ESSD?

• FAQ about standard SSDs

- What is the I/O performance of a standard SSD?
- What scenarios are standard SSDs ideal for?
- Can I replace a basic disk with a standard SSD?
- How do I purchase a standard SSD? What are the pricing options for I/O optimized instances and standard SSDs?
- Can I upgrade a standard SSD after I purchase it?

- Why is an error returned when I attempt to mount the partitions of a standard SSD to an I/O optimized Linux instance?
- What must I be aware of before I add the mount information of basic disks or standard SSDs to Linux instances?

• FAQ about attaching and detaching disks

- What is a device name (mount point)?
- What is an independent disk?
- Can I attach a disk to multiple ECS instances?
- Do I need to partition and format a pay-as-you-go disk after I purchase and attach it to an ECS instance?
- Why am I unable to find the data disk that I purchased for a Linux instance?
- How many disks can be attached to a single ECS instance?
- Why am I unable to find the desired ECS instance when I attempt to attach a disk?
- Can I attach a disk to an ECS instance that resides within a different zone?
- Will data in a data disk be lost when I detach the disk?
- Can I det ach system disks?

FAQ about independent disks

- How is a separately created pay-as-you-go data disk billed?
- Lattached a separately created disk to an ECS instance. Why is the disk released when the instance is released?
- Can I attach a separately created pay-as-you-go data disk to a subscription instance?
- Can I detach a data disk from a subscription instance?
- I changed the configurations of an instance when I renewed the instance. Can I change the billing method of a subscription disk on the instance to pay-as-you-go within the remaining time of the current subscription period?

• FAQ about disk snapshots

- When I delete a disk, will its snapshots also be deleted?
- Why are some automatic snapshots on my disk missing?
- Can I use a snapshot to create an independent disk?

• FAQ about re-initializing disks

- I cannot access the data in a Linux data disk because an error occurred when I attached the disk. What do I do?
- o If I re-initialize a disk, are my snapshots retained?
- After I restarted the instance or re-initialized the system disk, data disks of a Linux instance cannot be found. What do I do?
- How do I re-attach data disks after I re-initialize the system disk of a Linux instance?

FAQ about resizing disks

- Are my snapshots retained if I replace a system disk?
- What must I be aware of before I replace a system disk?
- How do I resize a system disk?

- Can I shrink a system disk after I resize it by extending its capacity?
- What EBS devices can be resized when they are used as system disks? Do regional limits apply to this operation?
- Can the system disks of both subscription and pay-as-you-go ECS instances be resized?
- What is the capacity range of a system disk?
- I changed the configurations of an instance when I renewed the instance. Can I specify a new size for the system disk when I replace the system disk?
- How do I create a disk from a snapshot of a data disk on an ECS instance to resize the data disk without data loss?
- What do I do if the "Bad magic number in super-block while trying to open /dev/xvdb1" error message is returned when I resize a disk of a Linux instance?
- What do I do if I have resized a disk offline but do not want to restart the instance?

• FAQ about partitions

- o Can I partition a data disk for data storage?
- For a disk with multiple partitions, are snapshots created for the entire disk or only for a specific partition?
- What must I be aware of before I re-partition a disk?
- What is the relationship between data writing and partitioning and formatting?

FAQ about rolling back disks

- I rolled back a data disk by using a snapshot after I re-partitioned the disk. How many partitions are available in the disk?
- When I roll back a disk, an error message similar to the following one is returned: "A disk can be
 rolled back only when the instance to which the disk is attached has been stopped and the disk
 has no snapshots being created. If the operating system of the current ECS instance has been
 replaced, the snapshot taken before the operating system is replaced cannot be used to roll back
 the new system disk." What do I do?

Other FAQ

- How do I migrate data from the system disk of a Linux instance?
- How do I copy data across instances?

What is an SCU?

SCUs are subscription storage resource plans that can be used to offset the pay-as-you-go bills of storage resources such as disks. Compared with disks purchased together with subscription Elastic Compute Service (ECS) instances, SCUs used together with pay-as-you-go disks offer a better combination of cost-effectiveness and flexibility. For more information, see Overview.

Which EBS devices can use SCUs?

SCUs can be used to offset the pay-as-you-go bills of eligible storage resources. Take note of the following items:

- SCUs can be used to offset the bills of enhanced SSDs (ESSDs), standard SSDs, ultra disks, and basic disks. SCUs cannot be used to offset the bills of local disks.
- SCUs can be used to offset the bills of Capacity Apsara File Storage NAS file systems and Performance NAS file systems. SCUs cannot be used to offset the bills of Extreme NAS file systems or

Infrequent Access (IA) storage media.

- SCUs can be used to offset the bills of snapshots.
- SCUs can be used to offset the bills of Object Storage Service (OSS) Standard, Infrequent Access, and Archive storage classes.

Can an SCU be used alone?

No, SCUs cannot be used alone. They must be matched to and used with pay-as-you-go disks to offset bills.

What rules apply when you use SCUs?

SCUs offset the pay-as-you-go bills of disks based on a variety of deduction factors. For more information, see Usage rules.

How are SCUs billed?

SCUs are billed based on their capacity. SCU capacity prices vary with regions.

What is an ESSD?

An ESSD is an ultra-high performance disk provided by Alibaba Cloud. ESSDs use 25 Gigabit Ethernet and remote direct memory access (RDMA) technologies to deliver up to 1 million random read/write IOPS and reduce one-way latency. For more information, see ESSDs.

What specifications do ESSDs have?

ESSDs are categorized into different performance levels and have different specifications. For information about ESSD performance levels, see ESSDs.

The performance of a storage device is closely related to the capacity of the device. A storage device that has a larger capacity provides higher data processing capabilities. All ESSDs have the same I/O performance per unit of capacity. However, the performance of ESSDs increases linearly with its capacity until the maximum performance per disk at the PL is reached.

PL	ESSD capacity range (GiB)	Maximum IOPS	Maximum throughput (MB/s)
PLO	40~32,768	10,000	180
PL1	20~32,768	50,000	350
PL2	461~32,768	100,000	750
PL3	1,261~32,768	1,000,000	4,000

What are the similarities and differences between ESSDs, standard SSDs, and ultra disks?

- Similarities: All three categories of disks are based on a distributed EBS architecture, provide high reliability and scalability, and support snapshots and data encryption.
- Differences: ESSDs have the best performance of the three disk categories. For more information, see ESSDs and EBS performance.

How is the performance level of an ESSD measured?

The performance level of an ESSD is proportional to its storage capacity. An ESSD that has a larger capacity delivers better performance. ESSDs have better performance than standard SSDs. For more information, see ESSDs.

How do I test the performance of an ESSD?

You can use fio (flexible IO tester) to perform a stress test on an ESSD. For more information, see Test the IOPS performance of an ESSD.

What is the relationship between the storage performance of an ESSD and the storage performance of the instance to which the ESSD is attached?

The storage I/O performance of specific instances is proportional to the specifications of instance types. The higher specifications an instance type has, the higher IOPS and throughput the instance type delivers.

For example, when you create an instance of the g7se storage-enhanced general-purpose instance family and attach ESSDs to the instance, the following situations may occur:

- If the total performance of ESSDs does not exceed the storage I/O performance of the instance type, the total performance of the ESSDs prevails.
- If the total performance of ESSDs exceeds the storage I/O performance of the instance type, the storage I/O performance of the instance type prevails.

Assume that you create an instance of the ecs.g7se.xlarge instance type that can deliver up to 60,000 IOPS. If you attach a 2 TiB ESSD to the instance and the ESSD can deliver up to 101,800 IOPS, the maximum IOPS of the instance is 60,000, instead of 101,800.

For information about the performance and specifications of the g7se instance family, see Instance family.

How are ESSDs billed?

ESSDs support the subscription and pay-as-you-go billing methods. For more information, see the Pricing tab of the Elastic Compute Service product page.

Which instance families do ESSDs support?

For information about the instance families that ESSDs support, see Instance family.

What tools can I use to test the performance of EBS devices?

For information about how to test the performance of EBS devices, see Test the performance of EBS devices.

Why does my instance go down when I use fio to test the I/O performance of the instance?

You can use fio to test the I/O performance of an instance by testing raw disk partitions or file systems. If you perform this test on raw disk partitions, the metadata of the file systems in the raw disk partitions may be damaged. As a result, you cannot access files in the partitions and the instance even goes down. This problem does not occur when you use fio to perform the test on file systems.

What must I consider when I select zones to create disks and then attach the disks to ECS instances?

A pay-as-you-go disk can be attached only to an ECS instance that resides within the same zone.

- For high-availability applications, we recommend that you create data disks in different zones and attach the disks to ECS instances in those zones.
- For low-latency applications, we recommend that you create data disks in the same zone as ECS instances and attach the disks to the ECS instances.

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

For information about the common operations that you can perform on a disk, see the "Related operations" section in Disks.

How do I query the usage and free space of EBS devices?

You can log on to an ECS instance to query the usage and free space of EBS devices on the instance. You cannot query the usage and free space of EBS devices by using the ECS console or by calling ECS API operations.

Can I shrink a disk?

No, disks cannot be shrunk. If you want to shrink a disk that you purchased, we recommend that you create a disk of your desired size and attach it to the same instance as the original disk. Then, copy the data stored on the original disk to the new disk and release the original disk.

How do I release a subscription disk that has not expired?

Alibaba Cloud subscription data disks cannot be released before they expire. To release a subscription data disk, you must change its billing method to pay-as-you-go. Before you release the resulting pay-as-you-go data disk, make sure that you have backed up all important data stored on it. For more information, see Change the billing methods of a disk and Release a disk.

Note After the billing method of a data disk is changed from subscription to pay-as-you-go, you are billed for the data disk on an hourly basis. One hour after the disk is released, you are no longer charged for it. After the disk billing method is changed from subscription to pay-as-you-go, the refund amount is displayed in the ECS console. The coupons that have been used are not refundable.

What is I/O optimization? Can I upgrade an existing ECS instance to an I/O optimized instance?

I/O optimization provides better network capabilities and storage performance for instances and disks. For example, you can optimize the storage performance of a standard SSD by attaching the standard SSD to an I/O optimized instance.

You can call the ModifyInstanceSpec and ModifyPrepayInstanceSpec operations to upgrade non-I/O optimized instances to I/O optimized instances.

How does Alibaba Cloud deal with EBS resource contention?

Alibaba Cloud EBS is a multi-tenant storage service in which performance standards are defined based on instance types and disk specifications. For example, disks are provided in multiple categories such as ESSDs, standard SSDs, and ultra disks and have different performance specifications. You can purchase disks of different categories to meet your data storage needs. You can also resize disks or change their categories to modify disk performance specifications.

You can use CloudMonitor to monitor disks and collect their performance data. Then, you can determine based on the collected data whether the disks suit your business requirements. For more information about CloudMonitor, see Cloud Service Monitoring.

Alibaba Cloud EBS properly allocates infrastructure to provide storage and prevent resource contention and allows disk performance data to be monitored on an ongoing basis. When hardware fails or resource contention occurs, Alibaba Cloud schedules resources or makes repairs based on the damaged condition of infrastructure.

What is the I/O performance of a standard SSD?

For more information about the I/O performance of a standard SSD, see EBS performance.

What scenarios are standard SSDs ideal for?

Standard SSDs provide high performance and high reliability that make them ideal for I/O-intensive applications with high requirements on data reliability, such as MySQL, SQL Server, Oracle, PostgreSQL, and other small and medium-sized relational databases. They are also ideal for small and medium-sized development and testing environments that require high data reliability.

Can I replace a basic disk with a standard SSD?

No, standard SSDs cannot be used to replace basic disks because standard SSDs use SSDs as physical storage media.

How do I purchase a standard SSD? What are the pricing options for I/O optimized instances and standard SSDs?

For more information about pricing, see the Pricing tab of the Elastic Compute Service product page.

Can I upgrade a standard SSD after I purchase it?

Yes, you can upgrade or resize your standard SSDs. For more information, see Overview.

Why is an error returned when I attempt to mount the partitions of a standard SSD to an I/O optimized Linux instance?

In the Linux operating system, the mount points for standard SSDs are in the /dev/vd* format, and the mount points for basic disks are in the /dev/xvd* format. If you specify the mount point in the /dev/xvd* format in a command to mount a standard SSD partition, an error is returned. You must specify the mount point in the /dev/vd* format.

What must I be aware of before I add the mount information of basic disks or standard SSDs to Linux instances?

When you attach a data disk to a Linux instance and partition and format the disk, be aware that the device names of data disks for I/O optimized instances and those for non-I/O optimized instances are different. By default, the device names of data disks are assigned by the system based on the following naming conventions:

- I/O optimized instances:
 - The device names of data disks that are attached by using the Non-Volatile Memory Express
 (NVMe) protocol use the /dev/nvmeXn1 format. Examples: /dev/nvme1n1, /dev/nvme2n1, and /d
 ev/nvme3n1. For more information about disks that support NVMe, see Overview of disks that
 support NVMe.
 - The device names of data disks that are attached not by using the NVMe protocol use the /dev/vd [b-z]format. Examples: /dev/vdb, /dev/vdc, and /dev/vdd.
- The device names of data disks that are attached to non-I/O optimized instances use the /dev/xvd[b -z]format. Examples: /dev/xvdb, /dev/xvdc, and /dev/xvdd.

Notice We recommend that you use a universally unique identifier (UUID) to reference each new partition in /etc/fstab. The device names of other disks may change due to operations such as releasing a disk. If the device name of a data disk is referenced in /etc/fstab, your stored data may be affected if the device name changes.

If invalid information is specified in the mount -a command for a disk, the disk cannot be attached. To resolve the issue, perform the following steps:

- 1. Run the fdisk -1 command to view the information of the data disk.
- 2. Check whether the information of the data disk added to viewed. is the same as that you

Note Do not add duplicate mount information because this may cause a system startup failure.

- 3. Run the vim command to modify the /etc/fstab file.
- 4. Comment out or delete invalid information and add valid mount information.
- 5. Run the mount -a command to check whether the disk is attached.

For more information, see Partition and format a data disk on a Linux instance.

What is a device name (mount point)?

A device name (mount point) is the location of an ECS disk on the disk controller bus. The selected device name matches the disk device number in Linux. The selected device name matches the disk sequence number in the disk manager of Windows.

What is an independent disk?

An independent disk is a pay-as-you-go data disk that you separately purchase. An independent disk can be attached to or detached from an ECS instance that resides within the same zone. You must attach an independent disk to an instance, and partition and format the disk before you can use it. For more information, see Create a disk.

Can I attach a disk to multiple ECS instances?

No, a disk cannot be attached to multiple ECS instances. A disk can be attached only to a single ECS instance that resides within the same zone.

Do I need to partition and format a pay-as-you-go disk after I purchase and attach it to an ECS instance?

Yes, you must attach a pay-as-you-go disk to an ECS instance, and then partition and format the disk after you separately create the disk. For more information, see Partition and format a data disk on a Linux instance and Partition and format a data disk on a Windows instance.

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

If you separately create a pay-as-you-go data disk, you must attach it to the instance and partition it before you can view and use its storage space. For more information, see Partition and format a data disk on a Linux instance and Attach a data disk.

How many disks can be attached to a single ECS instance?

When disks are used as data disks,

Why am I unable to find the desired ECS instance when I attempt to attach a disk?

Check whether the ECS instance is released. If the ECS instance is not released, make sure that it resides within the same zone as the disk.

Can I attach a disk to an ECS instance that resides within a different zone?

No, a disk cannot be attached to an ECS instance that resides within a different zone. A pay-as-you-go disk can be attached only to an ECS instance that resides within the same zone.

Will data in a data disk be lost when I detach the disk?

- In Windows, we recommend that you stop all read and write operations on all file systems of the disk to ensure data integrity. Otherwise, the data that is being read or written will be lost when you detach the disk.
- In Linux, you must first log on to the ECS instance and run the **umount** command on the disk. Then, log on to the ECS console to detach the disk.

Can I detach system disks?

Yes, system disks can be detached.

How is a separately created pay-as-you-go data disk billed?

Pay-as-you-go data disks are billed on an hourly basis. If your account balance is insufficient, the services provided by the data disks are suspended.

I attached a separately created disk to an ECS instance. Why is the disk released when the instance is released?

This is because you have configured the disk to be released along with the instance. You can change this configuration by using the ECS console or by calling an API operation. For more information, see Release a disk.

Can I attach a separately created pay-as-you-go data disk to a subscription instance?

Yes, separately created pay-as-you-go data disks can be attached to subscription instances.

Can I detach a data disk from a subscription instance?

No, data disks cannot be detached from subscription instances. Data disks expire at the same time as the subscription instances to which they are attached, and are released together with the instances. To release a subscription data disk, you can convert it into a pay-as-you-go data disk and then detach and release it. For information about how to change the billing method of disks, see Change the billing methods of a disk.

I changed the configurations of an instance when I renewed the instance. Can I change the billing method of a subscription disk on the instance to pay-as-you-go within the remaining time of the current subscription period?

No, the billing method of the disk cannot be changed within the remaining time of the current subscription period. You can change the billing method of the disk from subscription to pay-as-you-go only by changing the configurations of the instance after the current subscription period ends.

When I delete a disk, will its snapshots also be deleted?

If you have enabled the **Delete Automatic Snapshots While Releasing Disk** feature for the disk, the automatic snapshots of the disk are deleted when you delete the disk. However, the manual snapshots are retained. You can change this setting any time. For more information, see <u>Delete automatic snapshots while releasing a disk</u>.

Why are some automatic snapshots on my disk missing?

When the number of snapshots reaches the upper limit, the earliest automatic snapshots are automatically deleted but manual snapshots are not affected.

? Note The automatic snapshot policy that is applied to a disk can be executed only after the disk is attached to an instance.

Can I use a snapshot to create an independent disk?

Yes, snapshots can be used to create independent disks. You can use an existing snapshot to create an independent pay-as-you-go disk. For more information, see Create a disk from a snapshot.

I cannot access the data in a Linux data disk because an error occurred when I attached the disk. What do I do?

Perform the following operations to troubleshoot the error for the Linux operating system:

- 1. Find the data disk and use one of the following methods to check whether the disk is attached to the corresponding ECS instance:
 - View the disk in the ECS console. For more information, see View the monitoring data of a disk.
 - Log on to the instance and run the fdisk -1 command to check whether the data disk partition information is correct. Run the df -h and mount | grep "<devpath>" commands

to view mount information.

- 2. Run the **cat** command to view the /etc/fstab file and check whether two disks are attached to the same directory.
 - If two disks are attached to the same directory, the second disk replaces the first disk. This
 causes data of the first disk to become inaccessible. We recommend that you attach one of the
 disks to a different directory.
 - If two disks are attached to different directories but their mount information shows that they
 reside within the same directory, run the ll command to check whether a connection exists
 between the two directories. If a connection exists between the two directories, run the mkdir
 command to create a directory, and attach one of the disks to the new directory. Then, check
 whether the data can be accessed.

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

- Problem description: All data in a directory such as /alidata is lost after you restart a Linux instance.
- Cause: After you run the df -h command, the command output shows that none of the data disk partitions is mounted to the directory.
- Solution: An I/O optimized instance is used in this example. If your instance is a non-I/O optimized instance, enter the device names of the disk partitions in the /dev/xvd*1 format in the mount command.
 - i. Run the fdisk -1 command to view the data disk partitions that are not mounted.
 - ii. Run the mount /dev/vdb1 /alidata command to mount the data disk partitions to the /alidata directory.
 - iii. Run the def -h command to check whether the data disk partitions are mounted to the directory.
 - iv. (Optional) Configure the /etc/fstab file for the system to automatically mount the disk partitions on next system startup to prevent this issue.

If I re-initialize a disk, are my snapshots retained?

Yes, both manual and automatic snapshots of the disk are retained.

After I restarted the instance or re-initialized the system disk, data disks of a Linux instance cannot be found. What do I do?

- Problem description: After you restart a Linux instance or re-initialize the system disk, you log on to the instance and run the df -h command. The command output shows that no data disks are found.
- Cause:
 - Restart of an instance: Mount information was not written to the /etc/fstab file before you
 restart the instance. As a result, data disks are not automatically attached after the instance
 restarts.
 - Re-initialization of the system disk: The /etc/fstab file is reset when the system disk is reinitialized. As a result, data disks are not automatically attached on instance startup.
- Solution: A non-I/O optimized instance is used in this example. If your instance is an I/O optimized instance, enter the device names of the disk partitions in the /dev/vd*1 format in the mount command.

- i. Run the mount /dev/xvdb1 command to remount the data disk partitions.
- ii. Run the **mount** command to check the file system of the /dev/xvdb1 partition.
- iii. Assume that the /dev/xvdb1 partition uses the ext3 file system. Run the following command to write the partition mount information to the /etc/fstab file:

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

iv. Restart the instance in the ECS console.

How do I re-attach data disks after I re-initialize the system disk of a Linux instance?

After you re-initialize the system disk of a Linux instance, data in the data disks remains unchanged, but the mount information of the data disks is lost. Assume that before the system disk is re-initialized, a data disk partition mounted to the instance is named /dev/vdb1, and its mount point is named /InitTest. Perform the following operations to remount the data disk partition after you restart the Linux instance:

- 1. Run the mount command to view the mount information of the data disk.
 - The command output does not contain the information of /dev/vdb1.
- 2. Run the fdisk -1 command to view the information of data disk partitions.
- 3. Run the cat /etc/fstab command to view the original mount point name of the /dev/vdb1 data disk partition.
- 4. Run the mkdir /InitTest command to recreate the mount point for the data disk partition.
 - For the /dev/vdb1 data disk partition, the new mount point name must be the same as the original one.
- 5. Run the mount /dev/vdb1 /InitTest command to remount the data disk partition.
- 6. Run the df -h command to check whether the data disk partition is mounted.
- 7. Perform the following operations to check whether the /dev/vdb1 data disk partition can be automatically mounted:
 - i. Run the umount /dev/vdb1 command to unmount the /dev/vdb1 data disk partition.
 - ii. Run the mount command to check mount information.
 - If the command output does not contain the information of /dev/vdb1, the partition is unmounted.
 - iii. Run the mount -a command to automatically mount /dev/vdb1.
 - iv. Run the mount command to check mount information.

If the command output contains the information of /dev/vdb1, the partition is automatically mounted.

Are my snapshots retained if I replace a system disk?

This depends on how the snapshots are created. Manual snapshots are retained, but automatic snapshots are deleted if the Delete Automatic Snapshots While Releasing Disk feature is enabled.

Note After a system disk is replaced, the disk ID changes. You cannot use the snapshots of the original system disk to roll back the new system disk.

What must I be aware of before I replace a system disk?

We recommend that you create snapshots of the system disk before you replace it. Make sure that the new system disk has a free space of 1 GiB or larger. If the free disk space is insufficient, the instance may not start normally after you replace the system disk.

How do I resize a system disk?

You can resize a system disk by using the ECS console or by calling the ResizeDisk operation.

Can I shrink a system disk after I resize it by extending its capacity?

No, system disks cannot be shrunk after they are resized by extending their capacity. We recommend that you resize the system disk based on your needs.

What EBS devices can be resized when they are used as system disks? Do regional limits apply to this operation?

Ultra disks, standard SSDs, and ESSDs can be resized when they are used as system disks. You can resize system disks in all regions.

Can the system disks of both subscription and pay-as-you-go ECS instances be resized?

Yes, the system disks of both subscription and pay-as-you-go ECS instances can be resized.

What is the capacity range of a system disk?

The capacity range of a system disk varies based on the operating system. For more information, see Overview.

I changed the configurations of an instance when I renewed the instance. Can I specify a new size for the system disk when I replace the system disk?

After you downgrade the configurations of a subscription instance when you renew the instance, you cannot resize its system disk until the new billing cycle starts.

How do I create a disk from a snapshot of a data disk on an ECS instance to resize the data disk without data loss?

If a data disk attached to an ECS instance cannot be resized without data loss due to a disk error, you can purchase a pay-as-you-go disk to temporarily store data from the original data disk, and then format the original data disk. Perform the following operations:

- 1. Create a snapshot for the current data disk (original data disk). For more information, see Create a snapshot of a disk.
- 2. Go to the disk buy page. Select the region and zone of the ECS instance to purchase a pay-as-you-go disk. Click Create from Snapshot. In the Select Snapshot dialog box, select the snapshot created in the previous step.

- 3. Log onto the ECS console and attach the new data disk that you purchased in the previous step to the ECS instance.
- 4. Log on to the ECS instance and run the **mount** command to attach the new data disk to the instance. For more information about how to attach a disk created from a snapshot, see Create a disk from a snapshot.
- 5. Check whether files in the new data disk are the same as those in the original data disk.
- 6. Run the fdisk command to delete the original partition table. Then, run commands such as fdisk and mkf s.ext 3 to re-partition and re-format the new data disk and resize it to the desired capacity. For more information, see Resize partitions and file systems of Linux data disks.
- 7. Run the cp -R command to copy all data in the new data disk back to the original data disk.

 You can add the --preserve=all parameter to retain the file attributes when you copy the files.
- 8. Run the umount command to detach the new data disk.
- 9. Detach the new data disk from the ECS instance and release the disk in the ECS console.

What do I do if the "Bad magic number in super-block while trying to open /dev/xvdb1" error message is returned when I resize a disk of a Linux instance?

- Problem description: When you run the e2fsck -f /dev/xvdb command to resize a disk of a Linux instance, the "Bad magic number in super-block while trying to open /dev/xvdb1" error message is returned.
- Cause: The disk to be resized is not partitioned.
- Solution: Run the e2fsck -f /dev/xvdb and resize2fs /dev/xvdb commands to resize the disk. Then, run the mount command to attach the disk.

What do I do if I have resized a disk offline but do not want to restart the instance?

If you have resized a disk offline but do not want to restart the instance, you can resize the disk online again to make the offline resizing operation take effect when the instance is in the Running state. You are charged for resizing the disk online. Resize the disk based on your business requirements.

For example, assume that you have resized a disk of an instance to 60 GiB in the ECS console, but you do not want to restart the instance, the disk resizing does not take effect. You can resize the disk online again in the ECS console. For example, you can resize the disk to 61 GiB. After the disk is resized, you can see that the capacity of the disk is resized to 61 GiB in the operating system.

Can I partition a data disk for data storage?

Yes, data disks can be partitioned for data storage. You can split a data disk into multiple partitions. We recommend that you use the system tool for partitioning.

For a disk with multiple partitions, are snapshots created for the entire disk or only for a specific partition?

Snapshots are created for the entire disk, instead of for a specific partition.

What must I be aware of before I re-partition a disk?

To ensure data security, we recommend that you create a snapshot of a disk before you re-partition the disk. This way, you can roll back the disk if you perform an invalid operation. For more information, see Create a snapshot and Roll back a disk by using a snapshot.

What is the relationship between data writing and partitioning and formatting?

A new disk or disk partition can be used only after it is initialized and has its data structure recorded on the disk. The goal of formatting is to create file systems. Therefore, when a file system is created on a disk, data of the file system is written to the disk. The amount of data written to disks during formatting varies based on the file system.

- In a Windows instance, you can use one of the following methods to format a data disk:
 - Quick formatting: This method allows you to allocate only file systems to partitions and rewrite the directory table. Quick formatting takes up less space than full formatting.
 - Full formatting: This method allows you to allocate files systems to partitions, rewrite the
 directory table, and scan for and mark damaged sectors. Additionally, during the formatting
 process, empty data blocks on the disk are filled in, which has the same effect as writing data to
 the entire disk. In this case, the size of the first full snapshot is approximately equal to the disk size.
- In a Linux instance, if no data is written to a disk after you format the disk, the size of the first snapshot depends on the format of file systems on the disk.

I rolled back a data disk by using a snapshot after I re-partitioned the disk. How many partitions are available in the disk?

When you roll back a data disk by using a snapshot, the disk enters the state it was in when the snapshot was taken. If the disk has not been re-partitioned when the snapshot was taken, only one partition is available in the disk.

When I roll back a disk, an error message similar to the following one is returned: "A disk can be rolled back only when the instance to which the disk is attached has been stopped and the disk has no snapshots being created. If the operating system of the current ECS instance has been replaced, the snapshot taken before the operating system is replaced cannot be used to roll back the new system disk." What do I do?

- Problem description: When you want to roll back a disk by using a snapshot, an error message similar
 to the following one is returned: "A disk can be rolled back only when the instance to which the disk
 is attached has been stopped and the disk has no snapshots being created. If the operating system
 of the current instance has been replaced, the snapshot taken before the operating system is
 replaced cannot be used to roll back the new system disk."
- Cause: The issue may be caused by an invalid disk attribute or disk state.
- Solution: You can troubleshoot this issue based on the instance state or snapshot state.
 - Check whether the instance to which the disk is attached is stopped.

You can roll back disks only when the instance to which the disks are attached is in the Stopped state. You can log on to the ECS console and check the state of the instance on the Instances page.

• Check whether the system disk of the instance associated with the snapshot is replaced.

If you selected an image to replace the system disk, a new system disk is automatically created from the new image, and the system disk ID changes. Therefore, you cannot use the snapshots taken for the original system disk to roll back the new system disk. However, you can create a custom image from one of these snapshots and then use the custom image to replace the system disk of the instance. This way, the instance returns to the state it was in when the snapshot was taken. For more information, see Create a custom image from a snapshot and Replace the operating system of an instance by using a non-public image.

• Check whether the disk to be rolled back has a snapshot being created.

To ensure data consistency, Alibaba Cloud does not allow users to roll back a disk when a snapshot is being created from the disk. On the Instance Details page, click the Snapshot tab and check the states of snapshots. A snapshot is being created if the Progress value is not 100% and the Status value is Progressing.

If you want to forcefully terminate the creation process of a snapshot to roll back the disk, select the snapshot and click **Delete**.

How do I migrate data from the system disk of a Linux instance?

Assume that you purchase a Linux instance and do not attach data disks to it. After the instance is used for a period of time, its system disk usage approaches 100% and can no longer meet your business needs. To resolve this issue, you can purchase a data disk and attach it to the instance. Then, run the **mv** command to migrate data from the system disk to the data disk.

How do I copy data across instances?

You can use one of the following methods to copy data across instances based on the operating system:

- Copy data between Linux instances.
 - Use the lrzsz tool.

Log on to the Linux instances, install the Irzsz tool, run the rz command to upload files to one Linux instance, and then run the sz command to download the files to the other Linux instance.

You can also run the sz command to download files to your computer and then run the rz command to upload these files to the other Linux instance.

Use the FTP service.

If you use the SFTP tool, we recommend that you use the root account to log on to instances and to upload or download files.

Run the wget command.

On one instance, compress a file or a folder and then save it to the web directory to generate a download URL. Then, run the **wget** command on the other Linux instance to download the file or folder.

• Copy data between a Linux instance and a Windows instance.

We recommend that you use the SFTP tool to download files from the Linux instance to your computer and then use the FTP service to upload the files to the Windows instance.

- Copy data between Windows instances.
 - Use the FTP service.

o Use TradeManager.