# Alibaba Cloud

Elastic Compute Service Block Storage

**Document Version: 20201009** 

(-) 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 Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
- 2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
- 3. The content of this document may be changed 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 Cloudauthorized 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 conventions**

| Style           | Description   | Example  |
|-----------------|---|--|
| <u>↑</u> Danger | 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> Warning</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    | Courier font is used for commands   | Run the cd /d C:/window command to enter the Windows system folder.  |
| Italic          | 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                                    | 06   |
|----|---|------|
|    | 1.1. Block Storage overview                               | 06   |
|    | 1.2. Disk overview  | 08   |
|    | 1.3. Enhanced SSDs  | 12   |
|    | 1.4. Local disks  | 15   |
|    | 1.5. Triplicate storage                                   | 21   |
| 2. | Cloud disks   | 23   |
|    | 2.1. Create a cloud disk                                  | 23   |
|    | 2.1.1. Create a disk                                      | 23   |
|    | 2.1.2. Create a disk from a snapshot                      | 26   |
|    | 2.2. Attach a data disk                                   | 28   |
|    | 2.3. Format a data disk                                   | 33   |
|    | 2.3.1. Format a data disk for a Linux instance            | - 33 |
|    | 2.3.2. Format a data disk for a Windows ECS instance      | 39   |
|    | 2.3.3. Partition and format a data disk larger than 2 TiB | 40   |
|    | 2.3.4. Create a file system on a raw data disk            | 46   |
|    | 2.4. Change the operating system                          | 48   |
|    | 2.4.1. Replace the system disk (non-public images)        | 48   |
|    | 2.4.2. Replace the system disk (public images)            | 51   |
|    | 2.5. Reinitialize a cloud disk                            | 56   |
|    | 2.5.1. Re-initialize a system disk                        | 56   |
|    | 2.5.2. Re-initialize a data disk                          | 58   |
|    | 2.6. Roll back a disk by using a snapshot                 | 61   |
|    | 2.7. Modify the performance levels of ESSDs               | 62   |
|    | 2.8. View the monitoring data of a cloud disk             | 64   |
|    | 2.9. Query the serial number of a disk                    | 66   |

| 2.10. Change billing methods of disks 6                        | 8 |
|--|---|
| 2.11. Detach or attach a system disk 6                         | 9 |
| 2.12. Detach a data disk 7                                     | 2 |
| 2.13. Release a disk 7   | 3 |
| 3.Encrypt disks 7  | 7 |
| 3.1. Encryption overview 7                                     | 7 |
| 3.2. Encrypt a system disk7                                    | 9 |
| 3.3. Encrypt a data disk 8                                     | 2 |
| 4.Resize cloud disks 8   | 5 |
| 4.1. Overview 8  | 5 |
| 4.2. Resize disks online for Linux instances                   | 6 |
| 4.3. Resize disks offline for Linux instances9                 | 2 |
| 4.4. Resize disks online for Windows instances9                | 8 |
| 4.5. Resize disks offline for Windows instances 10             | 4 |
| 4.6. Update the Red Hat VirtIO driver of Windows ECS instan 11 | 0 |
| 5.Performance 11   | 4 |
| 5.1. EBS performance   | 4 |
| 5.2. Storage I/O performance 11                                | 8 |
| 5.3. Performance tests on Block Storage                        | 5 |
| 5.4. Test the IOPS performance of an enhanced SSD 12           | 9 |
| 6.Block Storage FAQ 13   | 4 |

## 1.Block Storage overview

## 1.1. Block Storage overview

Block Storage is a high-performance, low-latency block storage service for Alibaba Cloud ECS. It supports random read and write operations. A Block Storage device is similar to a physical disk. You can format a Block Storage device and create a file system on it to meet the data storage requirements of your business.

cloud disk EBS local disk distributed storage architecture elastic block storage local storage

#### **Block Storage devices**

Alibaba Cloud provides a variety of Block Storage devices for ECS instances, such as cloud disks based on a distributed storage architecture, and local disks located on the physical machines where the ECS instances are hosted.

| Block Storage<br>device | Description  |
|-------------------------|--|
| Cloud disk              | Cloud disks are block-level storage devices provided by Alibaba Cloud for ECS instances. Cloud disks use a triplicate distributed mechanism and feature low latency, high performance, durability, and high reliability. Disks can be created, resized, and released at any time. For more information, see Disk overview.   |
|                         | Local disks are physical disks attached to physical machines that host ECS instances. Local disks provide local storage access capabilities for ECS instances. They are designed for business scenarios that require high storage I/O performance, a large amount of storage, and high cost-effective performance. Local disks feature low latency, high random IOPS and throughput, and excellent cost-effective performance.   |
| Local disk              | Note Local disks are attached to a single physical machine, which increases the risk of single point of failure (SPOF) because the reliability of data is dependent on the reliability of the physical machine. We recommend that you implement data redundancy at the application layer to ensure data availability and enhance business continuity. You can use deployment sets to distribute ECS instances across multiple physical machines to implement high availability and disaster recovery. For more information, see Local disks. |

### Service performance

For more information about the performance of each type of Block Storage devices, see EBS performance.

### Data security

Note Except for the data erasure mechanism, only cloud disks have the following features.

#### • Data reliability during read and write operations

Three copies of your business data are stored in the Block Storage cluster in the same zone to ensure 99.999999% data reliability during read and write operations. For more information, see Triplicate storage.

#### Proactive backup

You can create snapshots at regular intervals to enhance your data security. Snapshots are a backup service provided by Alibaba Cloud. They provide data backup capabilities for cloud disks to ensure that information such as logs and customer transactions are backed up. For more information, see Snapshot overview.

#### • Data erasure mechanism

When you delete a piece of data from cloud disks and Share Block Storage devices, it is completely erased from the distributed Block Storage system and can no longer be accessed by other users in any way. The following mechanisms are used to ensure that all data is erased:

- The storage system performs sequential writes to append data to an existing file at the underlying layer of a cloud disk. This mechanism makes full use of the high-bandwidth and low-latency features of sequential write to a physical disk. Based on the features of appending data to an existing file, if you delete the logical space of a 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 the logical space of a cloud disk, the storage system does not directly overwrite the data in the logical space, but modifies the mapping between the logical space and the physical space. This ensures that data that has been overwritten can no longer be read. Data fragments that result from delete or overwrite operations are forcibly and permanently deleted from the underlying physical disks.
- When a Block Storage device (cloud disk) is released, the storage system destroys the
  metadata of the device immediately to ensure that the data can no longer be accessed. At
  the same time, the physical storage space corresponding to the disk is recycled. The
  physical space must be cleared before it is re-assigned to store data. Before data is written
  to a new disk, the system returns only zero for all read requests.

#### • Data 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 with keys based on the industry-standard AES-256 algorithm. Data is automatically encrypted when it is transmitted from ECS instances to cloud disks and automatically decrypted when the data is read. For more information, see Encryption overview.

#### Billing

For more information about the billing methods and pricing of Block Storage devices, see Billing overview.

#### Limits

For more information about the limits and quotas of Block Storage devices, see the "Block Storage limits" section in Limits.

#### Differences among Alibaba Cloud storage services

Alibaba Cloud provides the following three data storage services: Block Storage, Object Storage Service (OSS), and Apsara File Storage NAS. The following table lists the differences among these data storage services.

| Data storage<br>service    | Description   |
|----------------------------|---|
| Block Storage              | A high-performance and low-latency block-level storage device provided by Alibaba Cloud for ECS instances. It supports random read and write operations, and can be attached to ECS instances as a system disk or data disk. You can partition and format a Block Storage device and create file systems on it in the same way as you do with a physical disk. Block Storage can meet the data storage requirements of most business scenarios.   |
| OSS                        | A huge storage space designed to store unstructured data on the Internet, such as images, audios, and videos. You can access the data stored in OSS anytime and anywhere by calling API operations. OSS is applicable to the business scenarios such as website construction, separation of dynamic and static resources, and acceleration of domain name access through CDN. For more information, see What is OSS?.   |
| Apsara File<br>Storage NAS | A storage space designed for storing a large amount of unstructured data that can be accessed based on standard file access protocols, such as the Network File System (NFS) protocol for Linux, and the Server Message Block (SMB, also called Common Internet File System (CIFS)) protocol for Windows. You can set permissions to allow different clients to access the same file at the same time. Apsara File Storage NAS is applicable to the business scenarios such as file sharing across departments in an enterprise, non-linear editing in radio and television industries, high-performance computing, and containerization. For more information, see What is Apsara File Storage NAS?. |

## 1.2. Disk overview

Disks are block-level data storage products provided by Alibaba Cloud for ECS and feature low latency, high performance, high durability, and high reliability. Disks use a distributed triplicate mechanism to ensure 99.9999999% data durability for ECS instances. If service disruptions occur (for example, due to hardware failures) within a zone, data within that zone is copied to an available disk in another zone to help ensure data availability.

#### 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 the 25 Gigabit Ethernet (25 GE) and remote direct memory access (RDMA)
  technologies. Each ESSD can deliver up to one million input/output operations per second
  (IOPS) and has low latency. For more information, see Enhanced SSDs.
  - We recommend that you use ESSDs for scenarios such as online transactional processing (OLTP) databases, NoSQL databases, and Elasticsearch, Logstash, and Kibana (ELK) log analysis.
- Standard SSDs: Standard SSDs are high-performance disks that feature consistent and 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 have medium random IOPS and high data reliability.

We recommend that you use ultra disks for scenarios such as development and testing business and system disks.

• Basic disks: Basic disks are the previous generation of disks and are unavailable for purchase.

For more information about disk performance, see EBS performance.

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

- System disks contain operating systems and can only be created along with instances. The lifecycle of a system disk is the same as that of the instance to which it 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 to only one ECS instance in the same zone.

The following table describes other limits.

| Item  | Limit   | How to apply for an extension on the limit |
|---|---|--|
| Permissions to create a pay-as-<br>you-go disk              | You must complete real-name verification before you can create a disk within a region in mainland China.  | None.                                      |
| Quota for pay-as-you-go disks in all regions for an account | This quota is calculated by using the following formula: Number of ECS instances across all regions × 5. A minimum of ten pay-as-you-go disks can be created in each account. If only one instance exists in your account, you can create a maximum of ten pay-as-you-go disks in all regions, instead of five pay-as-you-go disks. | Submit a ticket.                           |

| Item   | Limit   | How to apply for an extension on the limit |
|--|---|--|
| Quota for the capacity of pay-as-<br>you-go disks that are used as data<br>disks in an account | This limit is subject to ECS resource usage, region, and disk category. You can go to the Privileges & Quotas page in the ECS console to view this information. For more information, see View and increase quotas (new version). | Submit a ticket.                           |
| Quota for system disks on an instance  | 1   | None.                                      |
| Quota for data disks on an instance  | 16.   | None.                                      |
| Capacity of a basic disk   | 5 GiB to 2,000 GiB.   | None.                                      |
| Capacity of a standard SSD   | 20 GiB to 32,768 GiB.   | None.                                      |
| Capacity of an ultra disk  | 20 GiB to 32,768 GiB.   | None.                                      |
| Capacity of an enhanced SSD (ESSD)   | 20 GiB to 32,768 GiB.   | None.                                      |
| Capacity of a local SSD  | 5 GiB to 800 GiB.   | None.                                      |
| Total capacity of all local SSDs on an instance  | 1,024 GiB.  | None.                                      |
| Capacity of a system disk  | <ul> <li>Windows Server: 40 GiB to 500 GiB</li> <li>CoreOS and FreeBSD: 30 GiB to 500 GiB</li> <li>Linux systems excluding CoreOS: 20 GiB to 500 GiB</li> </ul>   | None.                                      |
| Permissions to attach new local disks to instances that are equipped with local disks          | Not allowed.  | None.                                      |
| Permissions to change<br>configurations of instances that<br>are equipped with local disks     | Only bandwidth configurations of instances that are equipped with local disks can be changed.   | None.                                      |
| Mount points of system disks   | /dev/vda  | None.                                      |
| Mount points of data disks   | /dev/vd[b-z]  | None.                                      |

## **Billing methods**

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

- Disks that are created along with subscription instances or created separately for subscription instances are billed on a subscription basis.
- Disks that are created along with pay-as-you-go instances or created separately for pay-as-you-go instances are billed on a pay-as-you-go basis. Storage capacity units (SCUs) can be used to offset bills of pay-as-you-go disks.

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

#### **Related operations**

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

| Operation   | Reference   |
|---|---|
| Attach an idle pay-as-you-go disk to<br>an ECS instance | <ol> <li>Attach a data disk</li> <li>Format the data disk based on the operating system:         <ul> <li>Partition and format a data disk larger than 2 TiB</li> <li>Format a data disk for a Linux instance</li> <li>Format a data disk for a Windows ECS instance</li> </ul> </li> </ol> |
| Create and use a data disk                              | <ol> <li>Create a disk</li> <li>Attach a data disk</li> <li>Format the data disk based on the operating system:         <ul> <li>Format a data disk for a Linux instance</li> <li>Format a data disk for a Windows ECS instance</li> </ul> </li> </ol>                                      |
| Encrypt data stored on a disk                           | For more information, see Encryption overview.  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 about resizing, see Overview.  For more information, see the following topics:  Resize disks online for Linux instances  Resize disks online for Windows instances   |
| Replace a system disk                                   | Replace the system disk (public images)   |
| Back up disk data                                       | <ul><li> Create a normal snapshot</li><li> Apply or disable an automatic snapshot policy</li></ul>  |
| Restore a disk to its initial status                    | Re-initialize a system disk   |

| Operation   | Reference  |
|---|--|
| Restore a disk to its status at a specific point of time  | Roll back a disk by using a snapshot   |
| Detach the damaged system disk of<br>an instance and attach the disk back<br>to the original instance after the disk<br>is repaired | Detach or attach a system disk   |
| Release an instance and retain data stored on its system disk   | <ul> <li>Detach or attach a system disk</li> <li>Enable or disable the Release Disk with Instance feature for disks while you create an ECS instance</li> <li>Enable or disable the Release Disk with Instance feature for a disk on the Disks page</li> </ul> |
| Release a subscription disk that is no longer needed  | <ol> <li>Change billing methods of disks</li> <li>Detach a data disk</li> <li>Release a disk</li> </ol>  |
| Release a pay-as-you-go disk that is no longer needed   | <ol> <li>Detach a data disk</li> <li>Release a disk</li> </ol>   |

## 1.3. Enhanced SSDs

This topic provides general information about Alibaba Cloud enhanced SSDs (ESSDs), such as performance levels (PLs), usage scenarios, and performance specifications. ESSDs integrate 25 Gigabit Ethernet and remote direct memory access (RDMA) technologies to reduce latency and deliver up to 1 million random IOPS.

## **Specifications**

Note PLO ESSDs have been put into public preview. You cannot modify the PL of PLO ESSDs during the public preview.

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

| Property of ESSD | PL  |     |     |     |
|------------------|-----|-----|-----|-----|
| Property of E33b | PL3 | PL2 | PL1 | PLO |

|   | PL  |   |   |   |
|---|---|---|---|---|
| Property of ESSD  | PL3   | PL2   | PL1   | PL0   |
| Performance description   | Ultra high<br>maximum<br>concurrent I/O<br>performance<br>and ultra low<br>I/O latency                                  | High maximum<br>concurrent I/O<br>performance<br>and low I/O<br>latency   | Moderate<br>maximum<br>concurrent I/O<br>performance<br>and low I/O<br>latency  | Moderate<br>maximum<br>concurrent I/O<br>performance<br>and low I/O<br>latency  |
| Capacity (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 input/output<br>operations per second<br>(IOPS) per disk  | 1,000,000   | 100,000   | 50,000  | 10,000  |
| Maximum throughput per<br>disk (MB/s)   | 4,000   | 750   | 350   | 180   |
| Formula for calculating<br>the IOPS per disk  | min{1,800 + 50 ×<br>Capacity,<br>1,000,000}   | min{1,800 + 50<br>× Capacity,<br>100,000}   | min{1,800 + 50<br>× Capacity,<br>50,000}  | min{1,800 + 12<br>× Capacity,<br>10,000}  |
| Formula for calculating the throughput per disk (MB/s)  | min{120 + 0.5 ×<br>Capacity, 4,000}   | min{120 + 0.5 ×<br>Capacity, 750}   | min{120 + 0.5 ×<br>Capacity, 350}   | min{100 + 0.25<br>× Capacity, 180}  |
| Example scenario  | Large and medium-sized relational databases for core businesses and NoSQL databases, and large SAP and Oracle databases | Medium-sized relational databases and NoSQL databases, medium-sized 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 |
| System or data disks<br>that are recommended<br>to be replaced with<br>ESSDs in recommended<br>business scenarios | Data disks of instance families with local SSDs and 16 or more vCPUs (i1, i2, and i2g)                                  | Data disks of<br>instance<br>families with<br>local SSDs (i1,<br>i2, and i2g)   | Standard SSDs   | System disks  |

For information about how to perform stress testing on ESSDs, see Test the IOPS performance of an enhanced SSD.

#### **Billing methods**

ESSDs support pay-as-you-go and subscription billing methods. For more information, see Create a disk and Create a subscription disk.

#### **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: Elasticsearch, Logstash, and Kibana (ELK) log analysis

#### Capacity and PLs of ESSDs

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

| PL  | Capacity of ESSD (GiB) | Maximum IOPS | Maximum throughput<br>(Mbit/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 he creates an ESSD in the ECS console, PL1 is the only available option. PL1 ESSDs have a maximum IOPS of 50,000.



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



### 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 g5se storage enhanced instance family, the higher the specification an instance type has, the higher storage IOPS and throughput it can deliver. For more information, see Storage I/O performance.

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

- If the total storage performance of ESSDs does not exceed the maximum storage performance that the instance type can deliver, the instance delivers the total storage performance of the ESSDs.
- If the total storage performance of ESSDs exceeds the maximum storage performance that the instance type can deliver, the storage performance of the instance is limited to the maximum storage performance that the instance type can deliver.

The actual storage performance differs among g5se instances of different instance types that have ESSDs at different PLs attached:

- Example 1: Alex creates a 16 GiB instance of the ecs.g5se.xlarge instance type. The maximum IOPS of this instance type is 60,000. Alex then attaches to the instance a PL2 ESSD whose capacity is 2,000 GiB and maximum IOPS is 100,000. 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.g5se.4xlarge instance type. The maximum IOPS of this instance type is 230,000. Alex then attaches to the instance three PL2 ESSDs. Each ESSD has a capacity of 2,000 GiB and a maximum IOPS of 100,000. The total maximum IOPS of these ESSDs is 300,000. The maximum IOPS of the instance is limited by the maximum IOPS of the instance type to 230,000.
- Example 3: Alex creates a 64 GiB instance of the ecs.g5se.4xlarge instance type. The maximum IOPS of this instance type is 230,000. Alex then attaches to the instance a PL3 ESSD whose capacity is 2,000 GiB and maximum IOPS is 101,800. The maximum IOPS of the instance is no longer 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

The instance families supported by ESSDs, see Instance families.

## 1.4. Local disks

Local disks are disks that are attached to the same physical machine that hosts their ECS instance. Local disks provide local storage and access for ECS instances. Local disks are cost-effective and provide high random IOPS, high throughput, and low latency.

#### **Disk categories**

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

Local disks are suitable for business scenarios that require high storage I/O performance, massive storage, and high cost-effective performance. Alibaba Cloud provides the following two categories of local disks:

#### NVMe SSDs

Instances of the following instance families are attached with NVMe SSDs:

- o Instance families with local SSDs: i2, i2g, i2ne, i2gne, and i1
- Compute optimized instance family with GPU capabilities: gn5
- Visualization and compute optimized instance family with GPU capabilities: ga1

For example, instance families with local NVMe SSDs are suitable for the following scenarios:

- Services such as online gaming, e-commerce, live streaming, and media services that run I/O-intensive applications and require high I/O performance and low latency.
- Services that require high 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.

#### SATA HDDs

Instances of the d1ne and d1 instance families are attached with SATA HDDs. These disks are the preferred storage media for industries with high requirements for big data computing, storage, and analysis, such as the Internet and financial industries. The disks are suitable for business scenarios that require massive storage and offline computing. These disks can fulfill the requirements of distributed computing services such as Hadoop, which have high requirements for storage performance, storage capacity, and internal bandwidth.

#### **NVMe SSD performance**

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

| NVMe SSD<br>performance metric | Single disk performance           | 2                               | Overall instance |  |
|--------------------------------|-----------------------------------|---------------------------------|------------------|--|
|                                | ecs.i2.xlarge and ecs.i2g.2xlarge | Other i2 and i2g instance types | performance*     |  |
| Maximum capacity               | 894 GiB                           | 1,788 GiB                       | 8 × 1,788 GiB    |  |
| Maximum read IOPS              | 150,000                           | 300,000                         | 1,500,000        |  |
| Maximum read<br>throughput     | 1 GB/s                            | 2 GB/s                          | 16 GB/s          |  |
| Maximum write throughput       | 0.5 GB/s                          | 1 GB/s                          | 8 GB/s           |  |
| Access latency                 | Microsecond (μs)                  |                                 |                  |  |

<sup>\*</sup>Overall instance performance in the table applies only to ecs.i2.16xlarge. This performance represents the highest performance level for local storage in the i2 instance family.

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

| NVMe SSD<br>performance<br>metric | ecs.i2ne.xlar<br>ge and<br>ecs.i2gne.2xl<br>arge | ecs.i2ne.2xla<br>rge and<br>ecs.i2gne.4xl<br>arge | ecs.i2ne.4xla<br>rge and<br>ecs.i2gne.8xl<br>arge | ecs.i2ne.8xla<br>rge and<br>ecs.i2gne.16x<br>large | ecs.i2ne.16xl<br>arge |
|-----------------------------------|--|---|---|--|-----------------------|
| Maximum<br>capacity               | 894 GiB  | 1,788 GiB   | 2 × 1,788 GiB                                     | 4 × 1,788 GiB                                      | 8 × 1,788 GiB         |
| Maximum<br>read IOPS              | 250,000  | 500,000   | 1,000,000   | 2,000,000  | 4,000,000             |
| Maximum<br>read<br>throughput     | 1.5 GB/s   | 3 GB/s  | 6 GB/s  | 12 GB/s  | 24 GB/s               |
| Maximum<br>write<br>throughput    | 1 GB/s   | 2 GB/s  | 4 GB/s  | 8 GB/s   | 16 GB/s               |
| Access<br>latency                 | Microsecond (μs)                                 |   |   |  |                       |

Note We recommend you use the latest version of Alibaba Cloud Linux 2 to ensure maximum throughput performance of disks for Linux instances. Otherwise, the maximum IOPS may not be available for Linux instances.

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

| NVMe SSD<br>performance<br>metric | Single disk performance         | Overall instance performance***     |  |
|-----------------------------------|---------------------------------|-------------------------------------|--|
| Maximum<br>capacity               | 1,456 GiB                       | 2,912 GiB                           |  |
| Maximum IOPS                      | 240,000                         | 480,000                             |  |
| Write IOPS**                      | min{165 × Capacity, 240000}     | 2 × min{165 × Capacity, 240000}     |  |
| Read IOPS**                       | mint 103 × capacity, 240000     | 2 × IIIII { 103 × Capacity, 240000} |  |
| Maximum read<br>throughput        | 2 GB/s                          | 4 GB/s                              |  |
| Read<br>throughput**              | min{1.4 × Capacity, 2000} MB/s  | 2 × min{1.4 × Capacity, 2000} MB/s  |  |
| Maximum write throughput          | 1.2 GB/s                        | 2.4 GB/s                            |  |
| Write<br>throughput**             | min{0.85 × Capacity, 1200} MB/s | 2 × min{0.85 × Capacity, 1200} MB/s |  |

| NVMe SSD<br>performance<br>metric | Single disk performance | Overall instance performance*** |
|-----------------------------------|-------------------------|---------------------------------|
| Access latency                    | Microsecond (µs)        |                                 |

<sup>\*\*</sup>The performance of a single NVMe SSD is calculated as follows:

- o Write IOPS: 165 IOPS per GiB, up to 240,000 IOPS
- o Write throughput: 0.85 MB/s per GiB, up to 1,200 MB/s

Note To obtain the standard performance data and measure the Quality of Service (QoS) of Alibaba Cloud local disks, you can test the bandwidth, IOPS, and latency of NVMe SSDs. For more information, see Performance tests on Block Storage.

#### **SATA HDD performance**

The following table describes the performance metrics of SATA HDDs.

| SATA HDD performance metric | Single disk performance | Overall instance performance**** |
|-----------------------------|-------------------------|----------------------------------|
| Maximum capacity            | 5,500 GiB               | 154,000 GiB                      |
| Maximum throughput          | 190 MB/s                | 5,320 MB/s                       |
| Access latency              | Millisecond (ms)        |                                  |

<sup>\*\*\*\*</sup>Overall instance performance in the table applies only to ecs.d1.14xlarge and ecs.d1ne.14xlarge. This performance represents the highest performance level for local storage in the d1 and d1ne 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.

#### **Instructions**

• Local disks are attached to a single physical server, which increases the risk of single point of failure (SPOF). The reliability of data stored on local disks depends on the reliability of the physical server.

<sup>\*\*\*</sup>Overall instance performance in the table applies only to ecs.i1.14xlarge. This performance represents the highest performance level for local storage in the i1 instance family.

☐ Warning For example, data stored on local disks may be lost when a hardware failure occurs. We recommend that you do not use local disks for long-term data storage.

- 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 to achieve high availability and disaster recovery.
   For more information, see Create a deployment set.
- If your applications do not have data reliability architecture, we recommend that you use cloud disks or the backup service in your ECS instances for data reliability.
   For more information, see Disk overview.
- After you purchase an ECS instance that is attached with a local disk, you must log on to the
  instance to partition and format the local disk. For more information, see Partition and format
  the local disk.
- The following operations are not supported on a local disk:
  - Create a separate local disk.
  - Use a snapshot to create a local disk.
  - o Attach a local disk.
  - Detach and release a local disk separately.
  - o Resize a local disk.
  - Reinitialize a local disk.
  - Create a snapshot for a local disk.
  - Use a snapshot to roll back a local disk.

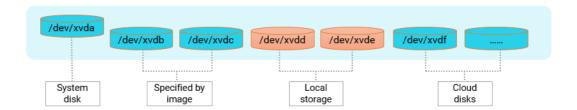
#### Disk initialization sequence

When you create an ECS instance that is equipped with local disks, all disks are initialized based on the following rules:

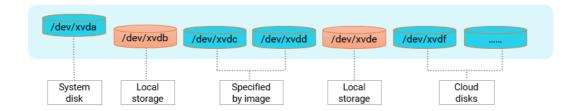
- Rule 1: If the specified image does not have data disk snapshots, the local disks are initialized prior to the cloud disks that are created along with the ECS instance.
- Rule 2: If the specified image has data disk snapshots, the initialization sequence of the data disks corresponds to the sequence of data disks retained in the snapshots. The initialization sequence of the remaining disks is subject to Rule 1.

For example, for a Linux image that contains snapshots of two data disks, the disks are initialized in the following sequence:

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



• If the original device names of the two data disks are /dev/xvdc and /dev/xvdd, Alibaba Cloud first allocates /dev/xvdc and /dev/xvdd to specified data disks in the image. The remaining device names are preferentially allocated to the local disks. The system disk is initialized first. Then, the initialization proceeds in the following sequence: local disk 1, data disk 1 specified in the image, data disk 2 specified in the image, local disk 2, cloud disk 1, cloud disk 2, and cloud disk N, as shown in the following figure.



### Lifecycle

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

### Impacts of instance operations on the 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 the local disk | 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 restore an instance.  | Erased                        | Released   |
| Release an instance.  | Erased                        | Released   |
| A subscription instance is stopped when it expires or a pay-as-you-go instance is stopped when there is an overdue payment. After that, the instance is released. | Erased                        | Released   |

| Instance operation   | Data stored on the local disk | Local disk |
|--|-------------------------------|------------|
| Renew an expired subscription instance.  | Retained                      | Retained   |
| Reactivate a pay-as-you-go instance that is stopped due to an overdue payment. | Retained                      | Retained   |

#### References

For more information about phased-out local SSDs, see Ephemeral SSDs.

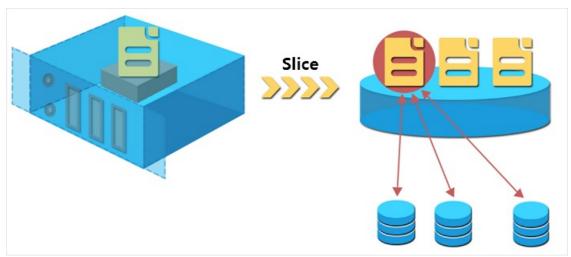
## 1.5. Triplicate storage

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%.

Alibaba Cloud triplicate storage distributed System

#### 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 ECS instance.

#### **Context**

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

| Billing method     | Limit  |
|--------------------|--|
| Pay-as-you-go disk | <ul> <li>The total number of pay-as-you-go data disks that you create in a region in your account cannot be greater than five times the total number of pay-as-you-go instances in that region in your account. For more information, see Limits.</li> <li>For more information, see Pay-as-you-go.</li> </ul>   |
| Subscription disk  | <ul> <li>When you create a subscription disk, you must attach it to a subscription instance. You cannot create standalone subscription disks.</li> <li>Subscription disks cannot be detached. They 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.</li> </ul> |

#### Additionally, you must take note of the following items:

- You cannot merge multiple disks by formatting them because they are independent of each other. 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 only back up data of a single disk. If you create a logical volume on multiple disks by using LVM, data discrepancies will occur when you roll back these disks.

#### **Procedure**

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
- 3. In the upper-right corner of the Disks page, click Create Disk.
- 4. On the disk creation page, configure parameters for the disk.



| Parameter      | Description  |
|----------------|--|
|                | Specifies whether to attach the created disks to ECS instances.  |
|                | <ul> <li>Not Attach: only creates disks and does not attach the disks to ECS<br/>instances.</li> </ul>   |
| Attach         | After you select this option, you can create only pay-as-you-go disks The disks and the ECS instances to which you want to attach the disks must be in the same zone. Proceed with caution when you specify Region.  |
|                | <ul> <li>Attach to ECS Instance: creates disks and attaches the disks to the<br/>specified ECS instance.</li> </ul>  |
|                | After you select this option, you must specify ECS Instance by selecting a destination region and an ECS instance.   |
|                | Specifies the billing method of the disks.   |
| Billing Method | <ul> <li>Pay-As-You-Go: Pay-as-you-go disks can be attached to subscription<br/>or pay-as-you-go instances.</li> </ul>   |
| <b>-</b>       | <ul> <li>Subscription: Subscription disks must be attached to subscription instances.</li> </ul>   |
|                | Select a disk category and specify the disk capacity.  |
|                | The following section describes other information:   |
| Storage        | <ul> <li>Performance Level: You can set performance levels only for<br/>enhanced SSDs (ESSDs). You can select a performance level based of<br/>the capacity of an ESSD. The disk performance varies depending on<br/>the capacity of the ESSD and the performance level. For more<br/>information, see Enhanced SSDs.</li> </ul> |
|                | <ul> <li>Create from Snapshot: You can select a snapshot to create disks.</li> <li>Then, the created disks contain the data from the snapshot.</li> </ul>  |
|                | For more information, see Create a disk from a snapshot.   |
|                | <ul> <li>Disk Encryption: Disk encryption is suitable for scenarios that require<br/>data security and regulatory compliance. If you select this option,<br/>data stored on the disk is automatically encrypted.</li> </ul>  |

| Parameter            | Description  |  |
|----------------------|--|--|
|                      | Select the number of disks that you want to purchase.  |  |
|                      | ? Note   |  |
|                      | <ul> <li>You can attach a maximum of 16 data disks to an ECS<br/>instance. When you purchase disks, note the number of data<br/>disks on the instance to which you want to attach the disks.</li> </ul>  |  |
| Quantity             | <ul> <li>Pay-as-you-go disks have capacity quotas. The total capacity that you purchase cannot exceed the disk capacity quota. You can use the following formula to calculate the total disk capacity: Total capacity = Capacity of a disk × Number of disks. After you configure the disk category, the purchased capacity and the remaining capacity quota are displayed on the disk creation page.</li> </ul> |  |
| ECS Terms of Service | Read and select ECS Terms of Service.  |  |
| Others (Optional)    | Specify the disk name, description, tags, and resource group for easy management.  |  |

- 5. Confirm the configuration information and disk fees, and then click Preview.
- 6. In the Preview dialog box, confirm the purchase information and click Create. After the disks are created, you can view the created disks on the Disks page. However, you cannot use the disks on ECS instances. If you want to use the disks, you must perform subsequent operations.

#### What's next

The following table describes the operations to be performed if you want to use the created disks.

| Scenario   | What to do next  |
|--|--|
| If you select Attach to ECS<br>Instance when you create the<br>disks | <ul> <li>Perform the following operations to partition and format the disks:</li> <li>Use the GUID Partition Table (GPT) format to partition the disks. Disks larger than 2 TiB in size support the GPT partition format. For more information, see Partition and format a data disk larger than 2 TiB.</li> <li>Use the Master Boot Record (MBR) format to partition the disks. Disks larger than 2 TiB in size do not support the MBR partition format. For more information, see Format a data disk for a Linux instance or Format a data disk for a Windows ECS instance.</li> </ul> |

| Scenario  | What to do next  |
|---|--|
|   | Perform the following operations to attach, partition, and format data disks:  1. Attach a data disk   |
|   | 2. Perform the following operations to partition and format the disks:   |
| If you select <b>Not Attach</b> when you create the disks | <ul> <li>Use the GUID Partition Table (GPT) format to partition the<br/>disks. Disks larger than 2 TiB in size support the GPT<br/>partition format. For more information, see Partition and<br/>format a data disk larger than 2 TiB.</li> </ul>  |
|   | <ul> <li>Use the Master Boot Record (MBR) format to partition the<br/>disks. Disks larger than 2 TiB in size do not support the MBR<br/>partition format. For more information, see Format a data<br/>disk for a Linux instance or Format a data disk for a<br/>Windows ECS instance.</li> </ul> |
|   | 0  |

#### **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 accidentally released disk to perform data extraction or fault analysis, you can create disks from snapshots of system or data disks.

snapshot new disk ECS Alibaba Cloud disk

#### **Prerequisites**

- A snapshot is created for a system or data disk. The snapshot ID is obtained. For more information, see Create a snapshot.
- The disks created in this example are billed on a pay-as-you-go basis. For more information about how to create a subscription disk, see Create a subscription disk.

Note You can also create disks from snapshots when you create an ECS instance. The disks have the same billing method as the instance.

#### **Context**

Take note of the following items when you create a disk from a snapshot:

- You cannot merge multiple disks by formatting them because they are independent of each other. 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 only back up data of a single disk. If you create a logical volume on multiple disks by using LVM, data discrepancies will occur when you

roll back these disks.

#### **Procedure**

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
- 3. In the upper-right corner of the Disks page, click Create Disk.
- 4. On the disk creation page, perform the following operations:
  - i. Select a region and a zone. If you want to attach the disk to an ECS instance, the disk and the instance must reside within the same region and zone.
  - ii. Configure the disk category, capacity, quantity, and name.

The following table describes some parameters on the disk creation page.

| Parameter            | Description  | Reference     |
|----------------------|--|---------------|
| Disk category        | Supports enhanced SSDs (ESSDs), standard SSDs, and ultra disks.  | Disk overview |
| Capacity             | <ul> <li>The following section describes disk capacities of different snapshot types:</li> <li>Normal snapshot: supports a disk capacity that is at least the snapshot size. If the disk capacity that you set is greater than the snapshot size, you must re-partition the disk to ensure that you can use the excess capacity.</li> </ul>  |               |
|                      | Notice If the snapshot size is less than 2,048 GiB and you want to set the disk capacity to a value greater than 2,048 GiB, ensure that the source disk of snapshot uses the GUID Partition Table (GPT) partition format. If the source disk does not use GPT, we recommend that you set the disk capacity to a value less than 2,048 GiB to avoid data loss that may occur during partitioning. For more information, see Partition and format a data disk larger than 2 TiB.   |               |
|                      | Local snapshot: supports only a capacity equal to<br>the snapshot size.  |               |
| Performance<br>Level | You can set performance levels for ESSDs based on their configured capacity. The performance levels include PL1, PL2, and PL3.   |               |
|                      | Note You can set performance levels only for ESSDs.  | Enhanced SSDs |
| Performance          | Notice If the snapshot size is less than 2,048 GiB and you want to set the disk capacity to a value greater than 2,048 GiB, ensure that the source disk of snapshot uses the GUID Partition Table (GPT) partition format. If the source disk does not use GPT, we recommend that you set the disk capacity to a value less than 2,048 GiB to avoid data loss that may occur during partitioning. For more information, see Partition and format a data disk larger than 2 TiB.  Local snapshot: supports only a capacity equal to the snapshot size.  You can set performance levels for ESSDs based on their configured capacity. The performance levels include PL1, PL2, and PL3. | overview      |

| Parameter               | Description  | Reference                                |
|-------------------------|--|--|
| Create from<br>Snapshot | Click <b>Create from Snapshot</b> and select a snapshot a system or data disk.   | N/A                                      |
| Disk Encryption         | Disk encryption is suitable for scenarios that require data security and regulatory compliance. If you select this option, data stored on the disk is automatically encrypted.                               | Encryption<br>overview                   |
| Disk Backup             | If you select this option, the automatic snapshot policy is executed and data stored on the disk is backed up on a regular basis.  | Automatic<br>snapshot<br>policy overview |
| Quantity                | The total number of pay-as-you-go data disks that you create under your account in a region cannot be greater than five times the total number of pay-as-you-go instances under your account in that region. | Limits                                   |

- iii. Bind one or more tags to the disk.
- iv. Add the disk to a resource group.
- v. Select ECS Terms of Service.
- vi. Confirm the configuration and Total fees, and click Preview. Then, follow the tips to create the disk.

#### Result

After the disk is created, go back to the Disks page and refresh the disk list. Find the new disk. Unattached is displayed in the Status column corresponding to the disk.

#### What's next

The disk can be used only after it is attached to an ECS instance. For a Linux instance, you must log on to the instance and run the mount command to mount the disk. For more information, see Attach a data disk.

#### **Related information**

CreateDisk

## 2.2. Attach a data disk

You can create a separate pay-as-you-go disk and then attach the disk to an ECS instance as a data disk.

#### **Prerequisites**

- The instance to which to attach the disk and the disk are in the same zone.
- The instance is in the Running (*Running*) or **Stopped** (*Stopped*) state, not in the **Locked** (*Locke d*) state.
- The disk is in the Unattached (Available) state.
- You have no overdue payments in your account.

#### **Context**

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. You do not need to manually attach the disks.
- You can attach a maximum of 16 data disks to an ECS instance.
- A disk can be attached only to a single instance.

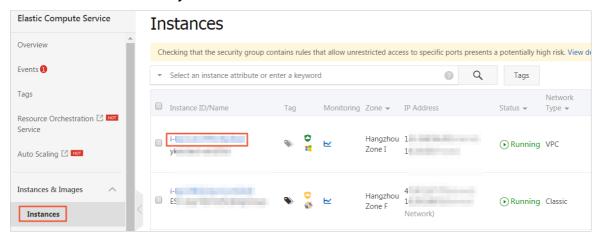
Note The attach operation refers to attaching disks to ECS instances in the console, instead of mounting file systems to the operating systems of ECS instances by running the mount command. For more information, see Format a data disk for a Linux instance and Format a data disk for a Windows ECS instance.

You can also attach disks on the Instances or Disks page.

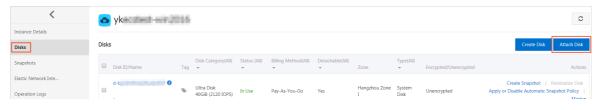
- If you want to attach multiple disks to an ECS instance, we recommend that you perform the operation on the Instances page. For more information, see Attach disks on the Instances page.
- If you want to attach multiple disks to different ECS instances, we recommend that you
  perform the operation on the Disks page. For more information, see Attach disks on the Disks
  page.

#### Attach disks on the Instances page

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the instance to which you want to attach disks and click the ID of the instance.



5. In the left-side navigation pane, click **Disks**. On the **Disks** page, click **Attach Disk** in the upper-right corner.



Attach Disk

Your Instance: EcsQuickStart (Zone: us-east-1b)

Available Mount Points:

\*Target Disk: Enter a disk ID

Release Disk with Instance

Delete Automatic Snapshots While Releasing Disk

OK

Note: After a disk is attached, you must log on to the instance, partition and format the disk, and mount new disk partitions. Tip: Partition and Format/Mount the Data Disk

Attach

Cancel

6. In the Attach Disk dialog box, configure the related parameters to attach the disk.

i. Select the disk that you want to attach and specify the release settings for the disk.

| Parameter  | Description  |  |
|--|--|--|
| Target Disk  | Select the disk that you want to attach.   |  |
|  | If you select this option, the disk will be automatically released when the instance is released. If you do not select this option, the disk will be retained when the instance is released.   |  |
| Release Disk with Instance                         | Note If the disk that you want to attach is a<br>system disk that you detached from another instance,<br>the instance specified by Release Disk with Instance<br>refers to the instance from which you detached the disk,<br>not the current instance. |  |
| Delete Automatic Snapshots<br>While Releasing Disk | If you select this option, automatic snapshots of the disk will<br>be released when the disk is released. We recommend that<br>you do not select this option to retain backup data.  |  |

- ii. Click OK.
- iii. Click Attach.

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

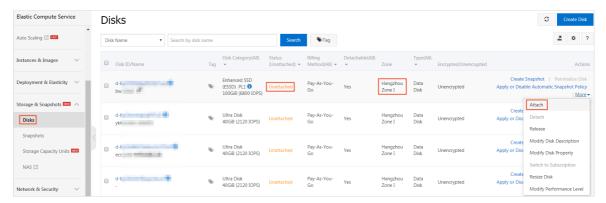
7. Create partitions and file systems after the disk is attached to the instance. The following table describes the operations to be performed based on different scenarios.

| Disk status ' | erating<br>stem | What to do next |
|---------------|-----------------|-----------------|
|---------------|-----------------|-----------------|

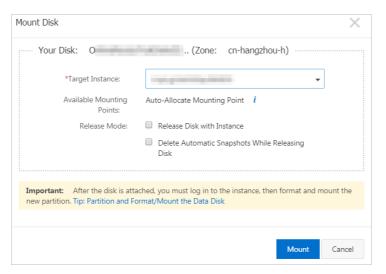
| Disk status                     | Operating system | What to do next  |
|---------------------------------|------------------|--|
| New empty<br>disks              | Linux            | <ul> <li>For information about operations on disks that are smaller than 2 TiB, see Format a data disk for a Linux instance.</li> <li>For information about operations on disks that are larger than 2 TiB, see Partition and format a data disk larger than 2 TiB.</li> </ul>       |
|                                 | Windows Server   | <ul> <li>For information about operations on disks that are smaller than 2 TiB, see Format a data disk for a Windows ECS instance.</li> <li>For information about operations on disks that are larger than 2 TiB, see Partition and format a data disk larger than 2 TiB.</li> </ul> |
| Disks created<br>from snapshots |                  | Connect to the instance and run the following command to mount the disk partitions that have available file systems:  mount <data disk="" partition=""> <mount point=""></mount></data>  |
|                                 |                  | N/A  |

## Attach disks on the Disks page

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
- 3. In the top navigation bar, select a region.
- 4. Find the disk in the Unattached state and choose More > Attach in the Actions column.



5. In the dialog box that appears, configure the following parameters.



i. Select the disk that you want to attach and specify the release settings for the disk.

| Parameter  | Description  |  |
|--|--|--|
| Target Instance                                    | Select an ECS instance to which you want to attach disks.  |  |
|  | If you select this option, the disk will be automatically released when the instance is released. If you do not select this option, the disk will be retained when the instance is released.   |  |
| Release Disk with Instance                         | Note If the disk that you want to attach is a<br>system disk that you detached from another instance,<br>the instance specified by Release Disk with Instance<br>refers to the instance from which you detached the disk,<br>not the current instance. |  |
| Delete Automatic Snapshots<br>While Releasing Disk | If you select this option, automatic snapshots of the disk will<br>be released when the disk is released. We recommend that<br>you do not select this option to retain backup data.  |  |

#### ii. Click Attach.

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

6. Create partitions and file systems after the disk is attached to the instance. The following table describes the operations to be performed based on different scenarios.

| Disk status | Operating system | What to do next  |
|-------------|------------------|--|
|             | Linux            | <ul> <li>For information about operations on disks that are smaller than 2 TiB, see Format a data disk for a Linux instance.</li> <li>For information about operations on disks that are larger than 2 TiB, see Partition and format a data disk larger than 2 TiB.</li> </ul> |

| New empty<br>Diskstatus                             | Operating system  | What to do next  |
|---|---|--|
|   | Windows Server  | <ul> <li>For information about operations on disks that are smaller than 2 TiB, see Format a data disk for a Windows ECS instance.</li> <li>For information about operations on disks that are larger than 2 TiB, see Partition and format a data disk larger than 2 TiB.</li> </ul> |
| Disks created from snapshots  Linux  Windows Server | Connect to the instance and run the following command to mount the disk partitions that have available file systems:  mount <data disk="" partition=""> <mount point=""></mount></data> |  |
|   | Windows Server  | N/A  |

#### **Related information**

#### References

AttachDisk

•

## 2.3. Format a data disk

## 2.3.1. Format a data disk for a Linux instance

After you attach a new data disk to an ECS instance, you must create at least one file system on the data disk and mount the file system. In this example, an I/O optimized instance that runs the Aliyun Linux 2.1903 LTS 64-bit operating system is used. A new 20 GiB data disk is attached to the instance. The device name of the data disk is /dev/vdb. A single Master Boot Record (MBR) partition is created on the data disk and an ext4 file system is mounted on the partition.

format Alibaba Cloud ECS Linux disk partition

### **Prerequisites**

Data disks that you have purchased separately are attached to ECS instances. If you purchased data disks with ECS instances, skip the attach operation. For more information, see Attach a data disk.

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

#### **Context**

The following procedure applies only to data disks up to 2 TiB in size. Data disks larger than 2 TiB must be partitioned into the GPT format. For more information, see Partition and format a data disk larger than 2 TiB.

By default, device names of data disks are assigned by the system. Device names of data disks for I/O optimized instances ascend from /dev/vdb to /dev/vdz. Device names of data disks for non-I/O optimized instances ascend from /dev/xvdb to /dev/xvdz.

The following section describes the risks of formatting a data disk:

- Disk partitioning and formatting are high-risk operations. Exercise caution. The procedure in this topic applies only to new data disks. If your data disk contains data, create a snapshot for the data disk to avoid data loss. For more information, see Create a normal snapshot.
- Only data disks can be partitioned. You cannot partition system disks. If you use a third-party tool to forcibly partition a system disk, unknown risks such as system failure and data loss may occur. You can extend partitions or add new partitions only for the system disk that has been extended. For more information, see Resize partitions and file systems of Linux system disks.



#### **Procedure**

- 1. Connect to an ECS instance. For more information, see Connect to an ECS instance.
- 2. Run the fdisk -l command to view all the data disks attached to the ECS instance.
  - **Note** If /dev/vdb is not displayed in the command output, the ECS instance does not have data disks. Check whether data disks are attached to the instance.
- 3. Run the following commands in sequence to create a single partition for the data disk:
  - i. Run the fdisk -u /dev/vdb command to partition the data disk.
  - ii. Enter 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 new partition.
  - iv. Enter p to set the partition as the primary partition.
    - Note When you create a single partition for a data disk, the partition must be a primary partition. If you want to create four or more partitions, enter e (extended) to create at least one extended partition.
  - v. Enter the partition number and press the Enter key. In this example, enter 1 to create one partition.
  - vi. Enter the number of the first available sector, or press the Enter key to use the default value 2048.
  - vii. Enter the number of the last sector. In this example, press the Enter key to use the default value.
  - viii. Enter p to view the intended partitions of the data disk.
    - ix. Enter w to start partitioning, and exit after partitioning is complete.

```
[root@ecshost~]# fdisk -u /dev/vdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0x3e60020e.
Command (m for help): p
Disk /dev/vdb: 21.5 GB, 21474836480 bytes, 41943040 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x3e60020e
Device Boot Start End Blocks Id System
Command (m for help): n
Partition type:
p primary (0 primary, 0 extended, 4 free)
e extended
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-41943039, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-41943039, default 41943039):
Using default value 41943039
Partition 1 of type Linux and of size 20 GiB is set
Command (m for help): p
Disk /dev/vdb: 21.5 GB, 21474836480 bytes, 41943040 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x3e60020e
Device Boot Start End Blocks Id System
/dev/vdb1 2048 41943039 20970496 83 Linux
Command (m for help): w
```

The partition table has been altered!

Calling ioctl() to re-read partition table.

Syncing disks.

4. Run the fdisk -lu /dev/vdb command to view the new partition. If information about /dev/vdb1 is displayed, the new partition is created.

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

Disk /dev/vdb: 21.5 GB, 21474836480 bytes, 41943040 sectors

Units = sectors of 1 \* 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk label type: dos

Disk identifier: 0x3e60020e

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

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

In this example, an ext4 file system is created. You can modify the mkfs.ext4 command to create other file systems. For example, if you want to share files between the Linux, Windows, and Mac operating systems, you can run the mkfs.vfat /dev/vdb1 command to create a Virtual File Allocation Table (VFAT) file system.

[root@ecshost~]# mkfs.ext4 /dev/vdb1

mke2fs 1.42.9 (28-Dec-2013)

Filesystem label=

OS type: Linux

Block size=4096 (log=2)

Fragment size=4096 (log=2)

Stride=0 blocks, Stripe width=0 blocks

1310720 inodes, 5242624 blocks

262131 blocks (5.00%) reserved for the super user

First data block=0

Maximum filesystem blocks=2153775104

160 block groups

32768 blocks per group, 32768 fragments per group

8192 inodes per group

Superblock backups stored on blocks:

32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,

4096000

Allocating group tables: done Writing inode tables: done

Creating journal (32768 blocks): done

Writing superblocks and filesystem accounting information: done

- 6. Write the information of the new partition to /etc/fstab to enable automatic partition mounting when the instance is started.
  - i. Run the cp /etc/fstab /etc/fstab.bak command to back up etc/fstab.

ii. Run the echo `blkid /dev/vdb1 | awk '{print \$2}' | sed 's/\"//g'` /mnt ext4 defaults 0 0 >> /etc/fst ab command to write the information of the new partition to the /etc/fstab file.

You can separately mount the data disk as a folder to store web pages. In this case, the mount portion in the command must be replaced with the intended mount point folder path.

# ? Note

- We recommend that you use a universally unique identifier (UUID) in /etc/fsta b to reference the new partition. You can run the blkid command to obtain the UUID of the new partition.
- 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 the /etc/fstab file.
- 7. Run the cat /etc/fstab command to view the information of the new partition in /etc/fstab.

```
[root@ecshost~]# cat /etc/fstab

# /etc/fstab

# Created by anaconda on Wed Dec 12 07:53:08 2018

# Accessible filesystems, by reference, are maintained under '/dev/disk'

# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info

# UUID=d67c3b17-255b-4687-be04-f29190d4**** / ext4 defaults 1 1

UUID=a4c73111-3a33-4569-a90c-b6d2e953**** /mnt ext4 defaults 0 0
```

8. Run the mount /dev/vdb1 /mnt command to mount the file system. If the information of the new file system is displayed after you run the df -h command, the file system is mounted.

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

Filesystem Size Used Avail Use% Mounted on
/dev/vda1 40G 1.6G 36G 5% /
devtmpfs 234M 0 234M 0% /dev
tmpfs 244M 0 244M 0% /dev/shm
tmpfs 244M 484K 244M 1% /run
tmpfs 244M 0 244M 0% /sys/fs/cgroup
tmpfs 49M 0 49M 0% /run/user/0
/dev/vdb1 20G 45M 19G 1% /mnt
```

# 2.3.2. Format a data disk for a Windows ECS

# instance

This topic describes how to create a partition of the Master Boot Record (MBR) format for a new data disk attached to a Windows ECS instance. Then, it describes how to mount an NTFS file system to the partition. You can create multiple partitions for a data disk based on your business needs.

format Alibaba Cloud ECS partition a disk in Windows

#### **Context**

The following risks may occur when you format a data disk:

- Partitioning and formatting data disks are high-risk operations. Exercise caution. The
  procedure in this topic applies only to new data disks. If your data disk contains data, create a
  snapshot for the data disk to avoid data loss. For more information, see Create a snapshot.
- You can partition only data disks that are attached to ECS instances. System disks cannot be
  partitioned. If you forcibly use a third-party tool to partition the system disk, unknown risks
  such as system failure and data loss may occur. You can only expand or add partitions for the
  system disk that has been extended. For more information, see Resize disks online for
  Windows instances.

#### **Procedure**

The following procedure applies only to data disks not greater than 2 TiB in size. For information about how to format a data disk greater than 2 TiB in size, see Partition and format a data disk larger than 2 TiB. In the following example, the ECS instance runs the Windows Server 2012 R2 64-bit operating system. A single MBR partition is created for a 20 GiB data disk that is attached to the ECS instance.

- 1. Connect to the instance. For more information, see Connect to an ECS instance.
- 2. On the Windows Server desktop, right-click the Start icon, and select Disk Management.
- 3. Find the unformatted data disk (Disk 2 in this example), which is in the Offline state.
- 4. Right-click the space around Disk 2, and select **Online**. After Disk 2 goes online, Disk 2 enters the **Not Initialized** state.
- 5. Right-click the space around Disk 2, and select Initialize Disk.
- 6. In the Initialize Disk dialog box, select Disk 2, and select one of the following partition formats:
  - Currently, MBR is most frequently used. This format only supports data disks up to 2 TiB in size with a maximum of four primary partitions. If you want to divide a disk into more than four partitions, you must use a primary partition as a partition to be expanded and create logical partitions in the primary partition.
  - GUID Partition Table (GPT) is a new partition format, which cannot be recognized by early Windows versions. The data disk size that GPT supports depends on the operating system and the file system. In Windows, GPT supports up to 128 primary partitions.
    - In this example, select MBR, and click OK.
- 7. In the Disk Management window, right-click the Unallocated space of Disk 2, and then select

New Simple Volume.

- 8. In the New Simple Volume wizard, perform the following operations:
  - i. Click Next.
  - ii. Specify the size of the simple volume. If you want to create only one primary partition, use the default value. Click **Next**.
  - iii. Select a drive letter. In this example, select F. Click Next.
  - iv. Specify the formatting settings, including file system, size of allocation unit, and volume label. Determine whether to enable Quick Formatting and File and Folder Compression. In this example, use the default settings. Click Next.
  - v. Start to create a simple volume. When the information shown in the following figure appears in the wizard, a new simple volume is created. Click Finish to close the New Simple Volume wizard.

#### Result

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



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

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

To partition and format data disks larger than 2 TiB, use the GPT partition format

This topic describes how to partition and format a data disk larger than 2 TiB on 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 Overview.

#### **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 Block Storage supports the master boot record (MBR) and GUID partition table (GPT) partition formats. MBR is applicable to data disks up to 2 TiB in size and can be used to create up to four primary partitions. If you want to partition a data disk larger than 2 TiB, 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 data on the disk, we recommend that you create and attach a new data disk to the instance. Use the GPT partition format to partition and format the new data disk, 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, we recommend that you use the partition tools, partition formats, 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 on a Windows instance

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

- 1. On the Windows Server desktop, press Win+R.
- 2. In the Run dialog box, enter *diskmgmt.msc* and click **OK**. The Disk Management window appears.
- 3. Find the disk to be partitioned and formatted. Disk 1 is used in this example. The disk is in the Offline state.
- 4. Right-click anywhere in the blank area next to Disk 1, and select Online. After Disk 1 comes online, it enters the Not Initialized state.
- 5. Right-click anywhere in 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 a primary partition, use the default value. Click Next. You can also divide Disk 1 into multiple partitions.
    - ? Note Theoretically, NTFS supports a maximum of 2 <sup>64</sup>-1 clusters in each volume. However, in Windows XP Pro, the number of clusters per NTFS volume is limited to 2 <sup>32</sup>-1. 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: Select the formatting settings including 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. Ensure 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 Window instance. The Windows Server 2012 R2 64-bit operating system is used in this example.

- 1. On the Windows Server desktop, right-click the Start icon and select Disk Management.
- 2. Find the disk to be converted. Disk 2 is used in this example.
- 3. Right-click a simple volume and select Delete Volume.
- 4. Right-click anywhere in the blank area next to Disk 2, and then 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 a primary partition, use the default value. Click Next. You can also divide Disk 2 into multiple partitions.
    - ? Note Theoretically, NTFS supports a maximum of 2 <sup>64</sup>-1 clusters in each volume. However, in Windows XP Pro, the number of clusters per NTFS volume is limited to 2 <sup>32</sup>-1. 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: Select the formatting settings including 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 on a Linux instance

This section describes how to use the parted and e2fsprogs tools to partition and format a data disk larger than 2 TiB on a Linux instance. The Alibaba Cloud Linux 2.1903 LTS 64-bit operating system is used in this example. Assume that the target data disk is a newly created 3 TiB empty disk and its device name is /dev/vdb.

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

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

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

1. Run the fdisk - command to check whether the data disk exists. The expected command output is similar to the following example. If different information is returned, the data disk is not attached to the instance.

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

- 2. Run the parted /dev/vdb command to start partitioning the disk.
  - i. Run the mklabel gpt command to convert the partition format from MBR to GPT.
  - ii. Run the mkpart primary 1 100% command to create a primary partition and specify the start and end sectors for the partition.
  - iii. Run the align-check optimal 1 command to check partition alignment.

Note If 1 not aligned is returned, the partition is not aligned. We recommend that you run the following commands and use the (<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 a new primary partition.

```
[root@ecshost~]# cat /sys/block/vdb/queue/optimal_io_size
[root@ecshost~]# cat /sys/block/vdb/queue/minimum_io_size
[root@ecshost~]# cat /sys/block/vdb/alignment_offset
[root@ecshost~]# cat /sys/block/vdb/queue/physical_block_size
```

iv. Run the print command to view the partition table.

(parted) mklabel gpt (parted) mkpart primary 1 100%

(parted) align-check optimal 1

1 aligned

(parted) print

Model: Virtio Block Device (virtblk)

Disk /dev/vdb: 3221GB

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

Partition Table: gpt

Disk Flags:

Number Start End Size File system Name Flags

1 17.4kB 3221GB 3221GB primary

- v. Run the quit command to exit the parted tool.
- 3. Run the partprobe command to enable the system to re-read the partition table.
- 4. Create a file system for the /dev/vdb1 partition. Run one of the following commands to create a file system.
  - o Create an ext4 file system.

[root@ecshost~]# mkfs -t ext4 /dev/vdb1

• Create an XFS file system.

[root@ecshost~]# mkfs -t xfs /dev/vdb1

# ? Note

- If the capacity of the data disk is 16 TiB, you must format the disk by using the correct version of e2fsprogs. For more information, see Appendix 1: Update e2fsprogs on a Linux instance.
- If you want to disable the lazy init feature of an ext4 file system to avoid impact on the I/O performance of data disks, see Appendix 2: Disable the lazy init feature on a Linux instance.
- 5. Run the mkdir /test command to create a mount point named /test.
- 6. Run the mount /dev/vdd1 /test command to mount the /dev/vdb1 partition to /test.
- 7. Run the df -h command to view the current disk space and usage.

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

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

Filesystem Size Used Avail Use% Mounted on
/dev/vda1 40G 6.4G 31G 18% /
devtmpfs 487M 0 487M 0% /dev

tmpfs 497M 0 497M 0% /dev/shm

tmpfs 497M 364K 496M 1% /run

tmpfs 497M 0 497M 0% /sys/fs/cgroup

tmpfs 100M 0 100M 0% /run/user/0
/dev/vdb1 2.9T 89M 2.8T 1% /test
```

- 8. (Recommended) Write new partition information to /etc/fstab so that partitions are automatically mounted each time the instance is started.
  - i. Run the cp /etc/fstab /etc/fstab.bak command to back up etc/fstab.
  - ii. Run the echo `blkid /dev/vdb1 | awk '{print \$2}' | sed 's/\"//g'` /test ext4 defaults 0 0 >> /etc/fs tab command to write new partition information to /etc/fstab.
    - ? Note 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 cat /etc/fstab command to check the information of /etc/fstab.

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

The procedure to partition and format a 3 TiB data disk is complete.

#### 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 /dev/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 rpm -qa | grep e2fsprogs command to check the current e2fsprogs version.

```
$sudo rpm -qa | grep e2fsprogs
e2fsprogs-libs-1.41.12-3
e2fsprogs-1.41.12-3
e2fsprogs-libs-1.39-33.1.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 the e2fsprogs page 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:

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

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

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

## Appendix 2: Disable the lazy init 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 initialize the file system:

```
mke2fs -0 64bit,has_journal,extents,huge_file,flex_bg,uninit_bg,dir_nlink,extra_isize -E lazy_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 initialize the file system. For example, the file system of a 32 TiB data disk may take 10 to 30 minutes to initialize. Enable or disable the lazy init feature to meet your business requirements.

# 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. Remotely connect to an ECS instance as a root user. For more information, see Overview.
- 2. Run the following command to view the name of the attached cloud disk.

fdisk -l

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

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

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

mkfs.ext4 /dev/vdb

Create an ext3 file system

mkfs.ext3 /dev/vdb

o Create an xfs file system

mkfs.xfs /dev/vdb

Create a btrfs file system

mkfs.btrfs /dev/vdb

4. (Optional)Create a mount directory, such as /media/vdb. If this step is omitted, you can also attach the disk to an existing directory.

mkdir /media/vdb

5. Attach the disk to the mount directory.

mount /dev/vdb /media/vdb

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

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

[root@ecshost ~]# df -h

Filesystem Size Used Avail Use% Mounted on

udev 3.9G 0 3.9G 0% /dev

tmpfs 798M 2.9M 795M 1% /run

/dev/vda1 40G 3.2G 35G 9% /

tmpfs 3.9G 0 3.9G 0% /dev/shm

tmpfs 5.0M 0 5.0G 0% /run/lock

tmpfs 3.9G 0 3.9G 0% /sys/fs/cgroup

tmpfs 798M 0 798M 0% /run/user/0

/dev/vdb 98G 61M 93G 1% /media/vdb

#### **Related information**

- Format a data disk for a Windows ECS instance
- Format a data disk for a Linux instance

# 2.4. Change the operating system

# 2.4.1. Replace the system disk (non-public images)

You can replace the operating system of an instance by replacing its system disk. To replace the system disk, you must assign a new system disk to the ECS instance. The original system disk is released and the system disk ID is updated. The disk category, the IP addresses of the instance, and MAC addresses of the elastic network interfaces (ENIs) that are bound to the instance remain unchanged. If you want to use a different operating system, you can replace the system disk to replace the operating system.

## **Prerequisites**

• After the system disk is replaced, the original system disk is released. We recommend that you create snapshots to back up disk data before you replace the system disk. For more information, see Create a snapshot.

Note Create snapshots during off-peak hours to minimize the impact on your business. It takes about 20 minutes to create the initial snapshot of a 40 GiB system disk. We recommend that you plan accordingly before you create a snapshot.

• The instance is in the **Stopped** state. For information about how to stop an instance, see **Stop** an instance.

#### Warning

- When you stop an instance, services that run on the instance may be interrupted. Proceed with caution.
- If the instance is a pay-as-you-go instance and is in the No Fees for Stopped Instances (VPC-Connected) state, it may not be able to start after you replace the system disk. We recommend that you select Retain Instance and Continue Charging After Instance Is Stopped when you stop the instance.

#### Context

Replacing a system disk is a high-risk operation. Before you perform this operation, read and understand the following precautions:

- You must redeploy the runtime environment on the new system disk. Your services may be stopped for an extended period of time.
- Snapshots of the original system disk cannot be used to roll back the new system disk.
- Manual snapshots are not released and can continue to be used to create custom images. If the Delete Automatic Snapshots While Releasing Disk feature is enabled on the original system disk, automatic snapshots are automatically deleted.
- For example, assume that the original and new operating systems are Linux systems, data disks are attached to the instance, and automatic partition mounting at system startup is enabled. After the system disk is replaced, the mounting configuration of the data disk partitions in the original system disk is lost. You need to manually update configurations in the /etc/fstab file. For more information, see What to do next.
- If you want to replace the system disk between different operating system families, take note of the following items:
  - For regions outside mainland China, you cannot replace a system disk with a disk of a different operating system family. The system supports only replacement among different Linux distributions or among different Windows Server versions.
  - Make sure that the hostname of the instance meets the requirements of the new operating system, and delete invalid characters. For more information about how to modify the hostname, see Modify the properties of an instance or ModifyInstanceAttribute.
  - When you replace a Windows Server operating system with a Linux operating system, you
    can select an SSH key pair for authentication. For this step, you must have an existing SSH
    key pair. For more information, see Create an SSH key pair.
  - Linux systems that use the default configuration do not recognize files in the NTFS format and Windows Server systems that use the default configuration do not recognize files in the ext3, ext4, or XFS format. You can select one of the following recommendations based on your data disks:
    - If your data disks do not contain important data, re-initialize the data disks and format them to file systems that are supported by the new operating system. For more information, see Re-initialize a data disk.
    - If your data disks contain important data, separately install software identifiers such as NTFS-3G when Windows Server is replaced with Linux, or Ext2Read and Ext2Fsd when Linux is replaced with Windows Server.
- If you want to replace the operating system to Windows Server, note that:
  - The system disk must have at least 1 GiB space reserved. Otherwise, the ECS instance may not be able to start after the system disk is replaced.
  - Microsoft no longer offers support for Windows Server 2003. To ensure your data security, we recommend that you do not use Windows Server 2003 for your ECS instances. Alibaba Cloud no longer provides Windows Server 2003 images. For more information, see Offline announcement of Windows Server 2003 system image.
  - If your instance is a non-I/O instance, you can only call the ReplaceSystemDisk operation to only replace to the specified Windows Server public images. For more information, see Replace the system disk (public images).

#### **Procedure**

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the target instance on the Instances page.
- 5. In the Actions column, choose More > Disk and Image > Replace System Disk.
- 6. In the message that appears, read the precautions for replacing the system disk, and then click **OK**.
- 7. On the Replace System Disk page, configure the following parameters:
  - i. Image Type: Select Custom Image, Shared Image, or Marketplace Image, and select an image version.
  - ii. System Disk: The category of the system disk cannot be changed. Configure the capacity of the system disk based on your business needs and the new image. The new capacity must be greater than or equal to the current capacity of the system disk. The minimum disk capacity is determined by the current disk capacity and the image type, while the maximum disk capacity is 500 GiB.

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

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

#### iii. Security enhancement:

- Windows Server instances only support password authentication.
- If the instance is an I/O optimized instance and the operating system is Linux, you can use a password or an SSH key pair for authentication.
- iv. Read and select ECS Service Terms.
- v. Price: The fees of the image and the system disk are included. For more information about prices of system disks, see Pricing.
- vi. Confirm settings and click Create Order.

#### Result

It takes about 10 minutes to replace the system disk. After the system disk is replaced, the instance automatically enters the Running state. Click the instance ID to go to the Instance Details page and view the new image.

#### What's next

After the system disk is replaced, you can perform the following operations:

- Optional. After the system disk is replaced, the automatic snapshot policy applied to the
  original 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.
- You must ensure that your snapshot quota is sufficient to configure an automatic snapshot policy for the new system disk. You can delete unneeded snapshots of the original system disk. For more information, see Delete a snapshot.
- Optional. For Linux instances only: Write the new partition information to the /etc/fstab file of the new system disk and mount the partitions. Data disks do not need to be formatted and partitioned. The following list provides the corresponding procedure. For more information about the specific commands, see Format a data disk for a Linux instance.
  - i. Recommended. Back up the /etc/fstab file.
  - ii. Write the new partition information to the /etc/fstab file.
  - iii. View the new partition information in the /etc/fstab file.
  - iv. Run the mount command to mount the partitions.
  - v. Run the df -h command to query the file system space and usage.

After the partitions are mounted, you can use the data disks without restarting the instance.

#### **Related information**

• ReplaceSystemDisk

# 2.4.2. Replace the system disk (public images)

Replace the image of an ECS instanceReplace the image of an ECS instance

You can replace the operating system of an instance by replacing its system disk. To replace the system disk, you must assign a new system disk to the ECS instance. The original system disk is released and the system disk ID is updated. The disk category, the IP addresses of the instance, and MAC addresses of the elastic network interfaces (ENIs) that are bound to the instance remain unchanged. If you want to use a different operating system, you can replace the system disk to replace the operating system.

## **Prerequisites**

- After the system disk is replaced, the original system disk is released. We recommend that you
  create snapshots to back up disk data before you replace the system disk. For more
  information, see Create a snapshot.
  - Note Create snapshots during off-peak hours to minimize the impact on your business. It takes about 20 minutes to create the initial snapshot of a 40 GiB system disk. We recommend that you plan accordingly before you create a snapshot.
- The instance is in the **Stopped** state. For information about how to stop an instance, see **Stop** an instance.

## Warning

- When you stop an instance, services that run on the instance may be interrupted. Proceed with caution.
- If the instance is a pay-as-you-go instance and is in the No Fees for Stopped Instances (VPC-Connected) state, it may not be able to start after you replace the system disk. We recommend that you select Retain Instance and Continue Charging After Instance Is Stopped when you stop the instance.

#### **Context**

Risks may arise when you replace a system risk. Before you perform this operation, read and understand the following precautions:

- You must redeploy the runtime environment on the new system disk. Your services may be stopped for an extended period of time.
- Snapshots of the original system disk cannot be used to roll back the new system disk.
- Manual snapshots are not released and can continue to be used to create custom images. If the Delete Automatic Snapshots While Releasing Disk feature is enabled on the original system disk, automatic snapshots are automatically deleted.
- For example, assume that the original and new operating systems are both Linux systems, data disks are attached to the instance, and automatic partition mounting at system startup is enabled. After the system disk is replaced, the mounting configuration of the data disk partitions in the original system disk is lost. You must manually update configurations in the /e tc/fstab file. For more information, see the What to do next section.
- If you want to replace the system disk between different operating system families, take note of the following items:
  - For regions outside mainland China, you cannot replace a system disk with a disk of a different operating system family. The system supports only replacement among different Linux distributions or among different Windows Server versions.
  - Make sure that the hostname of the instance meets the requirements of the new operating system, and delete invalid characters. For more information about how to modify the hostname, see Modify the properties of an instance or ModifyInstanceAttribute.
  - When you replace a Windows Server operating system with a Linux operating system, you
    can select an SSH key pair for authentication. For this step, you must have an existing SSH
    key pair. For more information, see Create an SSH key pair.
  - Linux systems that use the default configuration do not recognize files in the NTFS format and Windows Server systems that use the default configuration do not recognize files in the ext3, ext4, or XFS format. You can select one of the following recommendations based on your data disks:
    - If your data disks do not contain important data, re-initialize the data disks and format them to file systems that are supported by the new operating system. For more information, see Re-initialize a data disk.
    - If your data disks contain important data, separately install software identifiers such as NTFS-3G when Windows Server is replaced with Linux, or Ext2Read and Ext2Fsd when Linux is replaced with Windows Server.
- If you want to change the operating system to Windows Server, take note of the following items:

- The system disk must have at least 1 GiB space reserved. Otherwise, the ECS instance may not be able to start after the system disk is replaced.
- Microsoft no longer offers support for Windows Server 2003. To ensure your data security, we recommend that you do not use Windows Server 2003 for your ECS instances. Alibaba Cloud no longer provides Windows Server 2003 images. For more information, see Offline announcement of Windows Server 2003 system image.

#### **Procedure**

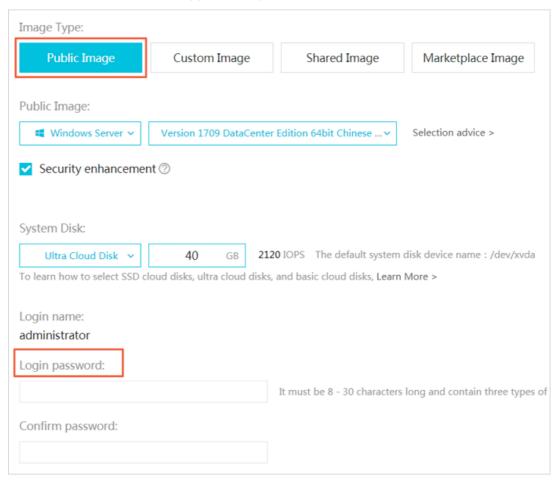
- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. On the Instances page, find the instance whose system disk you want to replace.
- 5. In the Actions column, choose More > Disk and Image > Replace System Disk.
- 6. In the message that appears, read the precautions for replacing the system disk, and then click **OK**.
- 7. On the Change Operating System page, configure the following parameters:
  - i. Image: Select Public Image and select an image version. For more information about how to replace an image of the instance with a non-public image, see Replace the system disk (non-public images).
  - ii. System Disk: The category of the system disk cannot be changed. Configure the capacity of the system disk based on your business needs and the new image. The new capacity must be greater than or equal to the current capacity of the system disk. The minimum disk capacity is determined by the current disk capacity and the image type, while the maximum disk capacity is 500 GiB.

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

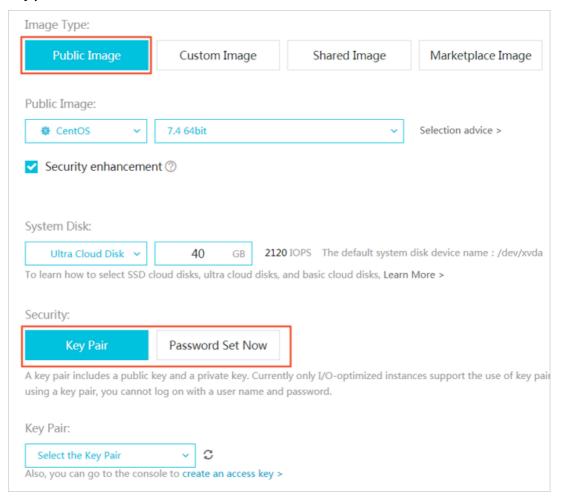
Note If you have renewed and downgraded your instance, the system disk capacity cannot be modified until the next subscription period starts.

iii. Security Settings: Specify the logon authentication method.

• Windows Server instances support only password authentication.



• If the instance is an I/O optimized Linux instance, you can use a password or an SSH key pair for authentication.



- iv. Read and select ECS Service Terms.
- v. Price: The fees of the image and the system disk are included. For more information about prices of system disks, see Pricing.
- vi. Confirm settings and click Create Order.

#### Result

It takes about 10 minutes to replace the system disk. After the system disk is replaced, the instance automatically enters the Running state. Click the instance ID to go to the Instance Details page and view the new image.

#### What's next

After the system disk is replaced, you can perform the following operations:

- Optional. After the system disk is replaced, the automatic snapshot policy applied to the
  original 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.
- You must ensure that your snapshot quota is sufficient to configure an automatic snapshot policy for the new system disk. You can delete unneeded snapshots of the original system disk. For more information, see Delete a snapshot.

- Optional. For Linux instances only: Write the new partition information to the /etc/fstab file of the new system disk and mount the partitions. Data disks do not need to be formatted and partitioned. The following list provides the corresponding procedure. For more information about the specific commands, see Format a data disk for a Linux instance.
  - i. Recommended. Back up the /etc/fstab file.
  - ii. Write the new partition information to the /etc/fstab file.
  - iii. View the new partition information in the /etc/fstab file.
  - iv. Run the mount command to mount the partitions.
  - v. Run the df -h command to query the file system space and usage.

After the partitions are mounted, you can use the data disks without restarting the instance.

#### **Related information**

ReplaceSystemDisk

# 2.5. Reinitialize a cloud disk

# 2.5.1. Re-initialize a system disk

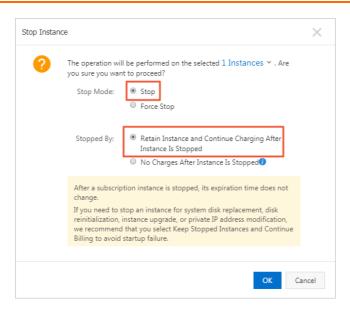
This topic describes how to re-initialize a system disk in the ECS console. After the system disk is re-initialized, it is restored to the status when it was created.

## **Prerequisites**

- Data is lost after the disk is re-initialized. You must create snapshots for the disk to back up data. For more information, see Create a normal snapshot.
- The instance must be in the Stopped state before you re-initialize the disk. Your business is interrupted when you stop your instance. Proceed with caution. For more information, see Stop an instance.

You must select **Retain Instance and Continue Charging After Instance Is Stopped** when you stop a pay-as-you-go instance.

Notice If you select No Charges After Instance Is Stopped, you may not be able to start the instance after you re-initialize the disk. For more information, see No Fees for Stopped Instances (VPC-Connected).



- You cannot re-initialize a disk that has local snapshots. You can create local snapshots only
  for enhanced SSDs (ESSDs). Therefore, if the disk is an ESSD, make sure that it does not have
  local snapshots.
- (Applicable to Linux instances only) If you want to use an SSH key pair for authentication
  when you re-initialize a system disk, you must have created or imported an SSH key pair. For
  more information, see Create an SSH key pair and Import an SSH key pair.

#### **Context**

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

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

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

**Note** If you re-initialize the system disk after you replace the system disk, the new system disk is re-initialized instead of the old one.

- If an automatic snapshot policy is applied to the system disk, the policy remains valid after the system disk is re-initialized.
- The IP address and system disk ID of the ECS instance are not changed.
- 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. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the ECS instance whose system disk you want to re-initialize, and click the instance ID

to go to the Instance Details page.

- 5. In the left-side navigation pane, click Disks.
- 6. Find the system disk, and click Reinitialize Disk in the Actions column.
- 7. In the Reinitialize Disk dialog box, configure the following parameters.

| Parameter                  | Description  |
|----------------------------|--|
|                            | Windows instances: Specify a logon password. You can use the old password or specify a new password.   |
|                            | Linux instances: Select Set SSH Key or Set Password.   |
| Security                   | <ul> <li>Set SSH Key: binds an SSH key pair to the instance. You can then log<br/>on the instance by using the SSH key pair.</li> </ul>  |
|                            | <ul> <li>Set Password: specifies a logon password. You can use the old<br/>password or specify a new password.</li> </ul>  |
| Security<br>Enhancement    | After you select Activate, security components are loaded to the instance free of charge. These components provide services such as webshell detection, remote logon reminder, and protection against brute-force attacks. |
| Instance Startup<br>Policy | If you select <b>Start Instance after Resetting Disk</b> , the instance automatically starts up after you re-initialize the system disk.   |

#### 8. Click Confirm.

#### What's next

• (Applicable to Linux instances only) If data disks have been attached to the ECS instance before you re-initialize its system disk, you must recreate the mount point information and mount file systems to the instance. For more information, see How do I reattach data disks after I re-initialize the system disk of a Linux instance? in *Elastic Block Storage FAQ*.

Note After you re-initialize the system disk of a Linux instance, data within the data disks that are attached to the instance is not changed, but the mount point information is lost.

- After the system disk is re-initialized, you must redeploy and configure applications to restore your business as soon as possible.
- If you have created snapshots for the system disk, you can use the snapshots to create data disks and attach the data disks to the instance to obtain data in the original 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

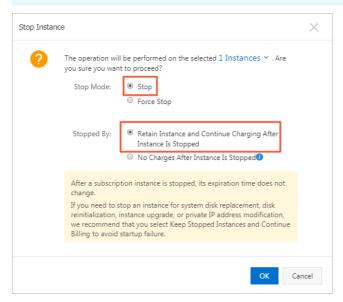
When a data disk is attached to an ECS instance, you can re-initialize the disk to restore it to the status when it was created.

# **Prerequisites**

- Data is lost after the disk is re-initialized. You must create snapshots for the disk to back up data. For more information, see Create a normal snapshot.
- 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 must be in the Stopped state before you re-initialize the disk. Your business is interrupted when you stop your instance. Proceed with caution. For more information, see Stop an instance.

You must select **Retain Instance and Continue Charging After Instance Is Stopped** when you stop a pay-as-you-go instance.

Notice If you select No Charges After Instance Is Stopped, you may not be able to start the instance after you re-initialize the disk. For more information, see No Fees for Stopped Instances (VPC-Connected).



- You cannot re-initialize a disk that has local snapshots. You can create local snapshots only for enhanced SSDs (ESSDs). Therefore, if the disk is an ESSD, make sure that it does not have local snapshots.
- (Applicable to Linux instances only) If you create an empty data disk and add a command in the /etc/fstab file to mount partitions of the data disk at the 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 operations:
  - 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 The /dev/vdb1 partition and the /InitTest mount point are used in this example. You can modify the command based on your actual conditions.

v. Press the Esc key to exit the edit mode. Then, 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.
  - The data disk stores the data recorded in the source snapshot if it was created from a snapshot.
- For a Windows instance, after you re-initialize its data disk, the data disk is ready for use without additional operations regardless of its original status.
- For a Linux instance:
  - If the 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 is created is lost.
  - If the 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. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the ECS instance for which you want to re-initialize a data disk, and click the instance ID to go to the Instance Details page.
- 5. In the left-side navigation pane, click Disks.
- 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 to Linux instances only) 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 Format a data disk for a Linux instance.
- After the data disk is re-initialized, you must redeploy and configure applications to restore your business as soon as possible.

#### **Related information**

• ReInitDisk

• 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.

Warning The rollback operation is irreversible. After a disk is rolled back, the data from the creation of the snapshot to when the disk is rolled back is lost. To avoid data losses caused by incorrect operations, we recommend that you create a snapshot for the current disk before you roll back the disk.

- The disk has not been released.
- After you replace the system disk, 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.

Note A pay-as-you-go instance in a VPC may not be restarted after its disk is rolled back if the No Fees for Stopped Instances (VPC-Connected) feature is enabled. We recommend that you disable the No Fees for Stopped Instances (VPC-Connected) feature before you stop the instance. For more information, see No Fees for Stopped Instances (VPC-Connected).

## Roll back a disk by using a snapshot

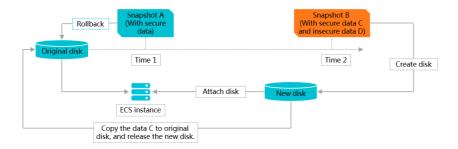
You can roll back a disk from the **Snapshots** or **Instances** page. The following example shows how to roll back a disk from the **Instances** page.

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the instance whose disk you want to roll back the disk and click Manage in the Actions column.
- 5. In the left-side navigation pane, click Snapshots.
- 6. Find the target snapshot and click Rollback 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 not affected. After the disk is rolled back, the disk is restored to the status of the entire disk at a certain point in time, instead of the status of a partition or a directory.
- 7. In the dialog box that appears, click OK.
  - Note If you select Start Instance Immediately after 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 need to synchronize the data 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.
- 3. Create a new 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. Remotely connect to an ECS instance. For more information, see Overview.
- 6. View the new disk.
  - Windows instances: The new disk is displayed in the system.
  - Linux instances: Run the mount command to mount the new disk.
- 7. Copy useful data from the new disk to the original disk.
- 8. (Optional)Release the new disk as needed.

#### **Related information**

ResetDisk

# 2.7. Modify the performance levels of ESSDs

You can upgrade the performance level of an ESSD in the ECS console. If you want to downgrade the performance level of a pay-as-you-go ESSD, you must call the ModifyDiskSpec operation. The new performance level takes effect immediately. You do not need to create a new ESSD or copy the data.

#### **Context**

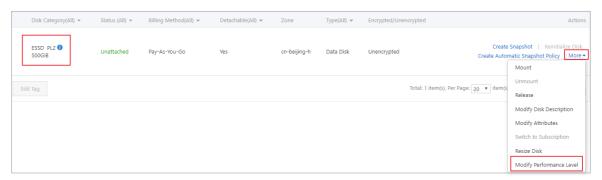
When you create an ECS instance, you can set an ESSD as the system disk or data disk. You can also create a separate ESSD. For more information about ESSDs, see Enhanced SSDs.

To modify the performance level of an ESSD, make sure that:

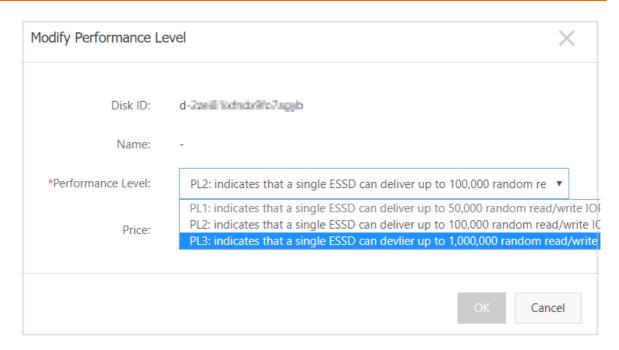
- Your account does not have overdue payments.
- If your ESSD is attached to a pay-as-you-go ECS instance, make sure that the instance is not in the expired state.
- You can modify the performance level of a new ESSD only after it enters the **Unattached** (*Avail able*) state.
- After the performance level is modified, the ESSD is billed based on the new performance level.

#### **Procedure**

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



5. In the Modify Performance Level dialog box that appears, select a new Performance Level. Click OK.



The performance level that you can select for an ESSD is determined by its storage capacity. If the performance level of an ESSD cannot be upgraded, resize the ESSD and then upgrade its performance level. For more information, see Resize cloud disks.

#### **Related information**

ModifyDiskSpec

# 2.8. View the monitoring data of a cloud disk

This topic describes how to monitor the IOPS and throughput of a cloud disk within a specified time range in the ECS console.

read/write monitoring IOPS throughput Alibaba Cloud ECS

#### **Context**

For information about the performance of different cloud disks, see EBS performance. The IOPS and throughput metrics are described as follows:

- IOPS: input/output operations per second, is the number of write or read operations that the block storage can perform per second. You must view IOPS data if transaction-intensive applications are employed on the disk.
- Throughput: measures the amount of data transferred per second. Unit: MB/s. You must view throughput data if applications requiring numerous sequential read/write operations are employed on the disk.

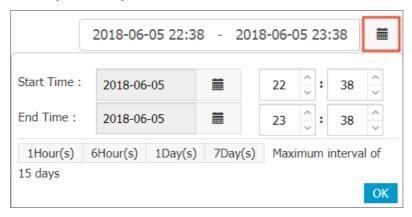
If you have installed CloudMonitor for the ECS instance to which the disk is attached, you can view the disk usage, IOPS, and throughput in CloudMonitor. For more information, see Host monitoring.

#### **Procedure**

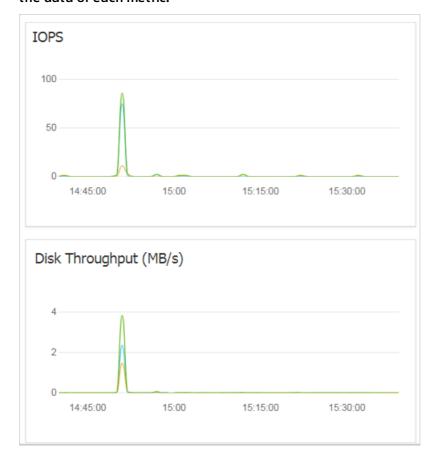
- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
- 3. In the top navigation bar, select a region.
- 4. Find the target cloud disk and click its ID to go to the Disk Details page.
- 5. In the left-side navigation pane, click Disk Monitoring.
- 6. On the Monitoring Information page, click the



icon to set the start and end time. You can view the monitoring information of the cloud disk for the past 15 days.



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



#### **Related information**

DescribeDiskMonitorData

# 2.9. 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, helping identify and distinguish different disks. This topic describes how to query the serial number of a disk in Linux and Windows.

#### Limits

Limits on using serial numbers:

- You can guery only the serial numbers of new disks.
- This feature can only be used for I/O optimized instances.
- This feature is supported by enhanced SSDs (ESSDs), standard SSDs, ultra disks, and basic disks.
- This feature is supported by local NVMe SSDs and SATA HDDs.

## Query the serial number of a disk in Linux

- 1. Remotely connect to an ECS instance. For more information, see Overview.
- 2. Run the lsblk command to query the names of Block Storage devices in 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 query the serial number of a Block Storage 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 ~]# udevadminfo --query=all --name=/dev/vda | grep ID_SERIAL

E: ID_SERIAL=bp1d4foh3ef8bntl****
```

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

- 1. Remotely connect to an ECS instance. For more information, see Overview.
- 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 specified logical volume (LV):

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

Example:

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 disk serial number:

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. Remotely connect to an ECS instance. For more information, see Overview.
- 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 LV:

Get-WmiObject -Class Win32\_LogicalDiskToPartition |select Antecedent, Dependent |fl

#### Example:

PS C:\Users\Administrator> Get-WmiObject -Class Win32\_LogicalDiskToPartition |select Antecede nt, 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 disk serial number:

Get-WmiObject -Class Win32\_PhysicalMedia | select Tag, Serialnumber

#### Example:

PS C:\Users\Administrator> Get-WmiObject -Class Win32\_PhysicalMedia |select Tag, Serialnumber

Tag Serialnumber

--- -----------\\.\PHYSICALDRIVEO bp1bet4g35opq6vq\*\*\*\*

# 2.10. Change billing methods of disks

After you create a disk, you can change its billing method between subscription and pay-as-you-go. This topic describes the methods to change billing methods, as well as the application 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 is a subscription disk that is created separately on the Disks page, you must pay for the disk before you can use it. For more information, see Subscription.
  - ? Note Subscription disks can be created separately 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 is a pay-as-you-go disk that is created separately 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 lists different operations that you can perform to change billing methods of disks.

| Direction of change              | Operation  | Application scope   | Time to take effect            |
|----------------------------------|--|---|--------------------------------|
| Subscription to<br>pay-as-you-go | Change the billing method of a subscription disk | This method can be used for subscription data disks that are attached to subscription instances. This method cannot be used to change the billing method of system disks.                     | From the next<br>billing cycle |
|                                  | Subscription to pay-as-you-go                    | This method can be used for system disks and data disks that are attached to subscription instances. The billing methods of the disks vary depending on the billing methods of the instances. | Immediately                    |
|                                  |  |   |                                |

| Direction of change                                | Operation                        | Application scope  | Time to take<br>effect |
|--|----------------------------------|--|------------------------|
| a subscription disk  Pay-as-you-go to subscription | billing method of a subscription | This method can be used for pay-as-<br>you-go data disks that are attached<br>to subscription instances. This<br>method cannot be used to change the<br>billing method of system disks.        | Immediately            |
|  | Pay-as-you-go to subscription    | This method can be used for system disks and data disks that are attached to pay-as-you-go instances. The billing methods of the disks vary depending on the billing methods of the instances. | Immediately            |

#### References

ModifyDiskChargeType

# 2.11. Detach or attach a system disk

You can detach a system disk from an ECS instance. If an instance fails to start because its system disk contains corrupted files, you can detach it 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.

Before you attach a 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 same 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
- Release the system disk separately
- Replace the system disk

- Resize the system disk
- Change the billing method of the disk

## Detach a system disk

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the target instance in the **Stopped** state. Click the instance ID to go to the **Instance Details** page.
- 5. In the left-side navigation pane, click Disks.
- 6. Select System Disk from the Type drop-down list. Choose More > Detach in the Actions column.
- 7. In the dialog box that appears, click OK.
  - Note If you clear Release Disk with Instance in the Detach Disk dialog box, the system disk will be retained as a data disk when its attached instance is released. The disk is then charged on a pay-as-you-go basis.
- 8. Check whether the system disk is detached.
  - i. On the Instances page, the instance is in the No System Disk state and is bound with a tag.
  - ii. In the left-side navigation pane, choose Storage & Snapshots > Disks.
  - iii. The system disk is converted to a data disk, the disk is bound with three tags, and the billing method remains unchanged.

The tags bound to the instance and the disk can help you search for resources. The following table describes the tags.

| Tag key                          | Description   | Tag value   |
|----------------------------------|---|---|
| acs:ecs:sourceSyste<br>mDiskId   | The system disk of the instance.  | The ID of the disk.   |
| acs:ecs:sourceInstan<br>ceId     | The instance to which the system disk is attached.  | The ID of the instance.   |
| acs:ecs:diskDeletePr<br>otection | The release protection feature is enabled for the disk. The disk cannot be manually released. | true  |
|                                  |   | The billing method of the disk before it is detached.                                 |
| acs:ecs:diskPayType              | The billing method of the disk  | <ul> <li>PrePaid:         subscription.</li> <li>AfterPay: pay-as- you-go.</li> </ul> |

## Attach a system disk

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the target instance in the No System Disk state. Click Attach Disk in the Actions column.
- 5. In the Attach Disk dialog box that appears, perform the following operations:
  - i. Copy the value of Source System Disk to the Target Disk field.
  - ii. Configure the parameters described in the following table for the system disk.

| Parameter   | Description   | Example                   |
|---|---|---------------------------|
| Login Credentials                                     | <ul> <li>When you attach the system disk, you must configure the Login Credentials parameter for the instance.</li> <li>Key Pair: Select this option for Linux instances only.</li> <li>Custom Password: You can select this option for both Windows Server and Linux instances. However, this option takes effects only when the default administrator and root usernames are used.</li> </ul> | Key Pair                  |
| Release Disk with<br>Instance                         | Specifies whether to retain the disk when the instance is released.  If this option is selected, the disk is also released when the instance is released.  If this option is not selected, the disk is retained as a data disk and billed on a pay-as-you-go basis.   | Do not select this option |
| Delete Automatic<br>Snapshots While<br>Releasing Disk | Specifies whether to retain the automatic snapshots of a disk when the disk is released.  If this option is selected, the automatic snapshots are released when the disk is released.  If this option is not selected, the automatic snapshots are retained when the disk is released.  | Select this option        |

- 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**

- DetachDisk
- AttachDisk

# 2.12. Detach a data disk

You can detach a data disk from an ECS instance if the billing method of the data disk is pay-as-you-go. You cannot detach a data disk which is used as a local disk.

detach a data disk Flexible attachment SCU ECS Alibaba Cloud

## **Prerequisites**

Ensure that you have performed the following operations before detaching a disk:

- The cloud disk has been attached to an instance and its Status is In Use.
- To ensure data integrity, we recommend that you stop reading or writing file systems of the data disk to avoid data loss before you detach the data disk.
- If file systems have been attached to the data disk, you must connect to the ECS instance and detach the data disk or partitions from the operating system, as shown in the following table.

| Operating system | Detach a data disk or partition   |  |
|------------------|---|--|
| Linux If y       | Run the umount [Data disk partition] command such as umount /dev/vdb1 .   |  |
|                  | If you have written the partition attachment information to the /etc/fstabfile, delete the partition attachment information. Otherwise, you cannot connect to the instance after the instance is restarted. |  |
| Windows Server   | In the Disk Management dialog box, right-click the disk name such as Disk 2 and select Offline.   |  |

#### **Context**

You can detach a data disk in the ECS console in the following ways:

- Detach a cloud disk on the Instance Details page. if you want to detach one or more disks from an instance.
- Detach a data disk on the Disks page if you know which disk you want to detach.

# Detach a cloud disk on the Instance Details page.

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
- 3. In the top navigation bar, select a region.
- 4. Find the target instance and click its ID to go to the Instance Details page.
- 5. In the left-side navigation pane, click Disks.
- 6. Find the target disk and choose **More** > **Release** in the **Actions** column. Only data disks that have the following attributes can be detached:

- Status must be In Use.
- Detachable must be Yes.
- Disk Category must be Data Disk.
- 7. In the message that appears, click OK.

Note You can repeat the preceding procedure to detach multiple data disks.

When the Status becomes Unattached, it indicates that the data disk has been detached.

#### Detach a data disk on the Disks page

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
- 3. In the top navigation bar, select a region.
- 4. Find the target disk and choose **More** > **Release** in the **Actions** column. Only data disks that have the following attributes can be detached:
  - o Status must be In Use.
  - Detachable must be Yes.
  - o Disk Category must be Data Disk.
- 5. In the message that appears, click OK.

When the Status becomes Unattached, it indicates that the data disk has been detached.

#### What's next

You can release cloud disks that are no longer needed. For more information, see Release a disk.

#### **Related information**

DetachDisk

## 2.13. 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 ECS instances are released together with the instances. You can disable the Release Disk with Instance feature to retain disks. This topic describes how to release a disk from the ECS console and how to enable or disable the Release with Instance feature.

release a disk delete a disk deletedisk Alibaba Cloud ECS

#### **Prerequisites**

- A snapshot of the disk is created to back up important data if the important data is stored on the disk. For more information, see Create a normal snapshot.
- The disk to be manually released and the pay-as-you-go data disk for which you want to enable or disable the Release Disk with Instance feature are 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 this disk are released along with the disk. You can disable the Delete Automatic Snapshots While Releasing Disk feature for a disk when you attach the disk. For more information, see Delete automatic snapshots while releasing a disk.

- Manual snapshots are not released together with a disk.
- You can manually release a disk or enable the Release Disk with Instance feature for the disk. The following table describes the releasing methods that different types of disks support.

| Disk type   | Subscription   | Pay-as-you-go   |
|-------------|--|---|
| System disk | By default, the Release Disk with Instance feature is enabled for subscription system disks. You can disable the feature.  | By default, the Release Disk with Instance feature is enabled for payas-you-go system disks. You can disable the feature.   |
| Data disk   | <ul> <li>By default, the Release Disk with<br/>Instance feature is enabled for<br/>subscription data disks. You can<br/>disable the feature.</li> <li>After subscription data disks are<br/>converted to pay-as-you-go data<br/>disks, they can be manually<br/>released.</li> </ul> | <ul> <li>Pay-as-you-go data disks can be released together with their attached instances. By default, the Release Disk with Instance feature is enabled for all pay-as-you-go data disks (except for pay-as-you-go data disks that are separately created). For separately created pay-as-you-go data disks, you must manually enable the feature.</li> <li>Pay-as-you-go data disks can be manually released.</li> </ul> |

- Consider the following items when you enable or disable the Release Disk with Instance feature for a system disk or a data disk:
  - With the Release Disk with Instance feature enabled, a disk is automatically released at the release of its attached ECS instance.
  - With the Release Disk with Instance feature disabled, a disk is retained as a pay-as-you-go data disk three days after the attached ECS instance expires, three days after a payment becomes overdue for the instance, or when the instance is manually released. Your account must complete real-name verification to ensure that disks can be retained.

Note The retained data disk is billed on a pay-as-you-go basis. You can log on to the Billing Management console and view the consumption details based on the disk ID.

#### Manually release a disk

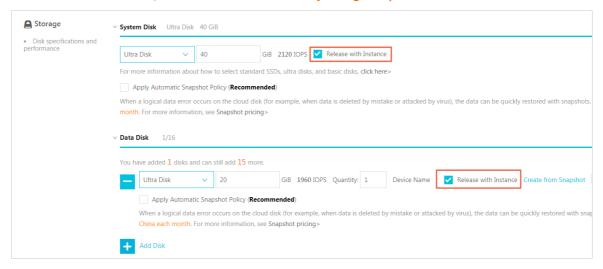
Perform the following steps to manually release a pay-as-you-go data disk:

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
- 3. In the top navigation bar, select a region.
- 4. Find the disk that you want to release. Choose More > Release in the Actions column.
- 5. In the message that appears, confirm the information and click Confirm Release.

# Enable or disable the Release Disk with Instance feature for disks while you create an ECS instance

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

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. Click Create Instance.
- 4. In the Basic Configurations step, scroll down to the Storage section and select or clear Release with Instance for data disks or the system disk. For more information about how to create an ECS instance, see Create an instance by using the provided wizard.

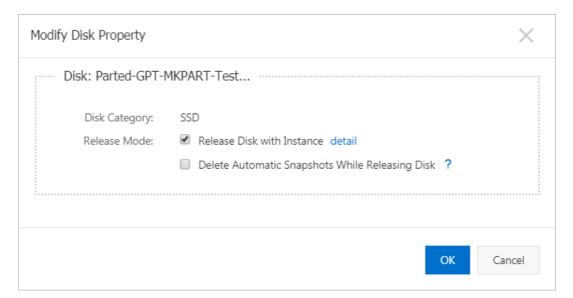


#### Enable or disable the Release Disk with Instance feature for a disk on the Disks page

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

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Storage & Snapshots > Disks.
- 3. In the top navigation bar, select a region.
- 4. Find the disk for which you want to enable or disable the feature. Choose More > Modify Disk Property in the Actions column.
- 5. In the Modify Disk Property dialog box that appears, select or clear Release Disk with Instance and click OK.

Note After you clear Release Disk with Instance, a disk is retained as a pay-as-you-go data disk three days after the attached ECS instance expires, three days after a payment becomes overdue for the instance, or when the instance is manually released.



### **Related information**

DeleteDisk

# 3.Encrypt disks

# 3.1. Encryption overview

This topic introduces the basic concepts of encryption. Encryption can help you secure your data stored in Alibaba Cloud ECS and comply with security standards. You can protect the privacy, autonomy, and security of data and ensure the security of your business data without the need to build or maintain key management infrastructure. Both system disks and data disks can be encrypted.

encryption data protection ECS Alibaba Cloud security compliance

#### **Feature description**

Alibaba Cloud ECS uses the industry-standard AES-256 algorithm to encrypt disks. Keys contain service keys and customer-created keys. You can obtain customer-created keys by using Bring Your Own Key (BYOK). Encryption and decryption have minimal impacts on ECS instance performance.

- Data in the operating system of the ECS instance is automatically encrypted when you create
  an ECS instance from an encrypted system disk or image. The data is automatically decrypted
  when it is read. For more information about the procedure, see Encrypt a system disk.
- The following types of data are automatically encrypted when you create an encrypted data disk and attach it to an ECS instance. The data is automatically decrypted when it is read. For more information about the procedure, see Encrypt a data disk.
  - o Static data stored on an encrypted disk.
  - Data transmitted between an encrypted disk and an instance. Data in the operating system is not encrypted again.
  - Data transmitted from an instance to a backend storage cluster.
  - All snapshots created from an encrypted disk. These snapshots have the same encryption key as the disk.
  - All disks created from encrypted snapshots.

#### Keys

A customer master key (CMK) contains metadata such as the key ID, creation date, description, rotation plan, lifecycle status, and the key material that is used to encrypt and decrypt data. If you use a CMK to encrypt a disk, all subsequent data written to the disk is also encrypted with the same CMK. The CMK is also used to encrypt all snapshots and the disks that are created from these snapshots. The CMK is used only in the memory of the host where your ECS instance is running. The CMK is not stored on any storage media in plaintext. CMKs are stored in the key management infrastructure provided by Key Management Service (KMS). CMKs provide strong physical and logical security protection against unauthorized access. The key management infrastructure of Alibaba Cloud conforms to the recommendations in NIST SP 800-57 and uses cryptographic algorithms that comply with FIPS Publication 140-2.

The following table describes the types of CMKs that you can use to encrypt disks.

| Туре   | Description  | Source  | Scenario  |
|--|--|---|---|
| Service key,<br>the first key<br>in the<br>following<br>figure               | When you activate KMS and use encryption for the first time within a region, KMS automatically creates a dedicated CMK for the ECS service within the region. The alias of the CMK is acs/ecs. Service keys cannot be deleted or disabled. | The default service CMK provided by KMS.  | You can use service keys when you want to enhance efficiency and convenience. For more information, see What is KMS?. |
| Customer-<br>created key,<br>the second<br>key in the<br>following<br>figure | The encryption keys that you create. You have full management permissions to create, rotate, and disable these keys, and define access control over them.  | <ul> <li>Source 1: The key that is created in KMS.</li> <li>Source 2: The key that is created in KMS from key materials imported through BYOK.</li> </ul> | You can use customer-created keys to improve operation flexibility and increase the number of keys.                   |

| Alias Name   | ID       | Protection Level | Created At | Next Rotation Date | Status                    | Actions                 |
|--------------|----------|------------------|------------|--------------------|---------------------------|-------------------------|
| EcsDocTest 2 | 2a513029 | SOFTWARE         | Feb 19, 20 | None               | <ul><li>Enabled</li></ul> | Enable   Disable   More |
| acs/ecs 1    | 315faac0 | SOFTWARE         | Mar 25, 20 | None               | Enabled                   | Enable   Disable   More |

#### **Billing**

The following table describes the billing information of encryption features and operations. Ensure that your account balance is sufficient. Otherwise, operations that incur costs may fail.

| Feature or operation  | Billing |
|---|---------|
| Encrypt system disks and data disks   | No      |
| Use service keys that are provided by KMS   | No      |
| CMKs that are created by you in KMS or created from key materials imported through BYOK.  | Yes     |
| Read and write operations on disks, such as mounting (mount) and unmounting (umount) of disks, creation of partitions, and formatting of file systems | No      |

#### Limits

- You can encrypt data disks such as enhanced SSDs, standard SSDs, ultra disks, and basic disks.
- You can encrypt system disks such as enhanced SSDs, standard SSDs, and ultra disks.
- You cannot encrypt local disks.
- You can encrypt system disks only when you are copying their custom images. You cannot

perform the following operations on encrypted system disks:

- o Convert encrypted images to unencrypted images.
- Copy encrypted images across regions.
- Share encrypted images.
- Export encrypted images.
- You cannot directly convert unencrypted disks to encrypted disks.
- You cannot directly convert encrypted disks to unencrypted disks.

# 3.2. Encrypt a system disk

When you copy an image, you can choose to encrypt the custom image. The system disk and data disks (if any) that are created from the custom image are automatically encrypted. The encryption key of the system disk and data disks is the same as that of the custom image. You can use the customer master key (CMK) that is automatically created in Key Management Service (KMS) or a custom key that you import (BYOK) as the encryption key.

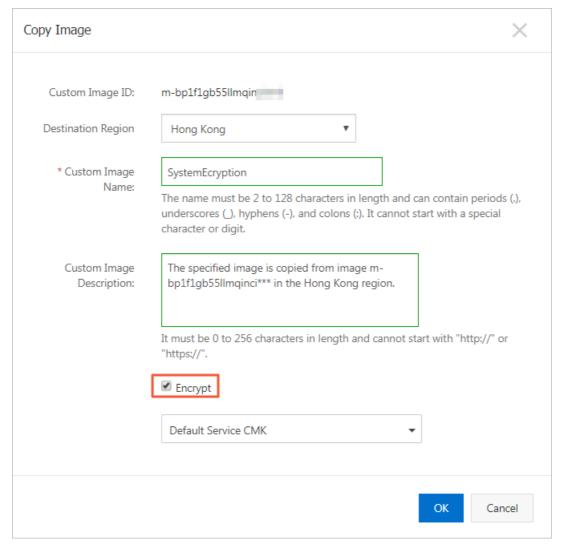
Alibaba Cloud data encryption security compliance ECS

#### **Context**

You can encrypt a system disk only by copying a custom image. For more information, see **Encryption overview**.

#### Encrypt a system disk when copying an image in the ECS console

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Images.
- 3. In the top navigation bar, select a region.
- 4. On the Images page, click the Custom Image tab.
- 5. Select the target image and click Copy Image in the Actions column.
  - ? Note If the size of your custom image is greater than 500 GiB, you are prompted to submit a ticket to complete the operation when you click Copy Image.
- 6. In the Copy Image dialog box, select Encrypt and then select a key from the drop-down list. When Encrypt is selected, Alibaba Cloud uses the managed service key Default Service CMK by default. You can also specify a CMK that you created in KMS as the encryption key. We recommend that you use a custom key as the encryption key.
  - Note If this is the first time that you select Encrypt, click Go to Authorize and select AliyunECSDiskEncryptDefaultRole to allow ECS to access your KMS resources. This procedure describes only how to configure the encryption setting when you copy a custom image. For more information about other configurations, see Copy custom images.



#### 7. Click OK.

After the encryption is complete, the KMS key that is used to encrypt the disk is automatically assigned with a fixed tag. The key of the tag is acs:ecs:disk-encryption, and the value of the tag is true. You can view the tag of the KMS key in the KMS console.

#### Encrypt a system disk by calling the CopyImage operation

The following example uses Alibaba Cloud CLI to call the Copylmage operation to specify KMSKeyld to encrypt the system disk.

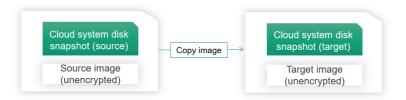
aliyun ecs Copylmage --RegionId cn-hongkong --ImageId m-bp155shrycg3s0\*\*\*\*\*\* --DestinationRegionId cn-shenzhen --Encrypted true --KmsKeyId e522b26d-abf6-4e0d-b5da-04b7\*\*\*\*\*3c --Tag.N.Key EcsDocu mentation

#### Convert the encryption state

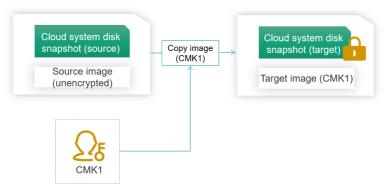
To determine the encryption state of the system disk and whether you need to change or select a new CMK, refer to the following points:

• If you do not select a CMK when copying an unencrypted image, the system disk that is

created from the new image is unencrypted.



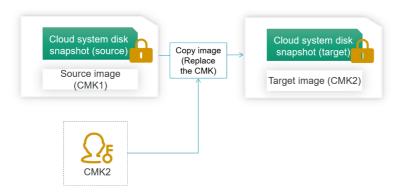
• If you select a CMK when copying an unencrypted image, the new image is encrypted. You must use this key to access instances that are created from the new image.



• If you do not select a CMK when copying an encrypted image, the new image is encrypted with the encryption key of the original image. You can use this key to access instances created from the new image.



• If you select a new CMK when copying an encrypted image, the new image is encrypted with the new CMK. You must use this key to access instances that are created from the new image.



#### What's next

You can use the copied image to create an instance or change the system disk:

Create an ECS instance

• Replace the system disk (non-public images)

#### **Related information**

- Copylmage
- CancelCopyImage

# 3.3. Encrypt a data disk

This topic describes how to encrypt a data disk. After you enable encryption on a data disk, both data in transit and data at rest on the disk is encrypted.

Alibaba Cloud data encryption security compliance ECS

#### Context

In this topic, data disks in Shared Block Storage and cloud disks are collectively called disks.

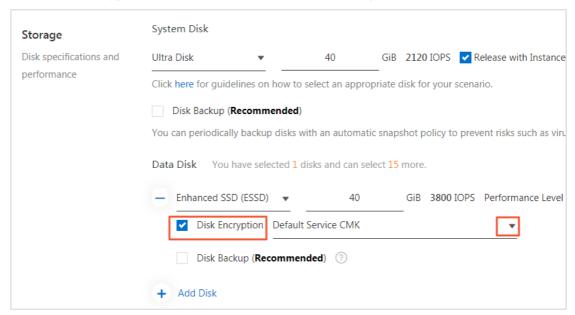
#### Encrypt a data disk when creating an instance

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. On the Instances page, click Create Instance.
- 4. On the Basic Configurations page, find the Storage section and perform the following steps.

Note This procedure describes only how to configure the encryption setting during instance creation. For more information about other configurations, see Create an instance by using the provided wizard.

- i. Click Add Disk.
- ii. Specify the disk category and capacity of the data disk.

#### iii. Select Disk Encryption and then select a key from the drop-down list.



After the encryption is complete, the KMS key that is used to encrypt the disk is automatically assigned with a fixed tag. The key of the tag is acs:ecs:disk-encryption, and the value of the tag is true. You can view the tag of the KMS key in the KMS console.

#### Encrypt a data disk when creating the disk

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

4.

5. In the Disk section, select Disk Encryption and then select a key from the drop-down list.

Note This procedure describes only how to configure the encryption setting during disk creation. For more information about other configurations, see Create a disk or Create a subscription disk.

Disk

Enhanced SSD (ESSD) ▼ 40 GIB 3800 IOPS Performance Level ②: PL1 (up to 50,000 IOPS per disk) ▼ Create from Snapshot ☑ Disk Encryption Default Service CMK



After the encryption is complete, the KMS key that is used to encrypt the disk is automatically assigned with a fixed tag. The key of the tag is acs:ecs:disk-encryption, and the value of the tag is true. You can view the tag of the KMS key in the KMS console.

#### Convert the encryption state

After a data disk is created, you cannot convert the encryption state of the data disk. The following table describes the operations to convert the encryption state.

| State conversion              | Method   | Windows Server                | Linux          |
|-------------------------------|--|-------------------------------|----------------|
| From unencrypted to encrypted | <ol> <li>Log on to the operating system of the ECS instance.</li> <li>Manually copy data on the unencrypted disk to a new encrypted disk.</li> </ol> | Run the robocopy              | Run the shell  |
| From encrypted to unencrypted | <ol> <li>Log on to the operating system of the ECS instance.</li> <li>Manually copy data on the encrypted disk to a new unencrypted disk.</li> </ol> | command in Command<br>Prompt. | command rsync. |

#### **Related information**

- Runinstances
- CreateInstance
- CreateDisk

# 4. Resize cloud disks

### 4.1. Overview

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

#### **Scenarios**

You can increase the storage capacity of an instance by using one of the following methods:

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

The following methods can be used to resize an existing disk.

| Method                   | Limit   | Procedure  |
|--------------------------|---|--|
| Resize a disk<br>online  | The instance to which the disk is attached must be in the Running (Running) state.  After the disk is resized, the new size takes effect automatically without the need to restart the instance.  | <ul> <li>Resize disks online for<br/>Linux instances</li> <li>Resize disks online for<br/>Windows instances</li> </ul>   |
| Resize a disk<br>offline | The instance to which the disk is attached must be in the Running (Running) or Stopped (Stopped) state.  After the disk is resized, you must restart its attached instance by using the ECS console or by calling the RebootInstance operation for the new size to take effect. | <ul> <li>Resize disks offline for<br/>Linux instances</li> <li>Resize disks offline for<br/>Windows instances</li> </ul> |

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

#### Size ranges of resized system disks

The sizes of resized system disks must be greater than the values before resizing, but less than or equal to 500 GiB. The following table describes the size ranges of resized system disks that attached to instances using different images.

| Image                     | Size range of a resized system disk (GiB)                   |
|---------------------------|---|
| CoreOS and FreeBSD        | [Max{30, the size of the system disk before resizing}, 500] |
| Other Linux distributions | [Max{20, the size of the system disk before resizing}, 500] |

| Image          | Size range of a resized system disk (GiB)                   |
|----------------|---|
| Windows Server | [Max{40, the size of the system disk before resizing}, 500] |

For example, the size of the system disk on a CentOS instance before resizing is 35 GiB. After the system disk is resized, the size of the new system disk must be greater than 35 GiB and less than or equal to 500 GiB.

#### Maximum sizes of resized data disks

The sizes of resized data disks must be greater than the values before resizing. The following table lists the maximum sizes of resized data disks of different categories.

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

## 4.2. Resize disks online for Linux instances

You can resize a system disk or data disk to extend its capacity. This topic describes how to resize a disk for a Linux instance without stopping the instance.

#### **Prerequisites**

The disk to be resized online and the instance to which the disk is attached meet the following requirements.

| Resource | Requirement |  |  |
|----------|-------------|--|--|
|----------|-------------|--|--|

| Resource | Requirement   |
|----------|---|
| Instance | <ul> <li>The instance is I/O optimized.</li> <li>The public image used by the instance supports online resizing of disks. For more information about images that support online resizing of disks, see Operating systems that support resizing online.</li> <li>Disks on instances of the ecs.ebmc4.8xlarge, ecs.ebmhfg5.2xlarge, and ecs.ebmg5.24xlarge instance types cannot be resized online.</li> <li>The instance is in the Running (Running) state.</li> <li>The Linux kernel version of the instance is not earlier than 3.6.0. You can run the uname -a command to check the kernel version.</li> <li>If the version of the Linux kernel is earlier than 3.6.0, you can refer to Procedure for instances with kernels earlier than 3.6.0 and Resize partitions and file systems of Linux data disks to extend partitions of a disk on the Linux instance.</li> <li>Note If your ECS instance does not meet the requirements to resize disks online, you can resize its disks offline. For more information, see Resize disks offline for Linux instances.</li> </ul>   |
| Disk     | <ul> <li>The disk is in the In Use (In Use) state.</li> <li>The disk is an enhanced SSD (ESSD), a standard SSD, or an ultra disk.</li> <li>After you renew a subscription instance and downgrade its configurations, you cannot resize the subscription disks attached to the instance for the remainder of the current billing cycle.</li> <li>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 limits" section in Limits.</li> <li>Notice If an existing partition of a disk is in the master boot record (MBR) format, the disk cannot be resized to 2 TiB or greater in size. If you want to resize a disk that has an MBR partition to more than 2 TiB, we recommend that you create a disk of more than 2 TiB, and format the disk to the GUID Partition Table (GPT) partitions. Then, copy the data from the MBR partition to the GPT partitions. For more information about how to format the disk to GPT partitions, see Partition and format a data disk larger than 2 TiB.</li> </ul> |

#### **Context**

The following table lists configurations of the instance and disks that are used in the examples of this topic.

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

| Resource    | Description   |
|-------------|---|
| System disk | /dev/vda: uses the MBR format and ext4 file system, and is resized from 40 GiB to 60 GiB.   |
| Data disk   | <ul> <li>/dev/vdb: uses the MBR format and ext4 file system, and is resized from 40 GiB to 60 GiB.</li> <li>/dev/vdc: uses the GPT format and xfs file system, and is resized from 40 GiB to 60 GiB.</li> </ul> |

#### Step 1: Create a snapshot

Create a snapshot for the disk to back up the data on the disk before you resize the disk.

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the instance whose disk you want to resize, and click the instance ID to go to the **Instance Details** page.
- 5. In the left-side navigation pane, click Disks.
- 6. Find the disk to be resized, and click Create Snapshot in the Actions column.
- 7. In the dialog box that appears, enter a snapshot name, bind a tag, and then click Create.
- 8. In the left-side navigation pane, click Snapshots to view the created snapshot. When 100% appears in the Progress column corresponding to the snapshot, the snapshot is created. You can continue to perform the following operations.

#### Step 2: Resize the disk in the console

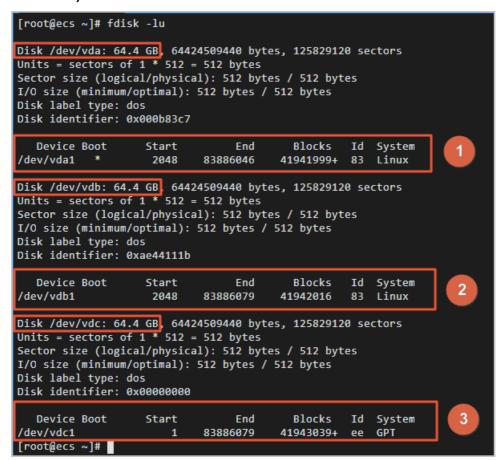
- 1. In the left-side navigation pane of the Instance Details page, click Disks.
- 2. Find the disk to be resized and choose More > Resize Disk in the Actions column.
  - To batch resize disks, log on to the ECS console by using your Alibaba Cloud account and choose Storage & Snapshots > Disks in the left-side navigation pane. On the Disks page, select the disks that you want to resize, and then click Resize Disk in the lower part of the page. Disks that are attached to the same instance cannot be batch resized.
- 3. On the Resize Disks page, select Online Resizing, and set the Size after Resize parameter. The specified Size after Resize value cannot be less than the current capacity.
- 4. Verify the price. Read and select ECS Service Terms, and then click Confirm.
- 5. Read the notes, and click I have read the notes. Resize to complete the payment.

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.

#### **Step 3: View the disk partitions**

Log on to the ECS instance to view the partition types (such as MBR and GPT) and file system types (such as ext4 and xfs) of the system disks and data disks. Subsequent resizing operations vary based on the types of the partitions and file systems.

- 1. Connect to the ECS instance. For more information, see Connect to a Linux instance by using VNC.
- 2. Run the fdisk -lu command to view the disks attached to the instance. The following figure shows the three partitions of a system disk (/dev/vda1) and two data disks (/dev/vdb1 and /dev/vdc1).



| No | Partition | Description  |
|----|-----------|--|
| •  | /dev/vda1 | A partition of the system disk. The Linux value of System indicates that the system disk is partitioned in the MBR format.       |
| 2  | /dev/vdb1 | A partition of a data disk. The Linux value of System indicates that the data disk is partitioned in the MBR format.             |
| 3  | /dev/vdc1 | A partition of a data disk. The <b>GPT</b> value of <b>System</b> indicates that the data disk is partitioned in the GPT format. |

3. Run the df -Th command to check the file system types of the existing partitions.

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

#### **Step 4: Resize partitions**

By checking the disk partitions, you can find that the partitions and file systems within the ECS instance have not been resized. This procedure describes how to resize the partitions of the resized disk within the ECS instance.

1. Install the gdisk tool on the ECS instance. You must perform this step if the partition is in the GPT format. Skip this step if the partition is in the MBR format.

```
yum install gdisk -y
```

2. Run the growpart /dev/vda 1 command to resize partitions. The partition of the system disk is resized in this example. Use a space to separate /dev/vda and 1 . To resize other partitions, modify the command accordingly.

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

#### Step 5: Resize file systems

This procedure describes how to resize the file system of a partition within an ECS instance.

- 1. Resize the file system within the ECS instance based on the file system type that you queried in Step 3.
  - o Resize the ext\* (such as ext4) file system: Run the resize2fs /dev/vda1 command.

• Resize the xfs file system: Run the xfs\_growfs /media/vdc command.

based on your partition names.

Note /media/vdc is the mount point of the /dev/vdc1 partition. Change the command based on the mount point of your partition.

2. Run the df -Th command to check the resizing results.

| [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 the disk is resized, check whether the new capacity of the disk is consistent with the expected results.

- If the resizing succeeds and the business programs in the instance can run properly, the resizing operation is complete.
- If the resizing fails, use the backup snapshot to roll back the disk.

#### Operating systems that support resizing online

The following Linux public images and custom images that are derived from the following public images support resizing online:

- Alibaba Cloud Linux: Alibaba Cloud Linux 2.1903 LTS 64-bit
- CentOS: CentOS 6.8 and later, CentOS 7.2 and later, and CentOS 8 and later
- Red Hat Enterprise Linux (RHEL): RHEL 6.9 and later, RHEL 7.4 and later, and 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

#### **FAQ**

• Problem description: When the growpart /dev/vda 1 command is run, the unexpected output in sfdisk --version [sfdisk, from util-linux 2.23.2] error is prompted.

#### Solution:

- i. Run the LANG=en\_US.UTF-8 command to switch the character encoding type of the ECS instance.
- ii. If the problem persists, run the reboot command to restart the ECS instance.
- iii. If the problem persists, run the localectl set-locale LANG=en\_US.UTF-8 command to change the local environment variable, and then restart the instance.
- Problem description: When the growpart /dev/vda 1 command is run, the -bash: growpart: command not found error is prompted.

#### Solution:

i. Run the uname -a command to check whether the version of the Linux kernel is 3.6.0 or later.

If the version of the Linux kernel is earlier than 3.6.0, you can refer to Procedure for instances with kernels earlier than 3.6.0 and Resize partitions and file systems of Linux data disks to extend partitions of a disk on the Linux instance.

- ii. Install the growpart tool based on the version of the Linux kernel.
  - CentOS 7 and later: Run the yuminstall -y cloud-utils-growpart command.
  - Debian 9 and later, or Ubuntu14 and later: Run the apt install -y cloud-guest-utils command.

#### Other scenarios for disk resizing

- If you want to create partitions on the resized disk, see Option 2: Add and format MBR partitions or Option 4: Add and format GPT partitions.
- If no partition is created on your data disk and a file system is created on the raw device, see Option 5: Resize the file system of a raw data disk to extend the file system.

#### **Related information**

- RebootInstance
- ResizeDisk
- AttachDisk

•

## 4.3. Resize disks offline for Linux instances

You can resize a system disk or a data disk to extend its capacity. If your Linux ECS instance does not support online resizing of disks, you can resize the disks on the instance offline. You must restart the instance after you resize a disk offline for the change to take effect. This will cause your business to be interrupted for a short period of time. We recommend that you resize disks offline at off-peak hours.

#### **Prerequisites**

The disk to be resized offline and the instance to which the disk is attached meet the following requirements.

| Resource | Requirement  |
|----------|--|
| Instance | <ul> <li>The instance is in the Running (Running) or Stopped (Stopped) state.</li> <li>The Linux kernel version of the instance is not earlier than 3.6.0. You can run the uname -a command to check the kernel version.</li> <li>If the version of the Linux kernel is earlier than 3.6.0, you can refer to Procedure for instances with kernels earlier than 3.6.0 and Resize partitions and file systems of Linux data disks to extend partitions of a disk on the Linux instance.</li> </ul> |

| Resource | Requirement  |  |  |  |
|----------|--|--|--|--|
|          | <ul> <li>The disk is in the In Use (In Use) state.</li> <li>The disk is an enhanced SSD (ESSD), a standard SSD, or an ultra disk.</li> <li>After you renew a subscription instance and downgrade its configurations, you cannot resize the subscription disks attached to the instance for the remainder of the current billing cycle.</li> </ul>  |  |  |  |
| Disk     | Notice If an existing partition of a disk is in the master boot record (MBR) format, the disk cannot be resized to 2 TiB or greater in size. If you want to resize a disk that has an MBR partition to more than 2 TiB, we recommend that you create a disk of more than 2 TiB, and format the disk to the GUID Partition Table (GPT) partitions. Then, copy the data from the MBR partition to the GPT partitions. For more information about how to format the disk to GPT partitions, see Partition and format a data disk larger than 2 TiB. |  |  |  |

#### Context

The following table lists configurations of the instance and disks that are used in the examples of this topic.

| Resource                   | Description   |
|----------------------------|---|
| Image used by the instance | Public image: Alibaba Cloud Linux 2.1903 LTS 64-bit   |
| System disk                | /dev/vda: uses the MBR format and ext4 file system, and is resized from 40 GiB to 60 GiB.   |
| Data disk                  | <ul> <li>/dev/vdb: uses the MBR format and ext4 file system, and is resized from 40 GiB to 60 GiB.</li> <li>/dev/vdc: uses the GPT format and xfs file system, and is resized from 40 GiB to 60 GiB.</li> </ul> |

#### Step 1: Create a snapshot

Create a snapshot for the disk to back up the data on the disk before you resize the disk.

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the instance whose disk you want to resize, and click the instance ID to go to the **Instance Details** page.
- 5. In the left-side navigation pane, click Disks.
- 6. Find the disk to be resized, and click Create Snapshot in the Actions column.
- 7. In the dialog box that appears, enter a snapshot name, bind a tag, and then click Create.
- 8. In the left-side navigation pane, click Snapshots to view the created snapshot. When 100%

appears in the **Progress** column corresponding to the snapshot, the snapshot is created. You can continue to perform the following operations.

#### Step 2: Resize the disk, and restart or start the instance in the console.

- 1. In the left-side navigation pane of the Instance Details page, click Disks.
- 2. Find the disk to be resized 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 then click **Resize Disk** in the lower part of the page. Disks that are attached to the same instance cannot be batch resized.
- 3. On the Resize Disks page, set the Size after Resize parameter. The specified Size after Resize value cannot be less than the current capacity.
- 4. Verify the price. Read and select ECS Service Terms, and then click Confirm.
- 5. Read the notes, and click I have read the notes. Resize to complete the payment.
- 6. Restart or start the instance in the ECS console.

#### ? Note

- If your instance is in the Running (Running) state, you must restart the instance in the console. If your instance is in the Stopped (Stopped) state, you must start the instance in the console.
- The disk resizing operation takes effect only after you restart or start the
  instance by using the console or by calling the RebootInstance operation. If the
  ECS instance is restarted from within the instance itself, the disk capacity will not
  change.

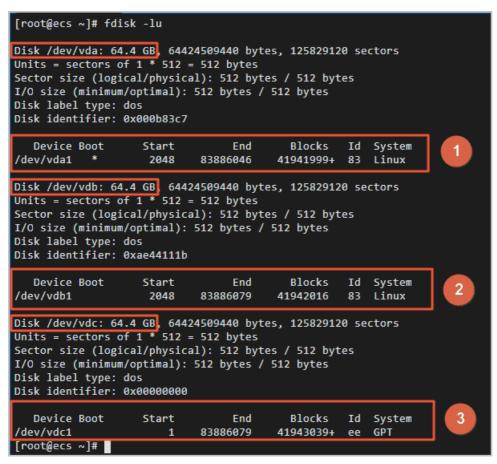
The following procedure demonstrates how to restart an instance.

- i. In the left-side navigation pane, choose Instances & Images > Instances.
- ii. Find the instance to be restarted, and choose More > Instance Status > Restart in the Actions column.
- iii. In the Restart Instance dialog box, click OK.

#### Step 3: View the disk partitions

Log on to the ECS instance to view the partition types (such as MBR and GPT) and file system types (such as ext4 and xfs) of the system disks and data disks. Subsequent resizing operations vary based on the types of the partitions and file systems.

- 1. Connect to the ECS instance. For more information, see Connect to a Linux instance by using VNC.
- 2. Run the fdisk -lu command to view the disks attached to the instance. The following figure shows the three partitions of a system disk (/dev/vda1) and two data disks (/dev/vdb1 and /dev/vdc1).



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

3. Run the df -Th command to check the file system types of the existing partitions.

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

#### **Step 4: Resize partitions**

By checking the disk partitions, you can find that the partitions and file systems within the ECS instance have not been resized. This procedure describes how to resize the partitions of the resized disk within the ECS instance.

1. Install the gdisk tool on the ECS instance. You must perform this step if the partition is in the GPT format. Skip this step if the partition is in the MBR format.

```
yum install gdisk -y
```

2. Run the growpart /dev/vda 1 command to resize partitions. The partition of the system disk is resized in this example. Use a space to separate /dev/vda and 1 . To resize other partitions, modify the command accordingly.

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

#### Step 5: Resize file systems

This procedure describes how to resize the file system of a partition within an ECS instance.

- 1. Resize the file system within the ECS instance based on the file system type that you queried in Step 3.
  - Resize the ext\* (such as ext4) file system: Run the resize2fs /dev/vda1 command.

# 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 a data disk, run the following command .

resize2fs /dev/vdb1

- Note /dev/vda1 and /dev/vdb1 are partition names. Change the commands based on your partition names.
- Resize the xfs file system: Run the xfs growfs /media/vdc command.
  - Note /media/vdc is the mount point of the /dev/vdc1 partition. Change the command based on the mount point of your partition.
- 2. Run the df -Th command to check the resizing results.

| [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 the disk is resized, check whether the new capacity of the disk is consistent with the expected results.

- If the resizing succeeds and the business programs in the instance can run properly, the resizing operation is complete.
- If the resizing fails, use the backup snapshot to roll back the disk.

#### **FAQ**

• Problem description: When the growpart /dev/vda 1 command is run, the unexpected output in sfdisk --version [sfdisk, from util-linux 2.23.2] error is prompted.

#### Solution:

- i. Run the LANG=en\_US.UTF-8 command to switch the character encoding type of the ECS instance.
- ii. If the problem persists, run the reboot command to restart the ECS instance.
- iii. If the problem persists, run the localectl set-locale LANG=en\_US.UTF-8 command to change the local environment variable, and then restart the instance.
- Problem description: When the growpart /dev/vda 1 command is run, the -bash: growpart: command not found error is prompted.

#### Solution:

- i. Run the uname -a command to check whether the version of the Linux kernel is 3.6.0 or later.
- ii. Install the growpart tool based on the version of the Linux kernel.
  - CentOS 7 and later: Run the yuminstall -y cloud-utils-growpart command.
  - Debian 9 and later, or Ubuntu14 and later: Run the apt install -y cloud-guest-utils command.

#### Other scenarios for disk resizing

- If you want to create partitions on the resized disk, see Option 2: Add and format MBR partitions or Option 4: Add and format GPT partitions.
- If no partition is created on your data disk and a file system is created on the raw device, see Option 5: Resize the file system of a raw data disk to extend the file system.

#### **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 instance without stopping the instance.

#### **Prerequisites**

The disk to be resized online and the instance to which the disk is attached meet the following requirements.

| Resource | Requirement  |
|----------|--|
| Instance | <ul> <li>The instance is I/O optimized.</li> <li>The instance does not run Windows Server 2003.</li> <li>Disks on instances of the ecs.ebmc4.8xlarge, ecs.ebmhfg5.2xlarge, and ecs.ebmg5.24xlarge instance types cannot be resized online.</li> <li>The instance is in the Running (Running) state.</li> <li>The version of the Red Hat virtio SCSI driver is later than 58011. For information about how to check the version of a Red Hat virtio SCSI driver and upgrade the driver, see Update the Red Hat VirtIO driver of Windows ECS instances.</li> </ul> |
|          | Note If your ECS instance does not meet the requirements to<br>resize disks online, you can resize its disks offline. For more<br>information, see Resize disks offline for Windows instances.   |

| Resource | Requirement  |
|----------|--|
|          | <ul> <li>The disk is in the In Use (In Use) state.</li> <li>The disk is an enhanced SSD (ESSD), a standard SSD, or an ultra disk.</li> <li>The disk is formatted as New Technology File System (NTFS).</li> <li>After you renew a subscription instance and downgrade its configurations, you cannot resize the subscription disks attached to the instance for the remainder of the current billing cycle.</li> <li>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 limits" section in Limits.</li> </ul> |
|          | Notice  If an existing partition of a disk is in the master boot record (MBR) format, the disk cannot be resized to 2 TiB or greater in size. If you want to resize a disk that has an MBR partition to more than 2 TiB, use one of the following methods:  • (Recommended) Create a disk of more than 2 TiB, and format   |
| Disk     | the disk to GUID Partition Table (GPT) partitions. Then, copy the data from the MBR partition to the GPT partitions. For more information about how to format the disk to GPT partitions, see Partition and format a data disk larger than 2 TiB.  Covert the MBR partition to a GPT partition within the instance. The risk of data loss may occur when you change the partition format. Proceed with caution. For more information, see Convert the partition format of a data disk on a Windows instance.   |
|          |  |

#### Context

The following table lists configurations of the instance and disks that are used in the examples of this topic.

| Resource                   | Description   |
|----------------------------|---|
| Image used by the instance | Windows Server 2012 R2 64-bit                           |
| System disk                | Is formatted as NTFS and resized from 40 GiB to 60 GiB. |

| Resource  | Description   |
|-----------|---|
| Data disk | Is formatted as NTFS and resized from 40 GiB to 60 GiB. |

#### Step 1: Create a snapshot

Create a snapshot for the disk to back up the data on the disk before you resize the disk.

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the instance whose disk you want to resize, and click the instance ID to go to the **Instance Details** page.
- 5. In the left-side navigation pane, click Disks.
- 6. Find the disk to be resized, and click Create Snapshot in the Actions column.
- 7. In the dialog box that appears, enter a snapshot name, bind a tag, and then click Create.
- 8. In the left-side navigation pane, click **Snapshots** to view the created snapshot. When **100**% appears in the **Progress** column corresponding to the snapshot, the snapshot is created. You can continue to perform the following operations.

#### Step 2: Resize the disk in the console

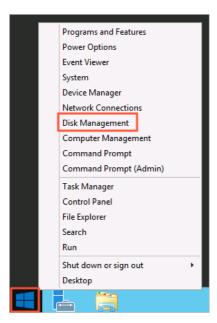
- 1. In the left-side navigation pane of the Instance Details page, click Disks.
- 2. Find the disk to be resized and choose More > Resize Disk in the Actions column.
  - To batch resize disks, log on to the ECS console by using your Alibaba Cloud account and choose **Storage & Snapshots** > **Disks** in the left-side navigation pane. On the Disks page, select the disks that you want to resize, and then click **Resize Disk** in the lower part of the page. Disks that are attached to the same instance cannot be batch resized.
- 3. On the Resize Disks page, select Online Resizing, and set the Size after Resize parameter. The specified Size after Resize value cannot be less than the current capacity.
- 4. Verify the price. Read and select ECS Service Terms, and then click Confirm.
- 5. Read the notes, and click I have read the notes. Resize to complete the payment.

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.

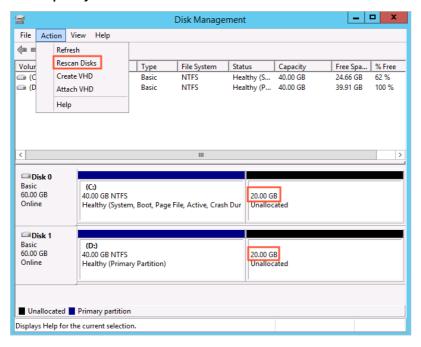
#### Step 3: Resize the file system of the system disk or data disk partition

After you resize a disk in the ECS console, the file systems of the disk partitions are not resized. You must connect to the instance to extend file systems. The disk is formatted as NTFS and will be resized from 40 GiB to 60 GiB. This procedure describes how to resize the file system of a system disk partition.

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



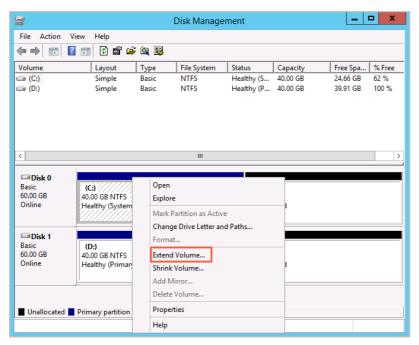
3. In the Disk Management dialog box, choose Action > Rescan Disks to view the unallocated disk capacity.



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

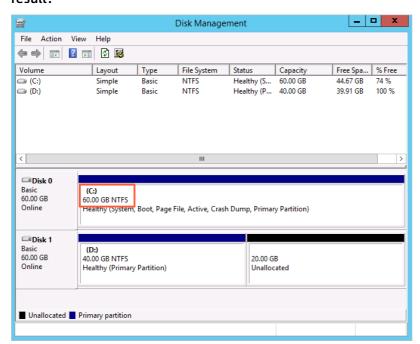
4. Right-click the blank space in the Disk O section and select Extend Volume...

Note If you are resizing the file system of a data disk partition, right-click the blank space in the data disk (for example, Disk 1) section to extend the file system.



5. Extend the volume as instructed in the Extend Volume Wizard dialog box.

The new space is automatically added to the original volume. The following figure shows the result.



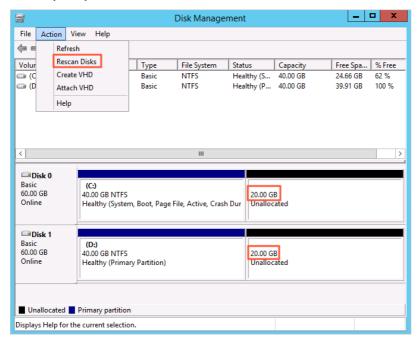
After the disk is resized, check whether the new capacity of the disk is consistent with the expected results.

- If the resizing succeeds and the business programs in the instance can run properly, the resizing operation is complete.
- o If the resizing fails, use the backup snapshot to roll back the disk.

#### (Optional) Step 4: Create new partitions on the data disk

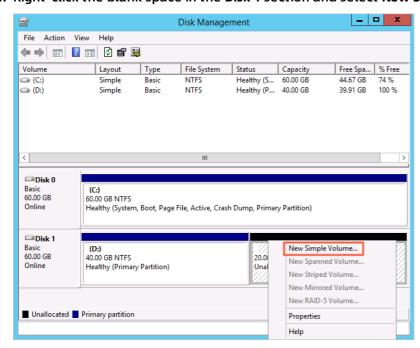
If you want to create new partitions in the new disk space, complete the following procedure.

- 1. Connect to the ECS instance. For more information, see Connect to a Windows instance by using VNC.
- 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 the **New Simple Volume Wizard** dialog box. A new 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 data disk to extend its capacity. If your Windows instance does not support online resizing of disks, you can resize the disks of the instance offline. After you resize a disk on an instance offline, you must restart the instance for the changes to take effect.

#### **Prerequisites**

The disk to be resized offline and the Windows instance to which it is attached meet the following requirements.

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

| Resource | Requirement   |
|----------|---|
|          | <ul> <li>The disk is in the In Use (In Use) state.</li> <li>The disk is an enhanced SSD, a standard SSD, or an ultra disk.</li> <li>The disk is formatted with an NTFS file system.</li> <li>After you renew a subscription instance and downgrade its configurations, you cannot resize the subscription disks attached to the instance during the remainder of the current billing cycle.</li> <li>When you resize a disk of a specific category, the specified new disk capacity must not exceed the maximum single-disk capacity allowed for the disk category. For more information, see the "Elastic Block Storage limits" section in Limits</li> </ul> |
|          | Notice  If an existing partition of a disk is in the master boot record (MBR) format, the disk cannot be resized to 2 TiB or greater in size. If you want to resize a disk that has an MBR partition to more than 2 TiB, use one of the following methods:  |
| Disk     | <ul> <li>(Recommended) Create a disk of more than 2 TiB, and format<br/>the disk to GUID Partition Table (GPT) partitions. Then, copy the<br/>data from the MBR partition to the GPT partitions. For more<br/>information about how to format the disk to GPT partitions, see<br/>Partition and format a data disk larger than 2 TiB.</li> </ul>  |
|          | <ul> <li>Covert the MBR partition to a GPT partition within the instance.         The risk of data loss may occur when you change the partition format. Proceed with caution. For more information, see Convert the partition format of a data disk on a Windows instance.     </li> </ul>  |
|          |   |
|          |   |
|          |   |

#### Step 1: Create a snapshot

Create a snapshot for the disk to back up the data on the disk before you resize the disk.

- 1. Log on to the ECS console.
- 2. In the left-side navigation pane, choose Instances & Images > Instances.
- 3. In the top navigation bar, select a region.
- 4. Find the instance whose disk you want to resize, and click the instance ID to go to the **Instance Details** page.
- 5. In the left-side navigation pane, click Disks.

- 6. Find the disk to be resized, and click Create Snapshot in the Actions column.
- 7. In the dialog box that appears, enter a snapshot name, bind a tag, and then click Create.
- 8. In the left-side navigation pane, click Snapshots to view the created snapshot. When 100% appears in the Progress column corresponding to the snapshot, the snapshot is created. You can continue to perform the following operations.

#### Step 2: Resize the disk, and restart or start the instance in the console.

- 1. In the left-side navigation pane of the Instance Details page, click Disks.
- 2. Find the disk to be resized 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 then click **Resize Disk** in the lower part of the page. Disks that are attached to the same instance cannot be batch resized.
- 3. On the Resize Disks page, set the Size after Resize parameter. The specified Size after Resize value cannot be less than the current capacity.
- 4. Verify the price. Read and select ECS Service Terms, and then click Confirm.
- 5. Read the notes, and click I have read the notes. Resize to complete the payment.
- 6. Restart or start the instance in the ECS console.

#### ? Note

- If your instance is in the Running (Running) state, you must restart the instance in the console. If your instance is in the Stopped (Stopped) state, you must start the instance in the console.
- The disk resizing operation takes effect only after you restart or start the instance by using the console or by calling the RebootInstance operation. If the ECS instance is restarted from within the instance itself, the disk capacity will not change.

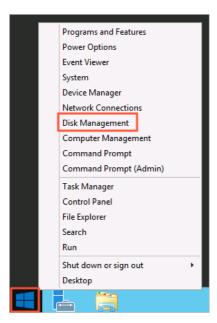
The following procedure demonstrates how to restart an instance.

- i. In the left-side navigation pane, choose Instances & Images > Instances.
- ii. Find the instance to be restarted, and choose More > Instance Status > Restart in the Actions column.
- iii. In the Restart Instance dialog box, click OK.

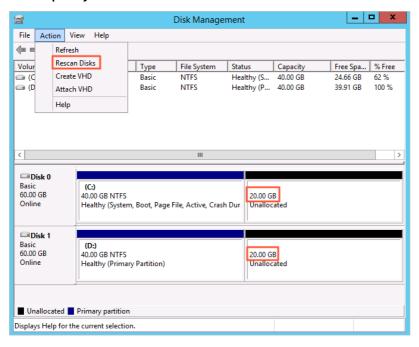
#### Step 3: Resize the file system of the system disk or data disk partition

After you resize a disk in the ECS console, the file systems of the disk partitions are not resized. You must connect to the instance to extend file systems. The disk is formatted as NTFS and will be resized from 40 GiB to 60 GiB. This procedure describes how to resize the file system of a system disk partition.

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



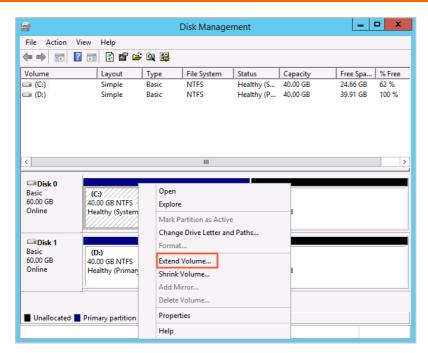
3. In the Disk Management dialog box, choose Action > Rescan Disks to view the unallocated disk capacity.



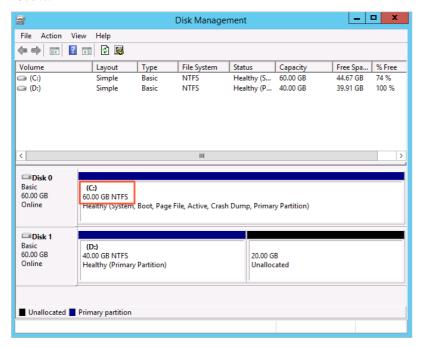
Disk 0 is the system disk, and Disk 1 is a data disk.

4. Right-click the blank space in the Disk O section and select Extend Volume...

Note If you are resizing the file system of a data disk partition, right-click the blank space in the data disk (for example, Disk 1) section to extend the file system.



Extend the volume as instructed in the Extend Volume Wizard dialog box.The new space is automatically added to the original volume. The following figure shows the result.



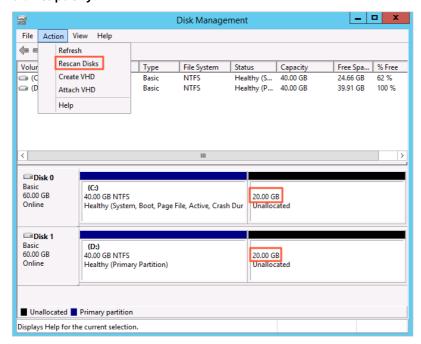
After the disk is resized, check whether the new capacity of the disk is consistent with the expected results.

- If the resizing succeeds and the business programs in the instance can run properly, the resizing operation is complete.
- o If the resizing fails, use the backup snapshot to roll back the disk.

#### (Optional) Step 4: Create new partitions on the data disk

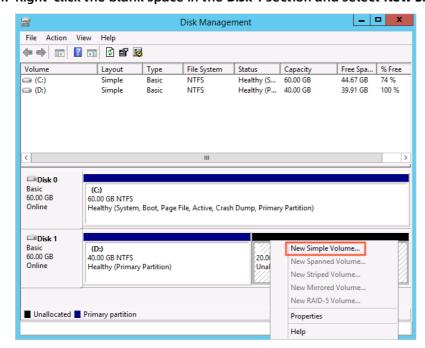
If you want to create new partitions in the new disk space, complete the following procedure.

- 1. Connect to the ECS instance. For more information, see Connect to a Windows instance by using VNC.
- 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 the **New Simple Volume Wizard** dialog box. A new partition is created in the new disk space.

#### **Related information**

- RebootInstance
- ResizeDisk
- AttachDisk

•

# 4.6. Update the Red Hat VirtIO driver of Windows ECS instances

You can resize the disks of Windows ECS instances online without having to restart the instances. You can also query the serial numbers of disks from the operating systems. If you want to resize the disks of a Windows ECS instance created before March 30, 2019 or query the serial numbers of the disks of an instance created before May 1, 2020, follow the procedure described in this topic to check whether the Red Hat VirtIO driver of the instances needs to be updated.

Alibaba Cloud resize a disk extend disk capacity resize a disk of a Windows instance

#### Context

- The Red Hat VirtIO driver is supported only on Windows Server 2008 and later.
- If an ECS instance has multiple data disks attached, the driver update process may take a few minutes to complete.

#### **Procedure**

You can follow the following procedure to update the Red Hat VirtIO driver of a Windows ECS instance:

- 1. Step 1. Check the driver version
- 2. Step 2. Download the driver
- 3. Step 3. Update the Red Hat VirtIO driver

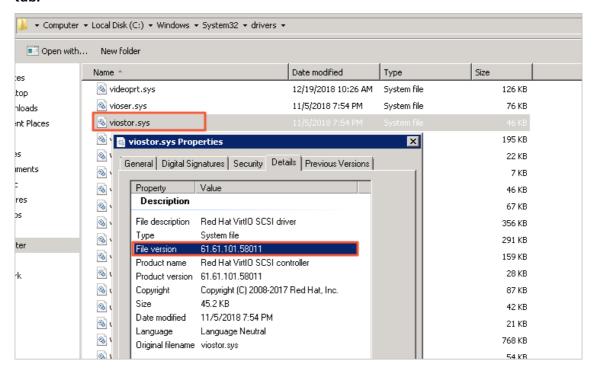
#### Step 1. Check the driver version

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

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

 $[System. Diagnostics. File Version Info] :: Get Version Info ("C:\Windows\System 32\drivers\viostor.s ys")$ 

- Method 2: Manually check the driver version
  - i. Connect to the Windows instance.
  - ii. Go to the system directory C:\Windows\System32\drivers.
  - iii. Right-click the *viostor.sys* file, select **Properties**, and 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      | Red Hat VirtIO SCSI driver version | Operation   |
|---------------|------------------------------------|---|
| Resize a disk | 58011 or later                     | Resize the disk online. For more information, see<br>Resize disks online for Linux instances. |
| online        | Earlier than 58011                 | Proceed to the next step.   |
| Query the     | 58015 or later                     | Query the serial number of the disk.  |
| of a disk     | Earlier than 58015                 | Proceed to the next step.   |

#### Step 2. Download the driver

Download and decompress the VirtIO driver package. The subsequent steps in this topic are based on the assumption that the decompressed driver package is located at C:\Users\Administrator\Desktop\virtioDriver. The following table lists the extracted folders corresponding to ECS instances of different operating system versions.

| Driver file (folder) name | Description                          |
|---------------------------|--------------------------------------|
| win7                      | Windows Server 2008 R2 and Windows 7 |

| Driver file (folder) name | Description   |
|---------------------------|---|
| Wlh                       | Windows Server 2008   |
| Win8                      | Windows Server 2012 and Windows Server 2012 R2                                |
| win10                     | Windows 10, Windows Server 2016, and Windows Server systems of later versions |
| amd64                     | 64-bit  |
| x86                       | 32-bit  |

#### Step 3. Update the Red Hat VirtIO driver

You can use one of the following methods to update the Red Hat VirtIO driver:

- Method 1: Use pnputil to add and install the driver
  - i. Open Command Prompt.
  - ii. Run the following command to add the driver package:

```
pnputil -i -a <path to virtio driver inf>
```

Make sure that you have extracted the .inf target file to the specified directory <path to vi rtio driver inf>. (Example: C:\Users\Administrator\Desktop\virtioDriver\win7\amd64\viost or.inf)

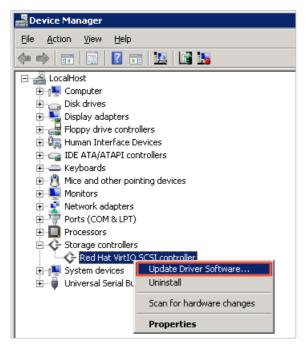
■ To update only the disk driver, run the following command:

■ To update all Red Hat VirtIO drivers, run the following command:

```
pnputil -i -a C:\Users\Administrator\Desktop\virtioDriver\win7\amd64\*.inf
```

- iii. Restart the operating system of the ECS instance for the driver update to take effect.
- Method 2: Manually add the driver
  - i. Open Device Manager.
  - ii. Right-click Red Hat VirtIO SCSI controller under Storage controllers and select Update Driver Software....

Note If multiple Red Hat VirtIO SCSI controller devices appear, you only need to update one of them.



- iii. Select Browse my computer for driver software.
- iv. Select Let me pick from a list of device drivers on my computer.
- v. Click Have Disk.
- vi. Select the driver file *viostor* in the corresponding folder, and follow the wizard to update the driver.
- vii. Restart the operating system of the ECS instance for the driver update to take effect.

#### What's next

• For more information about how to resize a disk online, see Resize a disk online.

# 5.Performance5.1. EBS performance

This topic describes the performance metrics and specifications of various Elastic Block Storage (EBS) devices, such as cloud disks and local disks.

#### Performance metrics

The key metrics used to measure EBS performance include Input/output operations per second (IOPS), throughput, and latency. Some EBS devices also have capacity requirements. For example, enhanced SSDs (ESSDs) of different performance levels (PLs) have different capacity ranges.

#### IOPS

The IOPS metric measures the number of read/write operations that can be performed per second. The metric indicates the performance of transaction intensive applications, such as databases. A standard SSD can implement the committed IOPS performance only when it is attached to an I/O optimized instance. Common IOPS performance metrics that are measured include sequential and random operations. The following table describes these performance metrics.

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

#### • Throughput

The throughput metric measures the size of data transferred per second. Unit: MB/s. This metric indicates the performance of applications that require a large number of sequential read/write operations, such as Hadoop offline computing applications.

#### Latency

Latency is the period of time required for an EBS device to process an I/O request. Unit: s, ms, or µs. High latency may cause performance degradation or errors in applications that have requirements for low latency.

• For latency-sensitive applications such as databases, we recommend that you use ESSDs, standard SSDs, or local SSDs.

 For applications that require high throughput but do not require low latency such as Hadoop offline computing applications, we recommend that you use ECS instances attached with local SATA HDDs of the d1 or d1ne instance family.

#### Capacity

Capacity is the size 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 the capacity as a metric to measure the performance of EBS devices, but the performance of EBS devices varies based on capacity. The larger the capacity of an EBS device is, the stronger its processing capabilities are. EBS devices of the same type have the same I/O performance per unit capacity. However, the performance of each disk increases based on capacity until it reaches the upper limit performance for that single disk. ESSDs of different capacity ranges have different PLs.

For information about how to test the performance of different types of EBS devices, see Performance tests on Block Storage Or Test the IOPS performance of an enhanced SSD.

#### **Disk performance**

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

| Performa<br>nce<br>category   | ESSD   |  |   |   | Standard<br>SSD                             | Ultra<br>disk                             | Basic<br>disk <sup>3</sup> |
|---|--|--|---|---|---|---|----------------------------|
| PL  | PL3  | PL2  | PL1   | PL0   | None  | None                                      | None                       |
| Maximum<br>capacity<br>of a<br>single<br>disk (GiB)                           | 1,261 to<br>32,768                             | 461 to<br>32,768                             | 20 to<br>32,768                             | 40 to<br>32,768                             | 32,768                                      | 32,768                                    | 2,000                      |
| Maximum<br>IOPS   | 1,000,000                                      | 100,000                                      | 50,000                                      | 10,000                                      | 25,000 <sup>1</sup>                         | 5,000                                     | Several<br>hundred<br>s    |
| Maximum<br>throughp<br>ut (MB/s)  | 4,000  | 750  | 350   | 180   | 300 <sup>1</sup>                            | 140                                       | 30 to 40                   |
| Formula<br>to<br>calculate<br>the IOPS<br>of a<br>single<br>disk <sup>2</sup> | min{1,800<br>+ 50 ×<br>Capacity,<br>1,000,000} | min{1,800<br>+ 50 ×<br>Capacity,<br>100,000} | min{1,800<br>+ 50 ×<br>Capacity,<br>50,000} | min{1,800<br>+ 12 ×<br>Capacity,<br>10,000} | min{1,800<br>+ 30 ×<br>Capacity,<br>25,000} | min{1,800<br>+ 8 ×<br>Capacity,<br>5,000} | None                       |

| Performa<br>nce<br>category   | ESSD                                      |   |   | Standard<br>SSD                          | Ultra<br>disk                           | Basic<br>disk <sup>3</sup>               |                 |
|---|---|---|---|--|---|--|-----------------|
| Formula<br>to<br>calculate<br>the<br>throughp<br>ut of a<br>single<br>disk<br>(MB/s) <sup>2</sup> | min{120 +<br>0.5 ×<br>Capacity,<br>4,000} | min{120 +<br>0.5 ×<br>Capacity,<br>750} | min{120 +<br>0.5 ×<br>Capacity,<br>350} | min{100 +<br>0.25 ×<br>Capacity,<br>180} | min{120 +<br>0.5 ×<br>Capacity,<br>300} | min{100 +<br>0.15 ×<br>Capacity,<br>140} | None            |
| Data<br>reliability   | 99.999999<br>9%                           | 99.999999<br>9%                         | 99.999999<br>9%                         | 99.99999<br>99%                          | 99.99999<br>99%                         | 99.99999<br>99%                          | 99.99999<br>99% |
| Average single- channel random write latency in milliseco nds (block size = 4 KB)                 | 0.2                                       |   |   | 0.3 to 0.5                               | 0.5 to 2                                | 1 to 3                                   | 5 to 10         |
| API<br>paramet<br>er value  | cloud_essd                                |   |   | cloud_ss<br>d                            | cloud_eff<br>iciency                    | cloud                                    |                 |

| Performa<br>nce<br>category | ESSD  | Standard<br>SSD  | Ultra<br>disk  | Basic<br>disk <sup>3</sup>  |
|-----------------------------|---|--|--|---|
| Scenario                    | <ul> <li>OLTP databases: relational databases such as MySQL, PostgreSQL, Oracle, and SQL Server databases</li> <li>NoSQL databases: non-relational databases such as MongoDB, HBase, and Cassandra databases</li> <li>Elasticsearch distributed logs: Elasticsearch, Logstash, and Kibana (ELK) log analysis</li> </ul> | I/O intensi ve applic ations Small and mediu m-sized relatio nal datab ases NoSQL datab ases | <ul> <li>Develo pment and test busine ss</li> <li>Syste m disks</li> </ul> | <ul> <li>Applic ations that are not freque ntly access ed or have low I/O loads</li> <li>Applic ations that require low costs and rando m I/O operations</li> </ul> |

<sup>&</sup>lt;sup>1</sup> 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 |

<sup>&</sup>lt;sup>2</sup> A standard SSD is used in the following examples to describe how to calculate the performance of a single disk:

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

#### Local disks

<sup>&</sup>lt;sup>3</sup> Basic disks are part of a previous generation of disks and are no longer available for purchase.

For information about the performance of local NVMe SSDs and SATA HDDs, see Local disks.

### 5.2. Storage I/O performance

Storage I/O performance, also known as storage read/write performance, is the performance that is implemented 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. IOPS  $\times$  I/O size = Throughput. You can select different Block Storage devices and instance types based on the I/O request features of the applications to achieve the best results:

- If the applications require a large amount of I/O, such as offline analysis and data warehousing applications, we recommend that you select the big data instance families that have higher throughput.
- If the applications require low-latency, random, and small-sized I/O, such as OLTP databases and enterprise-level applications such as SAP, we recommend that you select enhanced SSDs (ESSDs) and standard SSDs that have higher 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, hfc7, g6, c6, r6, and g5se. 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 for Alibaba Cloud ECS provides storage I/O performance isolation. Dedicated storage bandwidths are assigned to ECS instances and disks to avoid the 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 the new generation of instance family that provides storage I/O performance isolation:

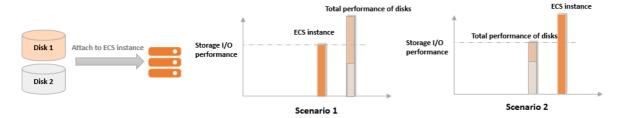
- Large and medium-sized database loads, such as Oracle, MySQL, SQL Server, PostgreSQL,
   Cassandra, and MongoDB databases.
- Enterprise-level applications, such as ERP and CRM.

#### Relationship between instance types and storage I/O performance

The storage I/O performance of an ECS instance varies depending on instance families and diskattached instance types. 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 60,000 instead of 81,800. For information about the performance levels of ESSDs, see Enhanced SSDs.

After you understand the relationship between instance storage performance and disk storage performance, you can choose instance types and Block Storage devices that meet your business needs based on the performance data of instances and disks. This prevents performance bottlenecks caused by improper configurations.

## Storage I/O performance of the new generation of enterprise-level instance families

The following table lists the storage I/O performance of the new generation of enterprise-level instance families. For information about other metrics of the instance families, see Instance families.

Note In this table, the maximum IOPS is measured at 4 KiB I/O size and the maximum throughput is measured at 1,024 KiB I/O size. For information about test methods, see Performance tests on Block Storage.

#### Storage I/O performance of hfg7

| Instance type   | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-----------------|--------------------------------|---------------------------------------|---|
| ecs.hfg7.large  | 20.0                           | 1.0                                   | 125   |
| ecs.hfg7.xlarge | 30.0                           | 1.5                                   | 187.5                                       |

| Instance type     | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-------------------|--------------------------------|---------------------------------------|---|
| ecs.hfg7.2xlarge  | 45.0                           | 2.0                                   | 250   |
| ecs.hfg7.3xlarge  | 60.0                           | 2.5                                   | 312.5                                       |
| ecs.hfg7.4xlarge  | 75.0                           | 3.0                                   | 375   |
| ecs.hfg7.6xlarge  | 90.0                           | 4.0                                   | 500   |
| ecs.hfg7.8xlarge  | 105.0                          | 5.0                                   | 625   |
| ecs.hfg7.12xlarge | 150.0                          | 8.0                                   | 1000  |
| ecs.hfg7.24xlarge | 300.0                          | 16.0                                  | 2000  |

### Storage I/O performance of hfc7

| Instance type     | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-------------------|--------------------------------|------------------------------------|---|
| ecs.hfc7.large    | 20.0                           | 1.0                                | 125   |
| ecs.hfc7.xlarge   | 30.0                           | 1.5                                | 187.5                                       |
| ecs.hfc7.2xlarge  | 45.0                           | 2.0                                | 250   |
| ecs.hfc7.3xlarge  | 60.0                           | 2.5                                | 312.5                                       |
| ecs.hfc7.4xlarge  | 75.0                           | 3.0                                | 375   |
| ecs.hfc7.6xlarge  | 90.0                           | 4.0                                | 500   |
| ecs.hfc7.8xlarge  | 105.0                          | 5.0                                | 625   |
| ecs.hfc7.12xlarge | 150.0                          | 8.0                                | 1000  |
| ecs.hfc7.24xlarge | 300.0                          | 16.0                               | 2000  |

### Storage I/O performance of hfr7

| Instance type    | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|------------------|--------------------------------|---------------------------------------|---|
| ecs.hfr7.large   | 20.0                           | 1.0                                   | 125   |
| ecs.hfr7.xlarge  | 30.0                           | 1.5                                   | 187.5                                       |
| ecs.hfr7.2xlarge | 45.0                           | 2.0                                   | 250   |
| ecs.hfr7.3xlarge | 60.0                           | 2.5                                   | 312.5                                       |

| Instance type     | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-------------------|--------------------------------|---------------------------------------|---|
| ecs.hfr7.4xlarge  | 75.0                           | 3.0                                   | 375   |
| ecs.hfr7.6xlarge  | 90.0                           | 4.0                                   | 500   |
| ecs.hfr7.8xlarge  | 105.0                          | 5.0                                   | 625   |
| ecs.hfr7.12xlarge | 150.0                          | 8.0                                   | 1000  |
| ecs.hfr7.24xlarge | 300.0                          | 16.0                                  | 2000  |

### Storage I/O performance of g6e

| Instance type    | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|------------------|--------------------------------|---------------------------------------|---|
| ecs.g6e.large    | 20.0                           | 1.0                                   | 125   |
| ecs.g6e.xlarge   | 40.0                           | 1.5                                   | 187.5                                       |
| ecs.g6e.2xlarge  | 50.0                           | 2.0                                   | 250   |
| ecs.g6e.4xlarge  | 80.0                           | 3.0                                   | 375   |
| ecs.g6e.8xlarge  | 150.0                          | 5.0                                   | 625   |
| ecs.g6e.13xlarge | 240.0                          | 8.0                                   | 1000  |
| ecs.g6e.26xlarge | 480.0                          | 16.0                                  | 2000  |

### Storage I/O performance of c6e

| Instance type    | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|------------------|--------------------------------|------------------------------------|---|
| ecs.c6e.large    | 20.0                           | 1.0                                | 125   |
| ecs.c6e.xlarge   | 40.0                           | 1.5                                | 187.5                                       |
| ecs.c6e.2xlarge  | 50.0                           | 2.0                                | 250   |
| ecs.c6e.4xlarge  | 80.0                           | 3.0                                | 375   |
| ecs.c6e.8xlarge  | 150.0                          | 5.0                                | 625   |
| ecs.c6e.13xlarge | 240.0                          | 8.0                                | 1000  |
| ecs.c6e.26xlarge | 480.0                          | 16.0                               | 2000  |

### Storage I/O performance of r6e

| Instance type    | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|------------------|--------------------------------|---------------------------------------|---|
| ecs.r6e.large    | 20.0                           | 1.0                                   | 125   |
| ecs.r6e.xlarge   | 40.0                           | 1.5                                   | 187.5                                       |
| ecs.r6e.2xlarge  | 50.0                           | 2.0                                   | 250   |
| ecs.r6e.4xlarge  | 80.0                           | 3.0                                   | 375   |
| ecs.r6e.8xlarge  | 150.0                          | 5.0                                   | 625   |
| ecs.r6e.13xlarge | 240.0                          | 8.0                                   | 1000  |
| ecs.r6e.26xlarge | 480.0                          | 16.0                                  | 2000  |

### Storage I/O performance of g6

| Instance type   | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-----------------|--------------------------------|------------------------------------|---|
| ecs.g6.large    | 10.0                           | 1.0                                | 125   |
| ecs.g6.xlarge   | 20.0                           | 1.5                                | 187.5                                       |
| ecs.g6.2xlarge  | 25.0                           | 2.0                                | 250   |
| ecs.g6.3xlarge  | 30.0                           | 2.5                                | 312.5                                       |
| ecs.g6.4xlarge  | 40.0                           | 3.0                                | 375   |
| ecs.g6.6xlarge  | 50.0                           | 4.0                                | 500   |
| ecs.g6.8xlarge  | 60.0                           | 5.0                                | 625   |
| ecs.g6.13xlarge | 100.0                          | 8.0                                | 1000  |
| ecs.g6.26xlarge | 200.0                          | 16.0                               | 2000  |

### Storage I/O performance of c6

| Instance type  | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|----------------|--------------------------------|------------------------------------|---|
| ecs.c6.large   | 10.0                           | 1.0                                | 125   |
| ecs.c6.xlarge  | 20.0                           | 1.5                                | 187.5                                       |
| ecs.c6.2xlarge | 25.0                           | 2.0                                | 250   |
| ecs.c6.3xlarge | 30.0                           | 2.5                                | 312.5                                       |

| Instance type   | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-----------------|--------------------------------|---------------------------------------|---|
| ecs.c6.4xlarge  | 40.0                           | 3.0                                   | 375   |
| ecs.c6.6xlarge  | 50.0                           | 4.0                                   | 500   |
| ecs.c6.8xlarge  | 60.0                           | 5.0                                   | 625   |
| ecs.c6.13xlarge | 100.0                          | 8.0                                   | 1000  |
| ecs.c6.26xlarge | 200.0                          | 16.0                                  | 2000  |

### Storage I/O performance of r6

| Instance type   | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-----------------|--------------------------------|------------------------------------|---|
| ecs.r6.large    | 10.0                           | 1.0                                | 125   |
| ecs.r6.xlarge   | 20.0                           | 1.5                                | 187.5                                       |
| ecs.r6.2xlarge  | 25.0                           | 2.0                                | 250   |
| ecs.r6.3xlarge  | 30.0                           | 2.5                                | 312.5                                       |
| ecs.r6.4xlarge  | 40.0                           | 3.0                                | 375   |
| ecs.r6.6xlarge  | 50.0                           | 4.0                                | 500   |
| ecs.r6.8xlarge  | 60.0                           | 5.0                                | 625   |
| ecs.r6.13xlarge | 100.0                          | 8.0                                | 1000  |
| ecs.r6.26xlarge | 200.0                          | 16.0                               | 2000  |

### Storage I/O performance of hfg6

| Instance type    | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|------------------|--------------------------------|---------------------------------------|---|
| ecs.hfg6.large   | 10.0                           | 1.0                                   | 125   |
| ecs.hfg6.xlarge  | 20.0                           | 1.5                                   | 187.5                                       |
| ecs.hfg6.2xlarge | 25.0                           | 2.0                                   | 250   |
| ecs.hfg6.3xlarge | 30.0                           | 2.5                                   | 312.5                                       |
| ecs.hfg6.4xlarge | 40.0                           | 3.0                                   | 375   |
| ecs.hfg6.6xlarge | 50.0                           | 4.0                                   | 500   |

| Instance type     | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-------------------|--------------------------------|---------------------------------------|---|
| ecs.hfg6.8xlarge  | 60.0                           | 5.0                                   | 625   |
| ecs.hfg6.10xlarge | 100.0                          | 8.0                                   | 1000  |
| ecs.hfg6.16xlarge | 120.0                          | 10.0                                  | 1250  |
| ecs.hfg6.20xlarge | 200.0                          | 16.0                                  | 2000  |

### Storage I/O performance of hfc6

| Instance type     | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-------------------|--------------------------------|---------------------------------------|---|
| ecs.hfc6.large    | 10.0                           | 1.0                                   | 125   |
| ecs.hfc6.xlarge   | 20.0                           | 1.5                                   | 187.5                                       |
| ecs.hfc6.2xlarge  | 25.0                           | 2.0                                   | 250   |
| ecs.hfc6.3xlarge  | 30.0                           | 2.5                                   | 312.5                                       |
| ecs.hfc6.4xlarge  | 40.0                           | 3.0                                   | 375   |
| ecs.hfc6.6xlarge  | 50.0                           | 4.0                                   | 500   |
| ecs.hfc6.8xlarge  | 60.0                           | 5.0                                   | 625   |
| ecs.hfc6.10xlarge | 100.0                          | 8.0                                   | 1000  |
| ecs.hfc6.16xlarge | 120.0                          | 10.0                                  | 1250  |
| ecs.hfc6.20xlarge | 200.0                          | 16.0                                  | 2000  |

### Storage I/O performance of hfr6

| Instance type    | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|------------------|--------------------------------|------------------------------------|---|
| ecs.hfr6.large   | 10.0                           | 1.0                                | 125   |
| ecs.hfr6.xlarge  | 20.0                           | 1.5                                | 187.5                                       |
| ecs.hfr6.2xlarge | 25.0                           | 2.0                                | 250   |
| ecs.hfr6.3xlarge | 30.0                           | 2.5                                | 312.5                                       |
| ecs.hfr6.4xlarge | 40.0                           | 3.0                                | 375   |
| ecs.hfr6.6xlarge | 50.0                           | 4.0                                | 500   |

| Instance type     | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage<br>bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-------------------|--------------------------------|---------------------------------------|---|
| ecs.hfr6.8xlarge  | 60.0                           | 5.0                                   | 625   |
| ecs.hfr6.10xlarge | 100.0                          | 8.0                                   | 1000  |
| ecs.hfr6.16xlarge | 120.0                          | 10.0                                  | 1250  |
| ecs.hfr6.20xlarge | 200.0                          | 16.0                                  | 2000  |

#### Storage I/O performance of g5se

| Instance type     | Maximum IOPS (K, 4<br>KiB I/O) | Maximum storage bandwidth (Gbit/s) | Maximum throughput<br>(MB/s, 1,024 KiB I/O) |
|-------------------|--------------------------------|------------------------------------|---|
| ecs.g5se.large    | 30.0                           | 1.0                                | 125   |
| ecs.g5se.xlarge   | 60.0                           | 2.0                                | 250   |
| ecs.g5se.2xlarge  | 120.0                          | 4.0                                | 500   |
| ecs.g5se.4xlarge  | 230.0                          | 8.0                                | 1000  |
| ecs.g5se.6xlarge  | 340.0                          | 12.0                               | 1500  |
| ecs.g5se.8xlarge  | 450.0                          | 15.0                               | 1875  |
| ecs.g5se.16xlarge | 900.0                          | 30.0                               | 3750  |
| ecs.g5se.18xlarge | 1000.0                         | 32.0                               | 4000  |

### 5.3. Performance tests on Block Storage

This topic describes how to use the fio tool on a Lunix instance to test performance of cloud disks and local disks, including IOPS, throughput, and latency.

the performance of cloud disks the performance of Block Storage the performance of the enhanced SSD performance metrics IOPS capacity throughput throughput

#### **Prerequisites**

A Block Storage device is created and attached to an ECS instance.

Note If you only want to test the performance of a Block Storage device of a specific type, we recommend that you use a newly created pay-as-you-go data disk. You can release the disk at any time after the test is completed.

#### Context

You can use other tools to test the performance of Block Storage devices, but the benchmark you obtain may be different. For example, tools such as dd, sysbench, and iometer may be affected by test parameter settings and file systems, which causes inaccurate results. The performance results in this topic are that of a Linux instance tested by using fio. These results are used as performance references for Alibaba Cloud Block Storage products. We recommend that you use the fio tool to test the performance of Block Storage 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 test the partition directly. Before you perform such a test, you must create a snapshot of the disk to back up your data. For more information about how to create a snapshot, see Create a normal snapshot. We recommend that you use newly purchased ECS instances that contain no data for the test to prevent data loss.

#### **Procedure**

- 1. Remotely connect to an ECS instance. For more information, see Connect to a Linux instance by using VNC.
- 2. Before you test a Block Storage device, ensure 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 continue the performance test.

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

3. Run the following commands to install the libaio and fio tools:

```
sudo yum install libaio -y
sudo yum install libaio-devel -y
sudo yum install fio -y
```

4. Switch the directory.

```
cd /tmp
```

5. Run the performance test commands. For more information about the commands, see the following section.

#### Commands used to test the performance of cloud disks

For more information about the IOPS test methods regarding the enhanced SSD, see Test the IOPS performance of an enhanced SSD.

Random write IOPS:

```
fio -direct=1 -iodepth=128 -rw=randwrite -ioengine=libaio -bs=4k -size=1G -numjobs=1 -runtime=100 
0 -group_reporting -filename=iotest -name=Rand_Write_Testing
```

#### • Random read IOPS:

fio -direct=1 -iodepth=128 -rw=randread -ioengine=libaio -bs=4k -size=1G -numjobs=1 -runtime=100 0 -group\_reporting -filename=iotest -name=Rand\_Read\_Testing

• Sequential write throughput (write bandwidth):

fio -direct=1 -iodepth=64 -rw=write -ioengine=libaio -bs=1024k -size=1G -numjobs=1 -runtime=1000 group\_reporting -filename=iotest -name=Write\_PPS\_Testing

Sequential read throughput (read bandwidth):

fio -direct=1 -iodepth=64 -rw=read -ioengine=libaio -bs=1024k -size=1G -numjobs=1 -runtime=1000 - group\_reporting -filename=iotest -name=Read\_PPS\_Testing

• Random write latency:

fio -direct=1 -iodepth=1 -rw=randwrite -ioengine=libaio -bs=4k -size=1G -numjobs=1 -group\_reporting -filename=iotest -name=Rand\_Write\_Latency\_Testing

• Random read latency:

fio -direct=1 -iodepth=1 -rw=randread -ioengine=libaio -bs=4k -size=1G -numjobs=1 -group\_reportin g -filename=iotest -name=Rand\_Read\_Latency\_Testingrandwrite -ioengine=libaio -bs=4k -size=1G -numjobs=1 -group\_reporting -filename=iotest -name=Rand\_Write\_Latency\_Testing

#### Commands used to test the performance of local disks

The following test commands are only applicable to local NVMe SSD disks.

• Random write IOPS:

fio -direct=1 -iodepth=32 -rw=randwrite -ioengine=libaio -bs=4k -numjobs=4 -time\_based=1 -runtim e=1000 -group\_reporting -filename=/dev/vdx -name=test

• Random read IOPS:

fio -direct=1 -iodepth=32 -rw=randread -ioengine=libaio -bs=4k -numjobs=4 -time\_based=1 -runtime =1000 -group\_reporting -filename=/dev/vdx -name=test

• Sequential write throughput (write bandwidth):

fio -direct=1 -iodepth=128 -rw=write -ioengine=libaio -bs=128k -numjobs=1 -time\_based=1 -runtime =1000 -group\_reporting -filename=/dev/vdx -name=test

Sequential read throughput (read bandwidth):

fio -direct=1 -iodepth=128 -rw=read -ioengine=libaio -bs=128k -numjobs=1 -time\_based=1 -runtime= 1000 -group\_reporting -filename=/dev/vdx -name=test

• Random write latency:

fio -direct=1 -iodepth=1 -rw=randwrite -ioengine=libaio -bs=4k -numjobs=1 -time\_based=1 -runtime =1000 -group\_reporting -filename=/dev/vdx -name=test

#### • Random read latency:

 $fio-direct=1-iodepth=1-rw=randread-ioengine=libaio-bs=4k-numjobs=1-time\_based=1-runtime=1000-group\_reporting-filename=/dev/vdx-name=test$ 

#### • Sequential write latency:

fio -direct=1 -iodepth=1 -rw=write -ioengine=libaio -bs=4k -numjobs=1 -time\_based=1 -runtime=100 0 -group\_reporting -filename=/dev/vdx -name=test

#### • Sequential read latency:

fio -direct=1 -iodepth=1 -rw=read -ioengine=libaio -bs=4k -numjobs=1 -time\_based=1 -runtime=1000 -group\_reporting -filename=/dev/vdx -name=test

#### fio parameter values

The following table uses the command for testing random write IOPS (*randwrite*) as an example to describe the parameter settings in fio commands.

| Parameter     | Description  |
|---------------|--|
| -direct=1     | Indicates that the I/O buffer is ignored during the test and data is written directly.   |
| -iodepth=128  | Indicates that when you use asynchronous I/O, a maximum of 128 I/O requests can be made at the same time.  |
| -rw=randwrite | Indicates that the read/write policy is random writes. Other valid values for rw:  • randread (random reads)  • read (sequential reads)  • write (sequential writes)  • randrw (random reads and writes) |

| Parameter                        | Description  |
|----------------------------------|--|
| -ioengine=libaio                 | <ul> <li>Indicates that libaio (the Linux-native asynchronous I/O facility) is used as the testing method. There are two ways for an application to use I/O:</li> <li>Synchronous         In synchronous I/O mode, a thread sends only one I/O request at a time and waits until the I/O request is completed. In this case, the iodepth value is always smaller than 1 for a single thread. You can increase the iodepth value by using multiple concurrent threads. The iodepth value will reach its upper limit when there are 16 to 32 threads running concurrently.     </li> <li>Asynchronous         In asynchronous I/O mode, a thread uses libaio to send multiple I/O requests at a time and waits until all these I/O requests are completed. Asynchronous I/O reduces the number of interactions and increases interaction efficiency.     </li> </ul> |
| -bs=4k                           | <ul> <li>Indicates that the size of each block for one I/O is 4 KB. If this parameter is not specified, the default value 4 KB is used.</li> <li>When IOPS is tested, we recommend that you set bs to a small value, such as 4k in this example command.</li> <li>When throughput is tested, we recommend that you set bs to a large value, such as 1024k in other example commands.</li> </ul>  |
| -size=1G                         | Indicates that the size of the testing file is 1 GiB.  |
| -numjobs=1                       | Indicates that the number of testing threads is 1.   |
| -runtime=1000                    | Indicates that the test duration is 1,000 seconds. If this parameter is not specified, the file of the size specified by -size is written in blocks of the size specified by -bs.  |
| -group_reporting                 | Indicates that in the testing results, the statistics are displayed for a group of threads, not for each individual thread.  |
| -filename=iotest                 | Indicates that the name of the test file is <i>iotest</i> .  |
| -<br>name=Rand_Write_Testi<br>ng | Indicates that the name of the test task is <i>Rand_Write_Testing</i> . You can specify any name you want.   |

# 5.4. Test the IOPS performance of an enhanced SSD

This topic describes how to test the IOPS performance of an enhanced SSD. The disk specifications and testing conditions will affect the test result. If you configure the testing conditions as in the following example, you can achieve an IOPS of one million when you perform a performance stress test on the enhanced SSD.

#### **Context**

**Testing conditions:** 

- Sample operation: random write (randwrite).
- Image: We recommend that you use a later version of a public Linux image provided by Alibaba Cloud, such as CentOS 7.4 64-bit, CentOS 7.3 64-bit, CentOS 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.g5se.18xlarge.
- Enhanced SSD: We recommend that you use a PL3 enhanced SSD. In this example, the device name of the enhanced SSD is /dev/vdb. For more information, see Enhanced SSDs.

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 test the partition directly. Before you test a raw disk, we recommend that you create a snapshot of the disk to back up your data. For more information about how to create a snapshot, see Create a normal snapshot. We recommend that you use a new ECS instance that contains no data for the test to prevent data loss.

#### **Procedure**

- 1. Remotely connect to an ECS instance. For more information, see Connect to a Linux instance by using VNC.
- 2. Run the following commands to install the libaio and fio tools:

```
sudo yum install libaio -y
sudo yum install libaio-devel -y
sudo yum install fio -y
```

3. Switch the directory.

cd /tmp

4. Create the test 100 w.sh script.

vim test100w.sh

5. Paste the following code to the *test100w.sh* script:

```
function RunFio
numjobs=$1 # The number of threads to be tested. In this example, the value is 10.
iodepth=$2 # The maximum number of concurrent I/O requests. In this example, the value is 64.
           # The size of the data block per I/O. In this example, the value is 4K.
           # The read and write policy. In this example, the value is randwrite.
filename=$5 # The name of the file to be tested. In this example, the value is /dev/vdb.
nr_cpus=`cat /proc/cpuinfo |grep "processor" |wc -l`
if [ $nr cpus -lt $numjobs ];then
   echo "Numjobs is more than cpu cores, exit!"
  exit -1
fi
let nu=$numjobs+1
cpulist=""
for ((i=1;i<10;i++))
do
  list=`cat /sys/block/vdb/mq/*/cpu_list | awk '{if(i<=NF) print $i;}' i="$i" | tr -d ',' | tr '\n' ','`
  if [ -z $list ];then
     break
  fi
  cpulist=${cpulist}${list}
spincpu='echo $cpulist | cut -d ',' -f 2-${nu}'
echo $spincpu
fio --ioengine=libaio --runtime=30s --numjobs=${numjobs} --iodepth=${iodepth} --bs=${bs} --rw=$
{rw} --filename=${filename} --time_based=1 --direct=1 --name=test --group_reporting --cpus_allo
wed=$spincpu --cpus_allowed_policy=split
echo 2 > /sys/block/vdb/queue/rq_affinity
sleep 5
RunFio 10 64 4k randwrite /dev/vdb
```

- 6. Modify the parameter settings in the test 100w.sh script based on your actual environment.
  - Set vdb to the actual device name of the enhanced SSD.
  - Modify *10, 64, 4k, randwrite,* and */dev/vdb* in RunFio 10 64 4k randwrite /dev/vdb as needed.
  - o If you want to proceed with this operation, set filename to a device name such as [/dev/vdb]. If you do not want to risk data loss, set filename to a file path such as [/mnt/test.image].
- 7. Test the performance of the enhanced SSD.

sh test100w.sh

If the IOPS=\*\*\* result appears, the performance test of the enhanced SSD is complete.

#### Details of the test 100w.sh script

• In the script, the following command sets the rq\_affinity parameter to 2:

echo 2 > /sys/block/vdb/queue/rq\_affinity

| Value of rq_affinity | Description  |
|----------------------|--|
| 1                    | Indicates that the block device forwards the received I/O Completion events to the vCPU group that submits the corresponding I/O requests. When multiple threads concurrently process I/O requests, the I/O Completion events may run on the same vCPU, which results in a performance bottleneck. |
| 2                    | Indicates that the block device forwards the received I/O Completion events to the vCPU that submits the corresponding I/O requests. The performance of each vCPU is fully delivered when multiple threads concurrently process I/O requests.  |

• The following command binds multiple jobs to different CPU cores:

 $fio -ioengine=libaio -runtime=30s -numjobs=\$\{numjobs\} -iodepth=\$\{iodepth\} -bs=\$\{bs\} -rw=\$\{rw\} -filename=\$\{filename\} -time\_based=1 -direct=1 -name=test -group\_reporting -cpus\_allowed=\$spincpu -cpus\_allowed\_policy=split$ 

Note Typically, a device has only one request queue. This unique request queue becomes a performance bottleneck when multiple threads concurrently process I/O requests. When a device is in multi-queue mode, it can have multiple request queues to process I/O requests and deliver the full performance of the backend storage service. If you have four I/O threads, you must bind them to the CPU cores that correspond to different request queues. This allows you to leverage the full capabilities of multi-queue mode to improve storage performance.

| Parameter           | Description  | Example  |
|---------------------|--|----------|
| numjobs             | The number of I/O threads.   | 10       |
| /dev/vdb            | The device name of the enhanced SSD.   | /dev/vdb |
| cpus_allowed_policy | The cpus_allowed_policy and cpus_allowed parameters provided by fio to bind vCPUs. | split    |

The preceding command runs multiple jobs that are bound to different CPU cores and correspond to different queue IDs. To view the ID of the CPU core that is bound to a queue ID, perform the following steps:

- Run the ls /sys/block/vd\*/mq/ command to view the queue ID of the enhanced SSD whose device name is in the /dev/vd\* format.
- Run the cat /sys/block/vd\*/mq//cpu\_list command to view the ID of the CPU core that is bound to the queue of the enhanced SSD whose device name is in the /dev/vd\* format.

### 6.Block Storage FAQ

This topic provides answers to commonly asked questions about using Block Storage.

#### Enhanced SSD FAQ

- What is an enhanced SSD (ESSD)?
- What specifications do enhanced SSDs have?
- What are the similarities and differences between enhanced SSDs, standard SSDs, and ultra disks?
- How is the performance level of an enhanced SSD measured?
- How do I test the performance of an enhanced SSD and obtain the 1 million IOPS result?
- What is the relationship between the storage performance of an enhanced SSD and storage performance of the instance to which the enhanced SSD is attached?
- In which regions and zones are enhanced SSDs available?
- Which instance families can be attached with enhanced SSDs?

#### Common FAO

- 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 guery the usage and free space of Block Storage 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?

#### • Performance testing FAQ

- What tools can I use to test the performance of Block Storage devices?
- Why did my instance shut down when I used fio to test the I/O performance of the instance?
- How do I test the performance of an enhanced SSD and obtain the 1 million IOPS result?

#### • Standard SSD FAQ

- 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?
- Can I upgrade a standard SSD after I purchase it?
- Why is an error reported 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 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 one disk to multiple ECS instances at the same time?

- Do I need to partition and format a pay-as-you-go disk after I purchase it and attach it to an ECS instance?
- Why am I unable to find the data disk that I purchased for a Linux instance?
- How many disks can be attached to one ECS instance?
- Why am I unable to find the desired ECS instance when I attempt to attach a disk to it?
- Can I attach a disk to an ECS instance that is in a different zone?
- Will data in a data disk be lost when I detach the disk?
- Can I detach a system disk?

#### • Independent disk FAQ

- How is a separately purchased pay-as-you-go data disk billed?
- I attached a disk that I created separately to an ECS instance. Why is the disk released when the instance is released?
- Can I attach a separately purchased 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 a disk when I renewed it. Can I convert the billing method of the disk from subscription to pay-as-you-go within the remaining time of the current subscription?

#### • Disk snapshot FAQ

- When I delete a disk, will its snapshots also be deleted?
- Why are some automatic snapshots on my disk missing?
- Can I use a snapshot to create an independent disk?

#### FAQ about 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?
- Are my snapshots retained if I re-initialize a disk?
- Data disks of a Linux instance cannot be found after I restarted the instance or re-initialized the system disk. What do I do?
- How do I reattach data disks after I re-initialize the system disk of a Linux instance?

#### • FAQ about extending 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 extend a system disk?
- Can I shrink a system disk after I extend it?
- What Block Storage devices can be extended when they are used as system disks? Do any regional limits apply to this operation?
- Can the system disks of both subscription ECS instances and pay-as-you-go ECS instances be extended?
- What is the storage capacity range of a system disk?
- I changed the configurations of an instance when I renewed it. 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 to extend the data disk without risking 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 extend a disk of a Linux instance?

#### • Partition FAQ

- Can I partition a data disk for data storage?
- For a disk with multiple partitions, are snapshots created for the entire disk or only for a specific partition?
- What must I be aware of before I repartition a disk?
- What is the relationship between writing data and partitioning and formatting?

#### • FAQ about rolling back disks

- I rolled back a data disk by using a snapshot after I repartitioned the disk. How many partitions are available in the disk?
- An error message similar to the following one is returned when I attempt to roll back a disk:
   "A disk can be rolled back only when the attached instance 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 replacement 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 enhanced SSD (ESSD)?

An enhanced SSD is an ultra-high performance disk provided by Alibaba Cloud. Enhanced SSDs use 25 GE networks and the remote direct memory access (RDMA) technology to deliver up to 1 million random IOPS with low one-way latency. For more information, see Enhanced SSDs.

#### What specifications do enhanced SSDs have?

Enhanced SSDs have four specifications based on their maximum performance. For the latest information about enhanced SSD performance, see <a href="Enhanced SSDs">Enhanced SSDs</a>.

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

| PL  | Capacity of ESSD (GiB) | Maximum IOPS | Maximum throughput<br>(Mbit/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 enhanced SSDs, standard SSDs, and ultra disks?

- Similarities: All three categories of disks are based on a distributed Block Storage architecture, provide high reliability and scalability, and support snapshots and data encryption.
- Differences: Enhanced SSDs have the best performance among the three categories of disks. For more information, see Enhanced SSDs and EBS performance.

#### How is the performance level of an enhanced SSD measured?

The performance level of an enhanced SSD is proportional to its storage capacity. An enhanced SSD that has a larger storage capacity delivers higher performance. Compared with standard SSDs, enhanced SSDs have much better performance. For more information, see Enhanced SSDs.

### How do I test the performance of an enhanced SSD and obtain the 1 million IOPS result?

We recommend that you use fio (flexible IO tester) to perform a stress test on an enhanced SSD. For more information, see Test the IOPS performance of an enhanced SSD.

# What is the relationship between the storage performance of an enhanced SSD and storage performance of the instance to which the enhanced SSD is attached?

Instance type specifications impact the performance of instance-level storage. The higher specifications an instance type has, the higher IOPS and throughput it delivers.

For example, when you create an instance of the g5se storage optimized instance family with enhanced performance and attach enhanced SSDs to the instance:

- If the total storage performance of the enhanced SSDs does not exceed the maximum storage performance that the instance type can deliver, the instance delivers the total storage performance of the enhanced SSDs.
- If the total storage performance of the enhanced SSDs exceeds the maximum storage performance that the instance type can deliver, the storage performance of the instance is limited to the maximum storage performance that the instance type can deliver.

For example, you create a 16 GiB instance of the ecs.g5se.xlarge instance type that can deliver up to 60,000 IOPS. If you attach an enhanced SSD to the instance and the enhanced SSD has a storage capacity of 2 TiB with a maximum IOPS of 101,800, the maximum IOPS of the instance is 60,000, instead of 101,800.

For information about the performance and specifications of the g5se instance family, see <a href="Instance families">Instance families</a>.

#### In which regions and zones are enhanced SSDs available?

Enhanced SSDs are available in the following regions and zones:

- China (Hangzhou): Zone G, Zone H, and Zone I
- China (Shanghai): Zone E, Zone F, and Zone G
- China (Beijing): Zone F, Zone G, and Zone H
- China (Zhangjiakou-Beijing Winter Olympics): Zone A and Zone B

• China (Shenzhen): Zone D and Zone E

• China (Chengdu): Zone A

• China (Hong Kong): Zone C

• India (Mumbai): Zone B

• UK (London): Zone B

• Australia (Sydney): Zone B

#### Which instance families can be attached with enhanced SSDs?

Enhanced SSDs can be attached only to instances from the instance families that support 25 GE networks (g6, g5, ic5, c6, c5, r6, r5, and g5se), ECS Bare Metal Instance families (ebmhfg5, ebmc4, ebmg5, ebmgn6v, ebmgn6i, ebmc5s, ebmg5s, ebmr5s, and sccgn6), and enterprise-level heterogeneous computing instance families (vgn5i, gn6v, gn5, gn5i, gn4, ga1, f1, and f3). For more information, see Instance families.

#### What tools can I use to test the performance of Block Storage devices?

For details, see Performance tests on Block Storage.

## Why did my instance shut down when I used 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. However, performing this test on raw disk partitions may damage the metadata of the file systems in the raw disk partitions, which prevents you from accessing files in the partitions. This can also cause the instance to shut down. This problem does not occur when you use fio to test the I/O performance of the instance by testing 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 is in the same zone as the disk.

- For high-availability applications, we recommend that you create data disks in different zones and attach them to ECS instances.
- For low-latency applications, we recommend that you create data disks in the same zone as the ECS instances and attach them to the 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 Disk overview.

#### How do I query the usage and free space of Block Storage devices?

You can log on to an ECS instance to query the usage and free space of Block Storage devices in the instance. You cannot query the usage and free space of Block Storage devices by using the ECS console or by calling ECS API operations.

#### Can I shrink a disk?

No, you cannot shrink disks. 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 convert it into a pay-as-you-go data disk first. Before you release the resulted pay-as-you-go data disk, ensure that you have backed up all important data stored in it. For more information, see Change billing methods of disks and Release a disk.

Note The disks that have been converted from subscription billing to pay-as-you-go billing are billed by hour. Charges for pay-as-you-go data disks stop one hour after the disks are released. 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 it to an I/O optimized instance.

You can call the ModifyInstanceSpec and ModifyPrepayInstanceSpec operations to upgrade your non-I/O optimized instances to I/O optimized instances.

#### What is the I/O performance of a standard SSD?

For details, see EBS performance.

#### What scenarios are standard SSDs ideal for?

Standard SSDs provide high performance and high reliability. They are ideal for I/O-intensive applications, 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 which 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 the physical storage media.

#### Can I upgrade a standard SSD after I purchase it?

Yes, you can upgrade and scale up your standard SSDs. For more information, see Overview.

# Why is an error reported 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 reported. Specify a mount point in the /dev/vd\* format in the command to mount a partition of the standard SSD.

### What must I be aware of before I add mount information of basic disks or standard SSDs to Linux instances?

To attach a data disk to a Linux instance, you must format and partition the disk. When you perform this operation, note that /dev/xvdb1 is the mount point for a basic disk, and /dev/vdb1 is the mount point for an ultra disk, a standard SSD, or an enhanced SSD. If the information specified in the mount -a command is incorrect, the disk fails to be attached. To avoid this problem, perform the following steps:

- 1. Run the fdisk -l command to view information about the data disk.
- 2. Check whether the information of the data disk added to /etc/fstab is the same as that in Step 1.

Note Do not add duplicate mount information because it will cause a system startup failure.

- 3. Run the vim command to modify the /etc/fstab file.
- 4. Comment out or delete incorrect information and add correct mount information.
- 5. Run the mount -a command to check whether the disk is attached.

For more information, see Format a data disk for 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 and matches the disk sequence number in the disk manager in Windows.

#### What is an independent disk?

An independent disk is a pay-as-you-go data disk that you purchase separately. An independent disk can be attached to or detached from any ECS instance in the same zone as the disk. 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 one disk to multiple ECS instances at the same time?

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

### Do I need to partition and format a pay-as-you-go disk after I purchase it and attach it to an ECS instance?

Yes. After you purchase a pay-as-you-go data disk separately, you must attach it to an ECS instance and then partition and format the disk. For more information, see Format a data disk for a Linux instance and Format a data disk for a Windows ECS instance.

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

If you purchase a pay-as-you-go data disk separately, you must attach it to the instance and partition it before you can view and use its storage space. For more information, see Format a data disk for a Linux instance and Attach a data disk.

#### How many disks can be attached to one ECS instance?

When used as data disks, a maximum of 16 data disks can be attached to one ECS instance.

### Why am I unable to find the desired ECS instance when I attempt to attach a disk to it?

Check whether the ECS instance has been released. If no, ensure that it is in the same zone as the disk.

#### Can I attach a disk to an ECS instance that is in a different zone?

No, a pay-as-you-go disk can be attached only to an ECS instance in the same zone as the disk.

#### 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.
- In Linux, you must log on to the ECS instance and run the umount command on the disk. After the command is executed, log on to the ECS console to detach the disk.

#### Can I detach a system disk?

Yes, you can detach a system disk.

#### How is a separately purchased pay-as-you-go data disk billed?

A pay-as-you-go data disk is billed by the hour. Note that if your account balance is insufficient, the services of the data disk will be suspended.

### I attached a disk that I created separately 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 in the ECS console or by calling an API operation. For more information, see Release a disk.

# Can I attach a separately purchased pay-as-you-go data disk to a subscription instance?

Yes, you can attach a separately purchased pay-as-you-go data disk to a subscription instance.

#### Can I detach a data disk from a subscription instance?

No, you cannot detach data disks 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, convert it into a pay-as-you-go data disk first. Then detach and release the data disk. For information about how to change the billing method of disks, see Change billing methods of disks.

#### I changed the configurations of a disk when I renewed it. Can I convert the billing method of the disk from subscription to pay-as-you-go within the remaining time of the current subscription?

No, you can convert 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 also deleted when you delete the disk. However, the manual snapshots are retained. You can change this setting at 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, 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 steps to fix the error for the Linux operating system:

- 1. Find the data disk and use one of the following methods to check whether it is attached to the corresponding ECS instance:
  - Perform the check in the ECS console. For more information, see View the monitoring data of a cloud disk.
  - Log on to the instance, and run the partition information is correct. Run the df -h and mount | grep "<devpath>" commands to view the mount information.
- 2. Run the cat command to view the /etc/fstab file and check whether you have attached two disks to the same directory.
  - If two disks are attached to the same directory, the second disk will replace the first disk, causing data of the first disk to become inaccessible. Attach the second disk to a different directory.

If they are attached to different directories but the mount information shows that they
are in 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 new 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 an ECS Linux instance?

- Problem description: All data in a directory such as /alidata is lost after you restart an ECS 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: This solution uses an I/O-optimized instance as an example. If your ECS instance is a non-I/O optimized instance, enter the device names of the disk partitions in the /dev/xvd\*1 format in the mount command.
  - i. Run the fdisk-l 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 preceding directory.
  - iii. Run the df -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 during next system startup to avoid this problem.

#### Are my snapshots retained if I re-initialize a disk?

Yes, both manual and automatic snapshots in the disk are retained.

# Data disks of a Linux instance cannot be found after I restarted the instance or re-initialized the system disk. 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:
  - Restarting 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-initializing the system disk: The /etc/fstab file is reset after the system disk is re-initialized. As a result, data disks are not automatically attached during system startup.
- Solution:
  - 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 file system of the /dev/xvdb1 partition is ext3. Run the following command to write the partition mount information to the /etc/fstab file:

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

iv. Restart the instance in the ECS console.

### How do I reattach 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 that are attached to the instance remains unchanged, but the mount information of the data disks is lost. As a result, all the data disks become detached. 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 steps 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 information of /dev/vdb1.

- 2. Run the fdisk-l command to view information about 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 steps 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 the mount information.

If the command output does not contain 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 the mount information.

If the command output shows 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 current system disk before you replace it. Make sure that the new system disk has enough space available. The recommended available disk space is 1 GiB or larger. If disk space is insufficient, the instance may not start properly after you replace the system disk.

#### How do I extend a system disk?

You can extend a system disk in the ECS console or by calling the ResizeDisk operation.

#### Can I shrink a system disk after I extend it?

No, you cannot shrink a system disk after you extend it. We recommend that you extend the system disk reasonably based on your needs.

# What Block Storage devices can be extended when they are used as system disks? Do any regional limits apply to this operation?

Ultra disks, standard SSDs, and enhanced SSDs can be extended when they are used as system disks. You can extend system disks in all regions.

# Can the system disks of both subscription ECS instances and pay-as-you-go ECS instances be extended?

Yes, the system disks of both subscription ECS instances and pay-as-you-go ECS instances can be extended.

#### What is the storage capacity range of a system disk?

The capacity range of a system disk varies with the operating system. For more information, see Overview.

# I changed the configurations of an instance when I renewed it. 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 renewing it, you can extend its system disk only when the new billing cycle starts.

# How do I create a disk from a snapshot of a data disk to extend the data disk without risking data loss?

If a data disk cannot be extended without data loss due to a disk error, you can purchase a payas-you-go disk to temporarily store data from the original data disk, and then format the original data disk. To do so, perform the following steps:

- 1. Create a snapshot of the current data disk (original data disk). For more information, see Create a normal snapshot.
- 2. Go to the disk purchase page. Select the region and zone of the ECS instance to purchase a pay-as-you-go disk. Click Create from Snapshot. In the dialog box that appears, select the snapshot created in the previous step.
- 3. Log on to the ECS console and then attach the new data disk you purchased 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 mkfs.ext3 to repartition and reformat the new data disk, extending it to the target capacity. For more information, see Resize partitions and file systems of Linux data disks.
- 7. Run the cp -R command to copy all the data in the new data disk back to the original data disk.

You can add the --preserve=all setting to retain the file properties when you copy the files.

- 8. Run the umount command to detach the new data disk.
- 9. In the ECS console, detach the new data disk from the ECS instance, and release the disk.

# What do I do if the "Bad magic number in super-block while trying to open /dev/xvdb1" error message is returned when I extend a disk of a Linux instance?

- Problem description: When you run the Linux ECS instance, the "Bad magic number in super-block while trying to open /dev/xvdb1" error message is returned.
- Cause: The disk to be extended is not partitioned.
- Solution: Run the e2fsck-f/dev/xvdb and resize2fs/dev/xvdb commands to extend the disk.

  Then run the mount command to attach the disk.

#### Can I partition a data disk for data storage?

Yes, you can split a data disk into multiple partitions as needed. 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 repartition a disk?

To ensure data security, we recommend that you create a snapshot of a disk before repartitioning the disk. This way, you can roll back the disk if you perform an incorrect operation. For more information, see Create a snapshot and Roll back a disk by using a snapshot.

## What is the relationship between writing data 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 format: This method allows you to allocate only file systems to partitions and rewrite the directory table. Quick format takes up less space.

- Full format: This method allows you to allocate files systems to partitions, rewrite the
  directory table, and scan for and mark damaged sectors. In addition, during the formatting
  process, empty data blocks on the disk are filled in, which is equivalent to writing 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 repartitioned the disk. How many partitions are available in the disk?

When you roll back a data disk by using a snapshot, the disk returns to the status it was in when the snapshot was taken. If the disk has not been repartitioned when the snapshot was taken, only one partition is available in the disk.

An error message similar to the following one is returned when I attempt to roll back a disk: "A disk can be rolled back only when the attached instance 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 replacement cannot be used to roll back the new system disk." What do I do?

- Problem description: When you attempt 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
  attached instance 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
  replacement cannot be used to roll back the new system disk."
- Cause: The problem may be caused by an incorrect disk property or disk state.
- Solution: You can troubleshoot the problem based on the instance status or snapshot status.
  - $\circ~$  Check whether the instance to which the disk is attached has been 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 status of the instance on the Instances page.
  - Check whether the system disk of the instance associated with the snapshot has been replaced.

If you selected a new image to replace the system disk, a new system disk is automatically recreated 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 system disk (non-public images).

• 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. You can go to the Instance Details page. Click Instance Snapshots in the left-side navigation pane. Check the status of snapshots. A snapshot is being created if the Progress value is not 100% and the Status value is Progressing.

If you need to forcibly terminate the creation process of a snapshot to roll back the disk, select the snapshot and click **Delete**.

#### How do I migrate data from the system disk of a Linux instance?

Assume that you purchase an ECS Linux instance without attaching any 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 solve this problem, 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 choose data copy methods based on the operating system:

- Copy data between Linux instances
  - Use the lrzsz tool

Log on to the Linux instances, install the lrzsz tool, run the rz command to upload files to one Linux instance, and then run the sz command to download the files to the other Linux instance.

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

• Use the FTP 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.

o Use 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 local PC and then use the FTP service to upload the files to the Windows instance.

- Copy data between Windows instances
  - Use the FTP service
  - Use TradeManager