Dedicated Distributed Storage Service

Service Overview

Issue 01

Date 2019-04-30

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to “Vul. Response Process”. For details about the policy, see the following website:https://www.huawei.com/en/psirt/vul-response-process
For enterprise customers who need to obtain vulnerability information, visit:https://securitybulletin.huawei.com/enterprise/en/security-advisory

Contents

1 What Is DSS?

Dedicated Distributed Storage Service (DSS) provides you with dedicated storage pools which are physically isolated from other pools to ensure high security. With data redundancy and cache acceleration technologies, DSS delivers highly reliable, durable, low-latency, and stable storage resources. By flexibly interconnecting with various compute services, such as Dedicated Computing Cluster (DCC), Elastic Cloud Server (ECS) and Bare Metal Server (BMS), DSS is suitable for different scenarios, including high performance computing (HPC), online analytical processing (OLAP), and mixed loads.

Figure 1-1 DSS architecture

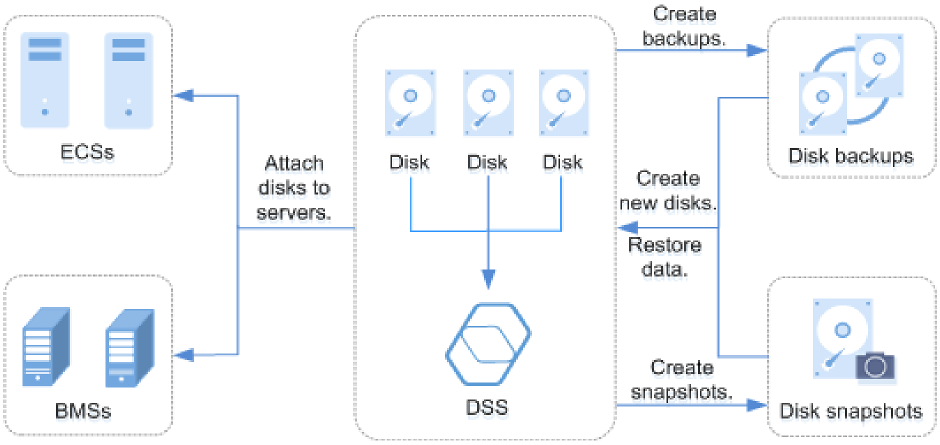

Advantages

- A variety of specifications
  - High I/O: Suitable for scenarios that require high performance, high read/write speed, and real-time data storage.
  - Ultra-high I/O: Excellent for read/write-intensive scenarios that require extremely high performance and read/write speed, and low latency.
- Elastic scalability
  - On-demand capacity improves resource utilization.
  - Linear performance increase can be achieved with capacity expansion.

- Security and reliability
  - Distributed storage with three data replicas ensures 99.9999999% durability.
  - System disks and data disks support data encryption with zero application awareness.
- Backup and restoration
  - Backups can be created for a DSS disk, and the backup data can be used to restore the disk data, maximizing data security and correctness and ensuring service security.

Differences Between DSS and EVS

Table 1-1 Differences between DSS and EVS

Service	Overall Introduction	Storage Category	Typical Application Scenarios	Performance
DSS	DSS provides exclusive physical storage resources for users. The storage pools are physically isolated, and data durability reaches 99.9999999%. Multiple types of compute services, including DCC, ECS and BMS, can be interconnected with DSS at the same time. DSS has abundant features to guarantee data security and reliability.	Dedicated storage pools, which means that storage pools are physically isolated and resources are exclusively used.	● Interconnection with compute services, such as ECS and BMS, in a dedicated cloud. ● Interconnection with compute services, such as ECS and BMS, in a non-dedicated cloud. ● Mixed load. DSS supports hybrid deployment of HPC, database, email, OA, and web applications. ● High-performance computing ● OLAP applications	● High I/O storage pool: The initial specification is 13.6 TB, which can be expanded to a maximum of 435.2 TB in 13.6 TB increments. The maximum IOPS is 1,500 IOPS/TB. ● Ultra-high I/O storage pool: The initial specification is 7.225 TB, which can be expanded to a maximum of 289 TB in 7.225 TB increments. The maximum IOPS is 8,000 IOPS/TB.

Service	Overall Introduction	Storage Category	Typical Application Scenarios	Performance
EVS	Elastic Volume Service (EVS) provides scalable block storage that features high reliability, high performance, and rich specifications for servers.	Shared storage pools	● Enterprise office applications ● Development and testing ● Enterprise applications, including SAP, Microsoft Exchange, and Microsoft SharePoint ● Distributed file systems ● Various databases, including MongoDB, Oracle, SQL Server, MySQL, and PostgreSQL	EVS disks start at 10 GB and can be expanded as required in 1 GB increments to a maximum of 32 TB.

Figure 1-2 Differences between DSS and EVS

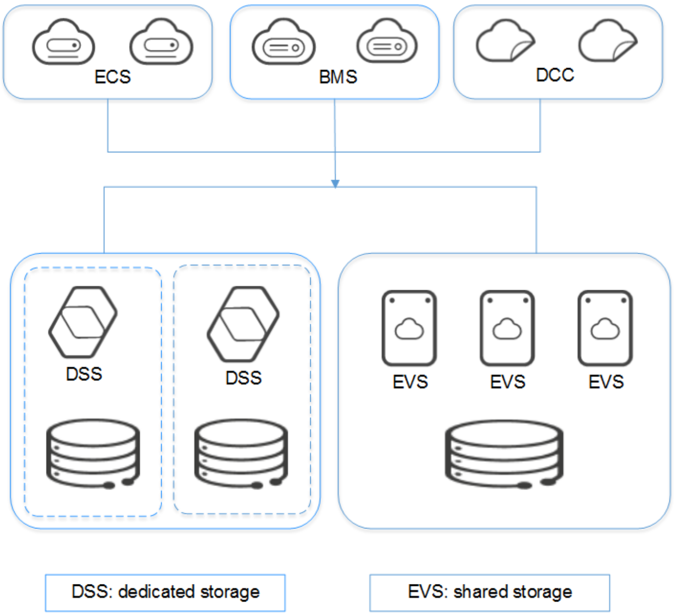

2 Region and AZ

Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified into universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides specific services for specific tenants.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers, to support cross-AZ high-availability systems.

Figure 2-1 shows the relationship between regions and AZs.

Figure 2-1 Regions and AZs

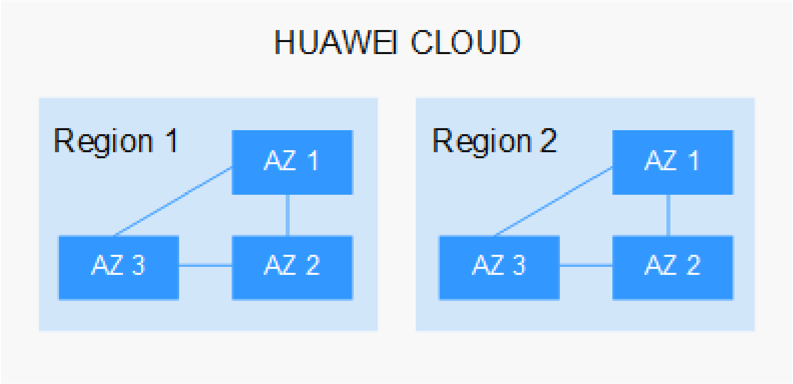

Huawei Cloud provides services in many regions around the world. You can select a region and an AZ based on requirements. For more information, see **Huawei Cloud Global Regions**.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  It is recommended that you select the closest region for lower network latency and quick access.

  – If your target users are in Asia Pacific (excluding the Chinese mainland), select the **CN-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.

  – If your target users are in Africa, select the **AF-Johannesburg** region.

  – If your target users are in Latin America, select the **LA-Santiago** region.

  NOTE

  The **LA-Santiago** region is located in Chile.

- Resource price

  Resource prices may vary in different regions. For details, see **Product Pricing Details**.

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For lower network latency, deploy resources in the same AZ.

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see **Regions and Endpoints**.

3 Storage Pool Types and Performance

DSS provides two types of storage pools, which differ in I/O performance and price. You can select the storage pool type based on your service requirements.

The disk type must be consistent with the storage pool type you selected.

Application Scenarios

- High I/O storage pool supports only high I/O disks. It can deliver a maximum of 1,500 IOPS per TB and a minimum of 1 to 3 ms read/write latency (single queue, 4 KiB data blocks). This type of storage pools is designed for mainstream high-performance, high-reliability applications, such as enterprise applications, large-scale development and testing, and web server logs.
- Ultra-high I/O storage pool supports only ultra-high I/O disks. It can deliver a maximum of 8,000 IOPS per TB and a minimum of 1 ms read/write latency (single queue, 4 KiB data blocks). This type of storage pools is perfect for read/write-intensive application scenarios, such as the distributed file systems in the HPC scenarios or NoSQL and relational databases in I/O-intensive scenarios.

Storage Pool Performance

Key metrics of the storage pool performance include read/write I/O latency, IOPS, and throughput.

- IOPS: Number of read/write operations performed per second
- Throughput: Amount of data read from and written into a storage pool per second
- Read/write I/O latency: Minimum interval between two consecutive read/ write operations

Table 3-1 Storage pool performance

Parameter	High I/O	Ultra-high I/O
IOPS	1,500 IOPS/TB	8,000 IOPS/TB

Parameter	High I/O	Ultra-high I/O
I/O read/write latency (single queue, 4 KiB data blocks)	1 ms to 3 ms	1 ms
Typical application scenarios	Common development and test environments	● Transcoding services ● I/O-intensive workloads – NoSQL – Oracle – SQL Server – PostgreSQL ● Latency-sensitive applications – Redis – Memcache

4 Storage Pool Capacity Description

Table 4-1 Storage pool capacity description

Type	Description
Requested Capacity	The capacity of the storage pool that you apply for.
Raw Capacity	The raw capacity of the storage pool that you apply for. The requested capacity of a storage pool is no less than 85% of its raw capacity.
Total Available Capacity	The total available capacity of a storage pool.
Allocated Capacity	The storage pool capacity that has been allocated. Includes the capacity allocated to: ● Volumes of VMs, bare metal servers, and containers ● Advanced services such as RDS ● Snapshots created during backup creation
Used Capacity	The storage pool physical capacity that has been used. Includes the capacity already used by: ● Volumes of VMs, bare metal servers, and containers ● Advanced services such as RDS ● Snapshots created during backup creation

Table 4-2 Storage pool capacity calculation example

Parameter	Capacity
Requested Capacity	27.2 TB
Raw Capacity	32 TB
Total Available Capacity	27.2 x 1024 GB = 27852 GB
Allocated Capacity	7330 GB
Used Capacity	432 GB

5 DSS Disks

DSS disks are essentially dedicated EVS disks, which can be used as scalable block storage for servers. With high reliability, high performance, and a variety of specifications, DSS disks can be used for distributed file systems, development and test environments, data warehouse applications, and HPC scenarios to meet diverse service requirements. Servers that DSS supports include ECSs and BMSs.

DSS disks are sometimes just referred to as disks in this document.

6 DSS Three-Copy Redundancy

What Is Three-Copy Redundancy?

The backend storage system of DSS employs three-copy redundancy to guarantee data reliability. With this mechanism, one piece of data is by default divided into multiple 1 MB data blocks. Each data block is saved in three copies, and these copies are stored on different nodes in the system according to the distributed algorithms.

DSS three-copy redundancy has the following characteristics:

- The storage system saves the data copies on different disks of different servers across cabinets, ensuring that services are not interrupted if a physical device fails.
- The storage system guarantees strong consistency between the data copies.

For example, for data block P1 on physical disk A of server A, the storage system backs up its data to P1'' on physical disk B of server B and to P1' on physical disk C of server C. Data blocks P1, P1', and P1'' are the three copies of the same data block. If physical disk A where P1 resides is faulty, P1' and P1'' can continue providing storage services, ensuring service continuity.

Figure 6-1 Three-copy redundancy

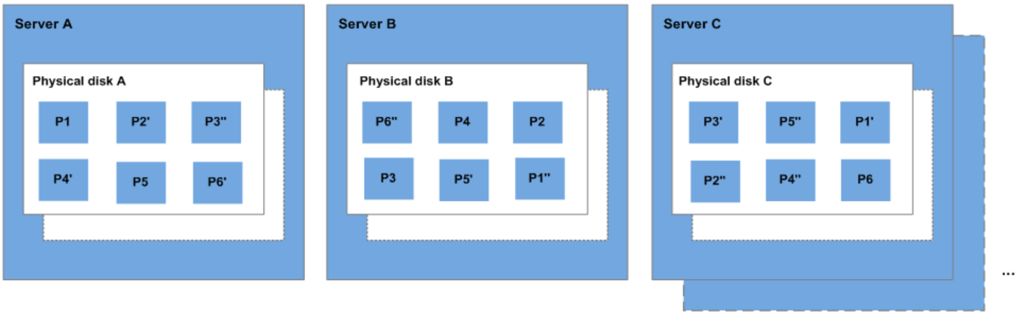

How Does Three-Copy Redundancy Keep Data Consistency?

Data consistency includes the following two aspects: When an application writes a piece of data to the system, the three copies of the data in the storage system must be consistent. When any of the three copies is read by the application later, the data on this copy is consistent with the data previously written to it.

DSS three-copy redundancy keeps data consistency in the following ways:

- Data is simultaneously written to the three copies of the data.

  When an application writes data, the storage system writes it to the three copies of the data simultaneously. In addition, the system returns the write success response to the application only after the data has been written to all of the three copies.

- Storage system automatically restores the damaged copy in the event of a data read failure.

  When an application fails to read data, the system automatically identifies the failure cause. If the data cannot be read from a physical disk sector, the system reads the data from another copy of the data on another node and writes it back to the original disk sector. This ensures the correct number of data copies and data consistency among data copies.

How Does Three-Copy Redundancy Rapidly Rebuild Data?

Each physical disk in the storage system stores multiple data blocks, whose copies are scattered on the nodes in the system according to certain distribution rules. When a physical server or disk fault is detected, the storage system automatically rebuilds the data. Since the copies of data blocks are scattered on different nodes, the storage system will start the data rebuild on multiple nodes simultaneously during a data restore, with only a small amount of data on each node. In this way, the system eliminates the potential performance bottlenecks that may occur when a large amount of data needs to be rebuilt on a single node, and therefore minimizes the adverse impacts exerted on upper-layer applications.

Figure 6-2 shows the data rebuild process.

Figure 6-2 Data rebuild process

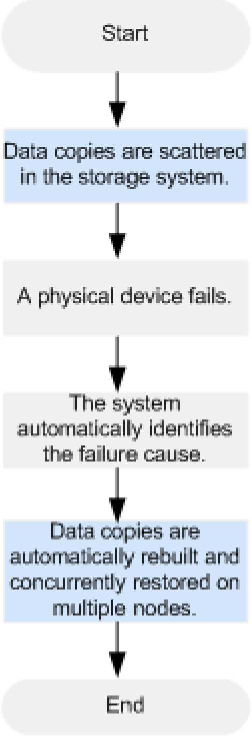
Start
Data copies are scattered in the storage system.
A physical device fails.
The system automatically identifies the failure cause.
Data copies are automatically rebuilt and concurrently restored on multiple nodes.
End

Figure 6-3 shows the data rebuild principle. For example, if physical disks on server F are faulty, the data blocks on these physical disks will be rebuilt on the physical disks of other servers.

Figure 6-3 Data rebuild principle

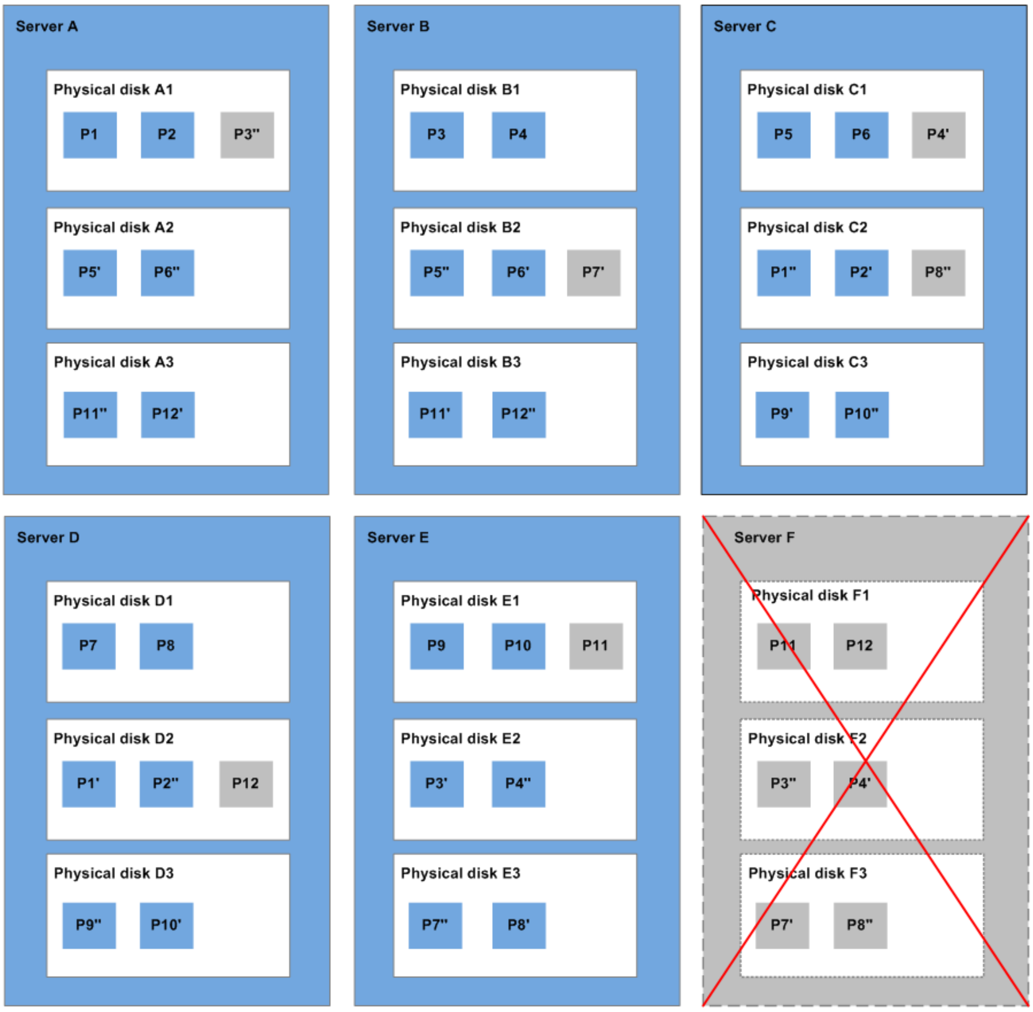

What Are the Differences Between Three-Copy Redundancy and Disk Backup?

Three-copy redundancy improves the reliability of the data stored on DSS disks. It is used to tackle data loss or inconsistency caused by physical device faults.

Whereas, backups are used to prevent data loss or inconsistency caused by misoperation, viruses, or hacker attacks. So you are advised to create backups to back up the DSS disk data on a timely basis.

7 Device Types and Usage Instructions

What Device Types Are Available?

There are two EVS device types: Virtual Block Device (VBD) and Small Computer System Interface (SCSI).

- VBD is the default EVS device type. VBD EVS disks support only basic read/write SCSI commands.
- SCSI EVS disks support transparent SCSI command transmission and allow the server OS to directly access the underlying storage media. Besides basic read/write SCSI commands, SCSI disks support advanced SCSI commands.

Device type is configured during purchase. It cannot be changed after the disk has been purchased.

Common Application Scenarios and Usage Instructions of SCSI EVS Disks

- BMSs support only SCSI EVS disks.
- Shared SCSI EVS disks: Shared SCSI EVS disks must be used together with a distributed file system or cluster software. Because most cluster applications, such as Windows MSCS, Veritas VCS, and Veritas CFS, require SCSI reservations, you are advised to use shared EVS disks with SCSI.

  SCSI reservations take effect only when shared SCSI EVS disks are attached to ECSs in the same ECS group. For more information about shared EVS disks, see **Shared Disks and Usage Instructions**.

Do I Need to Install a Driver for SCSI EVS Disks?

To use SCSI EVS disks, a cloud server must have a SCSI driver installed. If the SCSI driver is not pre-installed, you need to install it manually.

Check whether you need to manually install the driver based on the server type.

- Bare Metal Server (BMS)

  Both the Windows and Linux images for BMSs are pre-installed with the required SDI card driver. Therefore, no driver needs to be installed.
- KVM ECS

  You are advised to use SCSI EVS disks with KVM ECSs. Linux images and Windows images for KVM ECSs already have the required driver. Therefore, no driver needs to be installed for KVM ECSs.

NOTE

ECS virtualization types are categorized into KVM and Xen. For details, see **ECS Types**.

- Xen ECS

  Due to driver limitations, you are advised not to use SCSI EVS disk with Xen ECSs.

  However, a few images support SCSI EVS disks on Xen ECSs. For the supported images, see **Table 7-1**.

  NOTE

  After confirming that the OS images of Xen ECSs support SCSI EVS disks, determine whether you need to install the driver:

  - Public Windows images are preinstalled with the Paravirtual SCSI (PVSCSI) driver. Therefore, no driver needs to be installed.
  - Private Windows images are not preinstalled with the PVSCSI driver. You need to download and install it explicitly.

    For details, see **(Optional) Optimizing Windows Private Images** in the *Image Management Service User Guide*.
  - Linux images are not preinstalled with the PVSCSI driver. You need to obtain the source code of the open-source Linux driver at **https://github.com/UVP-Tools/SAP-HANA-Tools**.

Table 7-1 OSs supporting SCSI EVS disks

Virtualization Type	OS	
Xen	Windows	See the Windows images listed on the **Public Images** page. Log in to the management console, choose **Image Management Service**, click the **Public Images** tab, and select **ECS image** and **Windows** from the drop-down lists, respectively.
	Linux	● SUSE Linux Enterprise Server 11 SP4 64bit (The kernel version is 3.0.101-68-default or 3.0.101-80-default.) ● SUSE Linux Enterprise Server 12 64bit (The kernel version is 3.12.51-52.31-default.) ● SUSE Linux Enterprise Server 12 SP1 64bit (The kernel version is 3.12.67-60.64.24-default.) ● SUSE Linux Enterprise Server 12 SP2 64bit (The kernel version is 4.4.74-92.35.1-default.)

8 Shared Disks and Usage Instructions

DSS disks can be classified into non-shared disks and shared disks based on whether a disk can be attached to multiple servers. A non-shared disk can only be attached to one server, whereas a shared disk can be attached to multiple servers.

What Are Shared Disks?

Shared disks are block storage devices that support concurrent read/write operations and can be attached to multiple servers. Shared disks feature multiple attachments, high concurrency, high performance, and high reliability. A shared disk can be attached to a maximum of 16 servers. **Figure 8-1** shows its application scenario.

Currently, shared disks can be used as data disks only and cannot be used as system disks.

Figure 8-1 Application scenario of shared disks

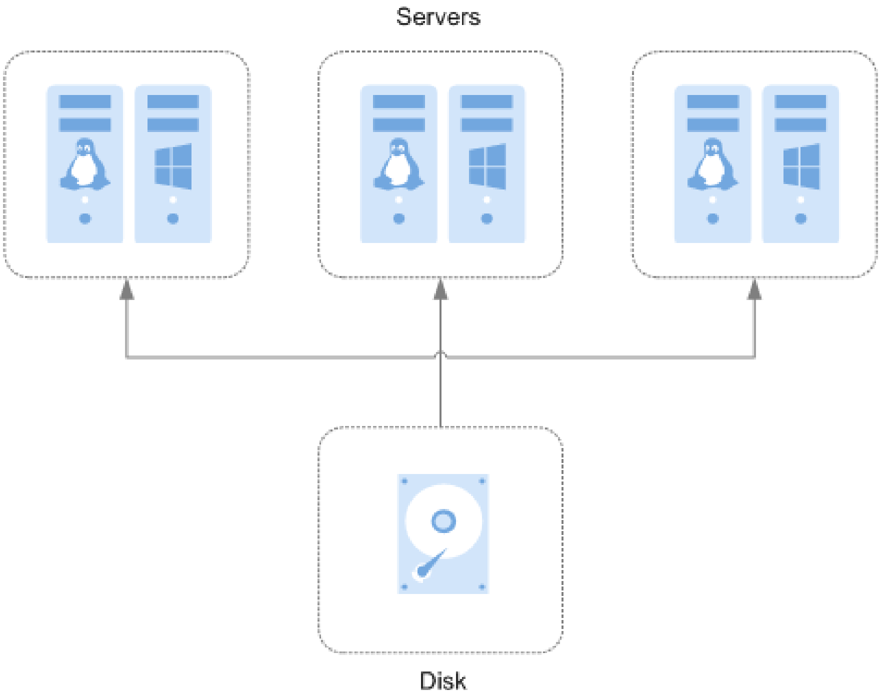

Application Scenarios and Precautions for Shared Disks

Shared disks are usually used for enterprise key applications that require cluster deployment and high availability (HA). These applications demand concurrent access to a disk from multiple servers. Before you attach a shared disk to multiple servers, the disk device type needs to be determined. The device type can be either VBD or SCSI.

Because most cluster applications, such as Windows MSCS, Veritas VCS, and Veritas CFS, require the usage of SCSI reservations, you are advised to use shared disks with SCSI. If a SCSI disk is attached to a Xen ECS for use, you must install the driver. For details, see **Device Types and Usage Instructions**.

You can create shared VBD disks or shared SCSI disks.

- Shared VBD disks: The device type of a newly created shared disk is VBD by default. Such disks can be used as virtual block storage devices, but do not support SCSI reservations. If SCSI reservations are required for your applications, create shared SCSI disks.
- Shared SCSI disks: These disks support SCSI reservations.

NOTICE

- To improve data security, you are advised to use SCSI reservations together with the anti-affinity policy of an ECS group. That said, ensure that the shared SCSI disk is only attached to ECSs in the same anti-affinity ECS group.
- If an ECS does not belong to any anti-affinity ECS group, you are advised not to attach shared SCSI disks to this ECS. Otherwise, SCSI reservations may not work properly, which may put your data at risk.

Concepts of the anti-affinity ECS group and SCSI reservations:

- The anti-affinity policy of an ECS group allows ECSs to be created on different physical servers to improve service reliability.

  For details about ECS groups, see **Managing ECS Groups**.
- The SCSI reservation mechanism uses a SCSI reservation command to perform SCSI reservation operations. If an ECS sends such a command to a disk, the disk is displayed as locked to other ECSs, preventing the data damage that may be caused by simultaneous read/write operations to the disk from multiple ECSs.
- ECS groups and SCSI reservations have the following relationship: A SCSI reservation on a single disk cannot differentiate multiple ECSs on the same physical host. For that reason, if multiple ECSs that use the same shared disk are running on the same physical host, SCSI reservations will not work properly. You are advised to use SCSI reservations only on ECSs that are in the same ECS group, thus having a working anti-affinity policy.

Advantages of Shared Disks

- Multiple attachments: A shared disk can be attached to a maximum of 16 servers.

- High-performance: When multiple servers concurrently access a shared ultra-high I/O disk, random read/write IOPS can reach up to 160,000.
- High-reliability: Shared disks support both manual and automatic backup, delivering highly reliable data storage.
- Wide application scenarios: Shared disks can be used for Linux RHCS clusters where only VBD disks are needed. Whereas, they can also be used for Windows MSCS and Veritas VCS clusters that require SCSI reservations.

Specifications of Shared Disks

Key metrics of the disk performance include read/write I/O latency, IOPS, and throughput.

- IOPS: Number of read/write operations performed by a disk per second
- Throughput: Amount of data read from and written into a disk per second
- Read/write I/O latency: Minimum interval between two consecutive read/write operations of a disk

  Single-queue access latencies (4 KiB data blocks) of different types of disks are as follows:

  – Common I/O: 5 ms to 10 ms
  – High I/O: 1 ms to 3 ms
  – Ultra-high I/O: 1 ms

Table 8-1 Disk performance data

Parameter	Common I/O	High I/O	Ultra-high I/O
Max. capacity	● System disk: 1,024 GB ● Data disk: 32,768 GB	● System disk: 1,024 GB ● Data disk: 32,768 GB	● System disk: 1,024 GB ● Data disk: 32,768 GB
Max. IOPS	2,200	5,000	33,000
Max. throughput	90 MB/s	150 MB/s	350 MB/s
Burst IOPS limit	2,200	5,000	16,000

Parameter	Common I/O	High I/O	Ultra-high I/O
Formula used to calculate disk IOPS **NOTE** Disk IOPS cannot exceed maximum IOPS. For example, the IOPS of an ultra-high I/O disk increases linearly as capacity grows (with a 50 IOPS increase for each GB added), but cannot exceed 33,000.	IOPS = Min. (2,200, 500 + 2 x Capacity)	IOPS = Min. (5,000, 1,200 + 6 x Capacity)	IOPS = Min. (33,000, 1,500 + 50 x Capacity)
API name **NOTE** This API name indicates the value of the **volume_type** parameter in the disk API. It does not represent the type of the underlying hardware devices.	SATA	SAS	SSD
Data durability	99.9999999%		
Number of servers that can be attached to	A shared disk can be attached to a maximum of 16 servers.		

NOTE

To test the performance of a shared disk, the following requirements must be met:

- The shared disk must be attached to multiple servers (ECSs or BMSs).
- If the shared disk is attached to multiple ECSs, these ECSs must belong to the same anti-affinity ECS group.

  If these ECSs fail to meet the anti-affinity requirement, the shared disk cannot reach the maximum performance.

Data Sharing Principle and Common Usage Mistakes of Shared Disks

A shared disk is essentially the disk that can be attached to multiple servers for use, which is similar to a physical disk in that the disk can be attached to multiple physical servers, and each server can read data from and write data into any space on the disk. If the data read/write rules, such as the read/write sequence and meaning, between these servers are not defined, data read/write interference between servers or other unpredictable errors may occur.

Though shared disks are block storage devices that provide shared access for servers, shared disks do not have the cluster management capability. You need to deploy a cluster system to manage shared disks. Common cluster management systems include Windows MSCS, Linux RHCS, Veritas VCS, and Veritas CFS.

If shared disks are not managed by a cluster system, the following issues may occur:

- Data inconsistency caused by read/write conflicts

  When a shared disk is attached to two servers (server A and server B), server A cannot recognize the disk spaces allocated to server B, vice versa. That said, a disk space allocated to server A may be already used by server B. In this case, repeated disk space allocation occurs, which leads to data errors.

  For example, a shared disk has been formatted into the ext3 file system and attached to server A and server B. Server A has written metadata into the file system in space R and space G. Then server B has written metadata into space E and space G. In this case, the data written into space G by server A will be replaced. When the metadata in space G is read, an error will occur.

- Data inconsistency caused by data caching

  When a shared disk is attached to two servers (server A and server B), the application on server A has read the data in space R and space G, then cached the data. At that time, other processes and threads on server A would then read this data directly from the cache. At the same time, if the application on server B has modified the data in space R and space G, the application on server A cannot detect this data change and still reads this data from the cache. As a result, the user cannot view the modified data on server A.

  For example, a shared disk has been formatted into the ext3 file system and attached to server A and server B. Both servers have cached the metadata in the file system. Then server A has created a new file (file F) on the shared disk, but server B cannot detect this modification and still reads data from its cached data. As a result, the user cannot view file F on server B.

Before you attach a shared disk to multiple servers, the disk device type needs to be determined. The device type can be either VBD or SCSI. Shared SCSI disks support SCSI reservations. Before using SCSI reservations, you need to install a driver in the server OS and ensure that the OS image is included in the compatibility list.

9 Disk Encryption

What Is Disk Encryption?

In case your services require encryption for the data stored on disks, EVS provides you with the encryption function. You can encrypt new disks. Keys used by encrypted disks are provided by the Key Management Service (KMS) of Data Encryption Workshop (DEW), which is secure and convenient. Therefore, you do not need to establish and maintain the key management infrastructure.

Keys Used for Disk Encryption

Keys provided by KMS include a Default Key and Custom Keys.

- Default Key: A key that is automatically created by EVS through KMS and named **evs/default**.

  It cannot be disabled and does not support scheduled deletion.
- Custom keys: Keys created by users. You can use existing keys or create new ones to encrypt disks. For details, see **Key Management Service** > **Creating a CMK** in the *Data Encryption Workshop User Guide*.

If you use a custom key to encrypt disks and this custom key is then disabled or scheduled for deletion, data cannot be read from or written to these disks or may never be restored. See **Table 9-1** for more information.

Table 9-1 Impact of custom key unavailability

Custom Key Status	Impact	How to Restore
Disabled	● For an encrypted disk already attached: Reads and writes to the disk are normal unless the disk is detached. Once detached, the disk cannot be attached again. ● For an encrypted disk not attached: The disk cannot be attached anymore.	Enable the custom key. For details, see **Enabling One or More Custom Keys**.
Scheduled deletion		Cancel the scheduled deletion for the custom key. For details, see **Canceling the Scheduled Deletion of One or More Custom Keys**.
Deleted		Data on the disks can never be restored.

NOTICE

You will be billed for the custom keys you use. If pay-per-use keys are used, ensure that you have sufficient account balance. If yearly/monthly keys are used, renew your order timely. Or, your services may be interrupted and data may never be restored as the encrypted disks become inaccessible.

Relationships Between Encrypted Disks and Backups

The encryption function can be used to encrypt system disks, data disks, and backups. The details are as follows:

- System disk encryption relies on images. For details, see the *Image Management Service User Guide*.
- The encryption attribute of an existing disk cannot be changed. You can create new disks and determine whether to encrypt the disks or not.
- When a disk is created from a backup, the encryption attribute of the new disk will be consistent with that of the backup's source disk.

Before you use the encryption function, EVS must be granted with the permission to access DEW. If you have the right to grant permissions, grant KMS access rights to EVS directly. If you do not have the permission, contact a user with the security administrator rights to add the security administrator rights for you. Then, grant KMS access rights to EVS. For details, see **Who Can Use the Encryption Feature?**

For how to create encrypted disks, see **Create a Disk**.

Who Can Use the Encryption Function?

- The security administrator (having Security Administrator permissions) can grant the KMS access rights to EVS for using the encryption function.

- When a user who does not have the Security Administrator permissions needs to use the encryption function, the condition varies depending on whether the user is the first one ever in the current region to use this function.
  - If the user is the first one ever in the current region to use this function, the user must contact a user having the Security Administrator permissions to grant the KMS access rights to EVS. Then, the user can use encryption.
  - If the user is not the first one ever in the current region to use this function, the user can use encryption directly.

From the perspective of a tenant, as long as the KMS access rights have been granted to EVS in a region, all the users in the same region can directly use the encryption function.

10 Disk Backup

What Is Disk Backup?

DSS implements the backup functions via Cloud Backup and Recovery (CBR). CBR allows you to create backups for disks on the console without stopping the server. If data is lost or damaged due to virus invasions, accidental deletions, or software/ hardware faults, you can use backups to restore data, guaranteeing your data integrity and security.

For more information, see the *Cloud Backup and Recovery Service User Guide*.

Backup Principles

See **CBR Service Overview** to learn about the backup principles.

Application Scenarios

Create and apply backup policies to schedule periodic backups for your disks. You can use the backup data to create new disks or restore to source disks.

Usage Instructions

For how to use disk backups, see **Cloud Backup and Recovery User Guide**.

11 DSS and Other Services

Figure 11-1 shows the related services.

Figure 11-1 DSS and other services

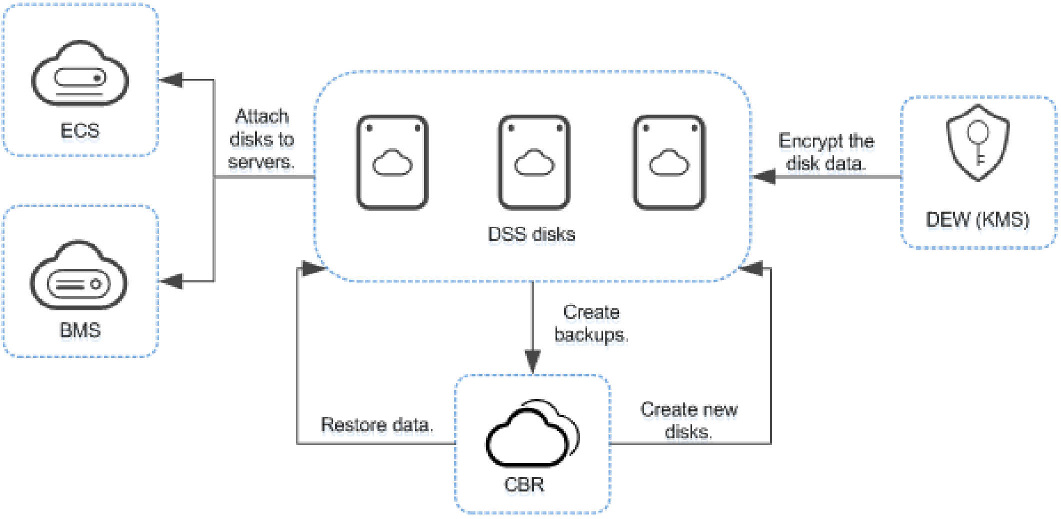

Table 11-1 Related services

Interactive Function	Related Service	Reference
Related services can use DSS disks and perform operations on the disks.	ECS	Purchasing an ECS Logging In to an ECS
	BMS	Creating a BMS and Attaching Data Disks
	CBR	Creating a Backup
	DEW	Creating a CMK See section "Creating a CMK" in the *Data Encryption Workshop User Guide*.

12 Billing

Billing Items

DSS is charged based on the storage pool type, capacity, and purchase quantity. For details, see **Pricing Details**.

Billing Modes

DSS supports yearly-based packages and does not support the pay-per-use billing mode.

Billing Involved in Configuration Modifications

- The storage pool and disk type cannot be changed.
- The capacity can be expanded only. For details, see **Expanding a Storage Pool** and **Expanding the Capacity of a Disk** .

13 Permissions

If you need to assign different permissions to employees in your enterprise to access your DSS resources, Identify and Access Management (IAM) is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you securely access your Huawei Cloud resources.

With IAM, you can use your Huawei Cloud account to create IAM users for your employees, and assign permissions to the users to control their access to specific resources. For example, some software developers in your enterprise need to use DSS resources but should not be allowed to delete them or perform any high-risk operations. In this scenario, you can create IAM users for the software developers and grant them only the permissions required for using DSS resources.

If your Huawei Cloud account does not require individual IAM users for permissions management, skip this section.

IAM can be used free of charge. You pay only for the resources in your account. For more information about IAM, see **IAM Service Overview**.

DSS Permissions

By default, new IAM users do not have permissions assigned. You need to add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services based on the permissions.

DSS is a project-level service deployed and accessed in specific physical regions. To assign DSS permissions to a user group, specify the scope as region-specific projects and select projects for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing DSS, the users need to switch to a region where they have been authorized to use this service. When accessing DSS, the users need to switch to a region where they have been authorized to use this service.

You can grant users permissions by using roles and policies.

- Roles: A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. This mechanism provides only a limited number of service-level roles for authorization. When using roles to grant permissions, you need to also assign other roles on which the permissions depend to take effect. However, roles are not an ideal choice for fine-grained authorization and secure access control.

- Policies: A type of fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control. For example, you can grant ECS users only the permissions for managing a certain type of ECSs. Most policies define permissions based on APIs. For the API actions supported by DSS, see **Permissions Policies and Supported Actions**.

Table 13-1 lists all the system-defined roles and policies supported by DSS.

Table 13-1 System-defined roles and policies supported by DSS

Role/Policy Name	Description	Type	Dependencies
DSS FullAccess	Full permissions for DSS. Users granted with this permission can create, expand, and query DSS resources.	System-defined policy	N/A
DSS ReadOnlyAccess	Read-only permission for DSS. Users granted with this permission can query DSS resources only.	System-defined policy	N/A

Table 13-2 lists the common operations supported by each system-defined policy or role of DSS. Select the policies or roles as required.

Table 13-2 Common operations supported by each system-defined policy or role of DSS

Operation	DSS FullAccess	DSS ReadOnlyAccess
Creating storage pools	√	×
Querying storage pools	√	√
Expanding storage pool capacities	√	×
Expanding disk capacity	√	×
Creating disks	√	×
Querying disks	√	√
Detaching disks	√	×
Deleting disks	√	×

Helpful Links

- **IAM Service Overview**
- **Creating a User and Granting DSS Permissions**
- **Permissions Policies and Supported Actions**

14 Constraints

This topic describes the constraints on using disks.

Table 14-1 Constraints on using disks

Scenario	Item	Restrictions
Creating disks	Device type	The device type of a disk cannot be changed after the disk has been created.
	Disk sharing	The sharing attribute of a disk cannot be changed after the disk has been created.
	Disk encryption	The encryption attribute of a disk cannot be changed after the disk has been created.
Attaching disks	Number of servers that a non-shared disk can be attached to	1
	Number of servers that a shared disk can be attached to	16
Expanding disk capacity	Capacity expansion	Disk capacities can be expanded only.
	Capacity expansion of non-shared disks	Some server OSs support the capacity expansion of non-shared, In-use disks.
	Capacity expansion of shared disks	A shared disk must be detached from all its servers before expansion. That is, the shared disk status must be **Available**.
	Expansion increment	1 GB
Detaching disks	System disk detachment	A system disk can only be detached offline, that is, its server must be in the **Stopped** state.

Scenario	Item	Restrictions
	Data disk detachment	A data disk can be detached online or offline, that is, its server can either be in the **Running** or **Stopped** state.
Deleting disks	-	● Only disks in the following statuses can be deleted: **Available**, **Error**, **Expansion failed**, or **Restoration failed**. ● Before you delete a shared disk, ensure that the disk has been detached from all its servers.
Disk capacity	Maximum capacity of a system disk	● High I/O: 1024 GB ● Ultra-high I/O: 1024 GB
	Maximum capacity of a data disk	● High I/O: 32768 GB ● Ultra-high I/O: 32768 GB
	Maximum capacity supported by the MBR partition style	2 TB
	Maximum capacity supported by the GPT partition style	18 EB

15 Change History

Released On	Description
2019-04-30	This issue is the first official release.