



Storage Resource Reutilization System

User Manual - Command Line

China Telecom Cloud Technology Co., Ltd.

Revision History

Version	Release Date	Description
4.0	March 24, 2026	<ol style="list-style-type: none"> 1. Support for the free edition. 2. Support for configuring server memory usage parameters for the HBlock service. 3. Support for suspending cloud LUNs. 4. Support for setting extended attributes of local LUNs via API. 5. Support for configuring authentication methods.
3.10	September 25, 2025	<ol style="list-style-type: none"> 1. Support target access permissions to implement permission management for clients and targets. 2. Support backup, generate data backup files independent of source LUNs based on snapshots. 3. Support for QoS policy to regulate traffic in terms of bandwidth and IOPS.
3.9	April 21, 2025	<ol style="list-style-type: none"> 1. Support for snapshot functionality to achieve rapid data backup and recovery. 2. Support for LUN cloning, for scenarios such as data replication and testing/validation.
3.8	February 14, 2025	<ol style="list-style-type: none"> 1. Update API signature method to enhance security. 2. Cloud LUNs support uploading data to S3-compatible object storage. 3. Support for setting the data storage directory of base services. 4. Add reclaim policy for targets, supporting automatic deletion after disassociation from LUNs. 5. Add performance parameter adjustment for scenarios such as intelligent computing, virtualization, and high-sensitivity scenarios. 6. Support for setting redundancy overlap, allowing data copies/fragments to be placed in the same fault domain.
3.7	August 8, 2024	<ol style="list-style-type: none"> 1. Support for configuring cluster topology. 2. Support for creating and managing multiple storage pools. 3. Support for room and rack level fault domains. 4. Support for setting cache pool for LUNs. 5. Support for base service migration.
3.6	June 3, 2024	<ol style="list-style-type: none"> 1. Support for uploading data from backend storage to OOS for storage LUNs and cache LUNs. 2. Hardware and HBlock monitoring data now support integration with Prometheus. 3. Optimize read/write performance. 4. Support for Loongson servers.
3.5	March 4, 2024	<ol style="list-style-type: none"> 1. Support for server and disk path level fault domains, and disk-level data services. 2. Support for specifying installation nodes for base services. 3. Support for capacity quota of the disk path, setting

		<p>the maximum amount of data that HBlock can write.</p> <ol style="list-style-type: none"> Support for targets to be discovered and connected by multiple clients. Support for setting the minimum number of write replicas for LUNs to improve data write security. Expand erasure coding (EC) N+M support range to satisfy $N+M \leq 128$.
3.4	July 12, 2023	<ol style="list-style-type: none"> LUN connections support one primary with multiple backups, improving business availability. Support for IPv6 environments. Provide a one-page Dashboard overview. Support for querying CHAP passwords via command line.
3.3	December 23, 2022	Support for safely removing servers.
3.2	September 26, 2022	<ol style="list-style-type: none"> Add "Basic Concepts" to the "Product Overview" chapter. Expand monitoring items to broaden coverage. Support for alarm and log management. Support for system events in the Events section. Add to the "Appendix" chapter: HBlock Service, System Event List, Monitoring Metrics, and Alarm List.
3.1	June 14, 2022	<ol style="list-style-type: none"> Support for multiple disk paths in standalone mode. Support for specifying servers when creating targets in cluster mode, and support target migration. Support for recording and querying user events.
3.0	January 18, 2022	<ol style="list-style-type: none"> Change command line to non-interactive mode. Support for WEB and API calling methods. LUN operations: Support for setting LUN high availability types and LUN write policies.
2.1	August 27, 2021	<ol style="list-style-type: none"> Support for ARM servers. Software license: Display allowed capacity when viewing licenses. LUN operations: Support for LUN active/standby switchover, i.e., switching between ACTIVE target and STANDBY target for LUNs.
2.0	May 28, 2021	<ol style="list-style-type: none"> Support for cluster deployment. Support for multi-replica and erasure coding (EC) data redundancy.

Contents

1 Product Overview	1
1.1 Product Definition	1
1.2 Application Scenarios	4
1.3 Basic Concepts	7
1.3.1 iSCSI	7
1.3.2 LUN	7
1.3.3 Snapshot	7
1.3.4 Clone LUN	8
1.3.5 Backup	8
1.3.6 iSCSI Target	8
1.3.7 iSCSI Target Portal	8
1.3.8 Storage Pool	9
1.3.9 QoS Policy	9
1.3.10 Cluster Topology	9
1.3.11 Fault Domain	9
1.3.12 Data Service	10
1.3.13 Monitoring	10
1.3.14 Event	10
1.3.15 Log Collection	10
1.3.16 Alarm	10
1.4 Usage Restrictions	12
1.5 Commercial Edition and Free Edition	13
1.6 Features Available After Software License Expiration (Trial or Subscription Mode)	16

1.7 Terms and Abbreviations	17
2 Server Side Deployment	19
2.1 Environmental Requirements	19
2.2 HBlock Deployment - Standalone Mode	21
2.2.1 Configuration Environment.....	21
2.2.2 Deploy HBlock.....	23
2.2.3 Example.....	26
2.3 HBlock Deployment - Cluster Mode	30
2.3.1 Configuration Environment.....	30
2.3.2 Deploy HBlock.....	33
2.3.3 Examples.....	36
3 Client Operations	41
3.1 Windows Client - Standalone Mode	41
3.2 Windows Client - Cluster Mode	46
3.3 Linux Client - Standalone Mode	50
3.3.1 Client Configuration.....	50
3.3.2 Examples.....	52
3.4 Linux Client - Cluster Mode	57
3.4.1 Client Configuration.....	57
3.4.2 Examples.....	61
4 Management Operations	69
4.1 Management Operation Command Line Format	69
4.2 Help Command	70
4.3 Install	73
4.4 Initialize	74
4.5 Software License	83
4.5.1 Import Software License.....	83

4.5.2 List Software Licenses.....	85
4.6 LUN Operations	88
4.6.1 Create a LUN.....	88
4.6.2 Create a Clone LUN.....	103
4.6.3 Flatten the Clone LUN.....	111
4.6.4 Delete a LUN.....	112
4.6.5 Expand a LUN.....	113
4.6.6 Edit a LUN.....	117
4.6.7 Set LUN Primary–Secondary Priority or Automatic Failover (Cluster Mode) .	121
4.6.8 Trigger Active/Standby Switchover of the Target Corresponding to the LUN (Cluster Mode).....	124
4.6.9 Recovering a LUN (Cache/Storage Mode LUN)	125
4.6.10 Suspend a LUN (Cache/Storage Mode LUN)	133
4.6.11 Resume a LUN from Restore Interruption or Suspension (Cache/Storage Mode LUN).....	135
4.6.12 Wipe a LUN (Local Mode LUN)	136
4.6.13 Query LUN Information.....	137
4.7 Snapshot Operations	152
4.7.1 Create a Snapshot.....	152
4.7.2 Modify a Snapshot.....	154
4.7.3 Roll Back a Snapshot.....	156
4.7.4 Delete a Snapshot.....	158
4.7.5 Query Snapshot Information.....	159
4.8 Consistency Snapshot Operation	162
4.8.1 Create a Consistency Snapshot.....	162

4.8.2	Modify a Consistency Snapshot.....	164
4.8.3	Roll Back a Consistency Snapshot.....	166
4.8.4	Delete a Consistency Snapshot.....	168
4.8.5	Query Consistency Snapshot.....	169
4.9	Backup Operation (Local LUN)	171
4.9.1	Export Backup File.....	171
4.9.2	Import Backup File.....	173
4.10	iSCSI Target Operations	175
4.10.1	Create an iSCSI Target.....	175
4.10.2	Delete an iSCSI Target.....	179
4.10.3	Set CHAP Authentication of the iSCSI Target.....	180
4.10.4	Migrate the iSCSI Target (Cluster Mode)	181
4.10.5	Edit Maximum Session Number per IQN Under the iSCSI Target.....	183
4.10.6	Modify the Remain Policy of the iSCSI Target.....	184
4.10.7	Set iSCSI Target AllowList.....	185
4.10.8	Delete an iSCSI Target Allowlist.....	188
4.10.9	Query Target Information.....	189
4.11	Storage Pool Operations (Cluster Mode)	193
4.11.1	Create a Storage Pool.....	193
4.11.2	Add Nodes to a Storage Pool	200
4.11.3	Modify a Storage Pool.....	202
4.11.4	Remove a Node from a Storage Pool.....	203
4.11.5	Delete a Non-base Storage Pool.....	204
4.11.6	Query Storage Pool Information.....	205

4.12 QoS Policy	211
4.12.1 Create a QoS Policy	211
4.12.2 Modify a QoS Policy	217
4.12.3 Associate the QoS Policy with Objects	223
4.12.4 Disassociate the QoS Policy from Objects	225
4.12.5 Delete a QoS Policy	227
4.12.6 Query QoS Policy Information	228
4.13 Cluster Topology (Cluster Mode)	233
4.13.1 Create a Topology Node	233
4.13.2 Modify Topology Node Information	234
4.13.3 Delete a Topology Node	235
4.13.4 Query Topology Information	237
4.14 Server Operations	239
4.14.1 Add a Server (Cluster Mode)	239
4.14.2 Modify Server Port Range	242
4.14.3 Set Server Target Portal IP	244
4.14.4 Set the Memory Available for HBlock on the Server	245
4.14.5 Set Server Default Disk Path (Standalone Mode)	247
4.14.6 Migrate Base Services on the Server (Cluster Mode)	248
4.14.7 Remove a Server (Cluster Mode)	250
4.14.8 Query Server Information	252
4.14.9 Add Disk Paths	262
4.14.10 Modify Capacity Quota of Disk Path	264
4.14.11 Remove a Disk Path	265

4.15 Query HBlock Information	266
4.16 Query HBlock Service Status	269
4.17 Monitoring	272
4.17.1 View Real-time Performance Data.....	272
4.17.2 Export Performance Data.....	275
4.18 Alarm	277
4.18.1 Query Alarm Information.....	277
4.18.2 Export Alarms.....	281
4.18.3 Manually Resolve the Alarm.....	282
4.18.4 Mute Alarm.....	283
4.18.5 Unmute Alarm.....	284
4.19 Event and Log Management	285
4.19.1 Query HBlock Events.....	285
4.19.2 Export HBlock Events.....	288
4.19.3 Initiate HBlock Log Collection.....	289
4.19.4 Query Collected Logs of HBlock.....	291
4.19.5 Delete Collection Logs.....	293
4.20 HBlock System Settings	294
4.20.1 Change Administrator Password.....	294
4.20.2 Email Settings.....	295
4.20.3 Remote Assistance Settings.....	299
4.20.4 Pushgateway Monitoring Configuration.....	302
4.20.5 Authentication Mode.....	310
4.20.6 Enable Pro Trial (Free Edition).....	312
4.21 Tune the Performance Parameters of HBlock (Cluster Mode).....	313
4.21.1 Tune the Performance Parameters of HBlock.....	313

4.21.2 View the Performance Tuning Configuration.....	314
4.22 Stop HBlock on the Server	315
4.23 Start HBlock on the Server	317
4.24 Restart HBlock on the Server	318
4.25 Uninstall HBlock	320
4.26 Query HBlock Version	321
4.27 Upgrade HBlock	322
4.27.1 Upgrade HBlock.....	322
4.27.2 Query the Upgrade Status.....	326
5 FAQs	328
5.1 Deployment Environmental Requirements	328
5.2 Common Operations	330
5.3 Upgrade and Expand Capacity	344
5.4 Data Migration	345
5.5 Data Security	346
5.6 Technical Support	347
5.7 Purchase Guide	348
5.8 Other FAQs	349
6 Troubleshooting	351
7 Appendix	352
7.1 HBlock Service	352
7.2 User Event List	353
7.3 System Event List	356
7.4 Monitoring Metrics	359
7.5 Alarm List	365
7.6 OOS Endpoint and Region	371
7.6.1 Object Storage Network.....	371
7.6.2 Object Storage Network 2.....	371
7.6.3 Hongkong Node.....	372
7.6.4 Other Region.....	372

7.7 Operating System Monitoring Metrics Which HBlock can Push.....	373
7.8 Cluster Topology File	377
7.9 iSCSI Target Allowlist File	381

1 Product Overview

1.1 Product Definition

HBlock is a Storage Resource Reutilization System (SRRS) independently developed by eSurfing Cloud. It is a lightweight storage cluster controller and a fully user-mode software-defined storage controller that converts commodity servers and managed storage resources into HA virtual storage disks. It provides distributed block storage services through the standard Internet Small Computer System Interface (iSCSI) protocol and the disks can be mounted to a local server (or other remote servers) for intensive use of resources. Meanwhile, HBlock is compatible with heterogeneous devices and applies to various scenarios. HBlock helps users tackle the challenges of IT architecture upgrade, reduce costs, and increase efficiency, and assists enterprise users in achieving green transformation.

In non-networked mode, HBlock can be seen as a substitute for local disk arrays for local data storage. In networking mode, HBlock can also serve as a bridge between local and cloud storage, automatically synchronizing all data to S3-compatible object storage, retaining only hot data locally to save local storage space, or retaining all data to ensure local I/O performance, achieving hybrid cloud storage.

HBlock can be installed on Linux operating systems as a common application without root permissions. You can deploy HBlock with other applications in a server in a hybrid manner to form high-performance and HA virtual hard disks for business. This way, HBlock directly re-uses storage resources without affecting user business or purchasing additional equipment!

The traditional hardware storage array provides low latency and high availability for each logical volume, but has poor horizontal scalability, high cost, and may form many "isolated data islands". That leads to high cost and low utilization of storage resource. Traditional distributed storage, while attractive, often suffers from issues such as complex deployment, poor performance, and low stability.

HBlock delivers storage arrays in a quite different approach:

- **Easy Installation:** The HBlock installation package is a zip file that can be installed on mainstream Linux operating systems running on commonly used 64-bit x86 servers, ARM servers or Loongson servers. It also supports physical servers, bare metal servers, and virtual machines. HBlock is completely decoupled with hardware drivers, so users can freely use the latest hardware on the market, with less vendor lock-in.

- **Green:** HBlock runs as a group of user-mode processes, without relying on any specific version of Linux kernel or distributions. It does not rely on or modify the operating system environment, neither does it monopolize the entire hard drive or interfere with the execution of any other processes. Therefore, HBlock can run in the same Linux operating system instance concurrently with other applications. On the one hand, it can help users improve the utilization of existing hardware resource, on the other hand, it also lowers the barriers for potential users to try HBlock - even a virtual machine is not needed.
- **High Utilization:** HBlock supports heterogeneous hardware. HBlock allows each Linux operating system instance in the cluster to have different hardware configurations, such as different numbers of CPUs, different sizes of memory, and different capacities of local hard drives. Therefore, it can improve the utilization of existing hardware resources.
- **High Performance:** HBlock adopts a distributed multi-controller architecture to provide the low latency and high availability just like traditional hardware storage arrays, as well as high scalability and high throughput just like traditional distributed storage. It can scale up from three servers to thousands of servers or shrink from thousands of servers to three servers one by one without service unavailability.
- **High Quality:** HBlock is designed to guarantee the data durability as long as the number of simultaneous disk faults in the cluster is not greater than the allowed faulty number of logical volumes in redundancy mode (for three-copy mode, the allowed number is two; for Erasure Code N+M mode, the allowed number is M). HBlock can ensure the service availability when any single server, link, or disk in the cluster experiences faults. HBlock is specifically designed for Chaos environment and can be applied to various complicated environments such as unstable networks, unstable computing power, and unstable disks. It has been sufficiently tested in complex and large-scale environments before release. HBlock supports snapshot and clone functions. Users can create time-point snapshots for logical volumes to quickly back up and restore data. The clone function enables quick generation of new volume copies based on snapshots, greatly improving the efficiency of data replication and test environment setup.

Meanwhile, as a bridge between local and cloud storage, HBlock also has the following features:

- **Data Cloud Migration:** HBlock can be integrated with S3-compatible object storage to create local LUNs (local mode) and cloud LUNs (cache mode and storage mode):
 - **Local mode LUNs:** All data is stored locally, which not only has the low latency and high availability of traditional hardware storage arrays, but also combines the high scalability and throughput of traditional distributed storage. The green feature greatly reduces the

cost of user deployment, while the heterogeneous hardware feature improves the utilization of existing hardware resources.

- Storage mode LUNs: All data is not only stored locally, but also asynchronously replicated to object storage, achieving high local performance and remote data disaster recovery.
- Cache mode LUNs: Recently read and written data will be cached locally to maximize performance, while the full amount of data will be saved on object storage to reduce costs, allowing a small local capacity to store massive amounts of data. This is particularly suitable for businesses that do not require high real-time performance, such as data backup and archiving, as well as businesses that require more writing and less access, such as document files, medical imaging, video surveillance, etc. HBlock can seamlessly connect local applications with cloud storage, enabling on-demand use of storage space and elastic expansion.
- **Consistency:** HBlock utilizes the atomic operations of object storage, which can truly ensure the consistency of data on the cloud (i.e., data on the cloud is always a snapshot of local data), and there will be no situation where the entire business cannot be restored due to inconsistency of data on the cloud, thus ensuring data security.

1.2 Application Scenarios

- **Utilization of Existing Resources**

HBlock features broad compatibility and uniformly manages and integrates idle storage space in various servers into storage pools, and provides high-performance and HA virtual disks to other hosts through the iSCSI protocol. To meet the increasing demand for storage capacity raised by rapid service growth and solve the waste of resources due to idle server resources, HBlock provides rapid deployment and expansion solutions that improve the usage of storage resources without extra cost, and support rolling service updates. This way, HBlock meets the ever-changing requirements for capacity and performance of future business operations.

- **HA Storage for Application Systems**

HBlock manages the physical disks of application nodes and mounts the HA virtual disks to the application nodes. It makes it easier for applications to achieve HA and re-utilize the storage resources of application nodes. Furthermore, no additional procurement for the storage hardware is required, which reduces the total cost of ownership (TCO) for users.

- **Independent Management and Control of New Resources**

When using HBlock to manage the storage resource pools that you have created, you will be able to have actual control over storage servers. This means you can use HBlock to manage storage resources and deploy other applications on your hardware to give full play to the hardware. Traditional integrated software and hardware storage products or distributed storage solutions require exclusive access to devices, and users can only perform limited operations on the management console. The storage cluster managed by HBlock enables users to comprehensively manage and control resource pools and enhances operating flexibility. You can select servers with any specifications for resource pool upgrading or expansion without limitations of vendor lock-in. You can also flexibly choose suitable hardware based on your business needs and budgets, improving the cost-effectiveness of investment.

- **Hybrid Cloud Storage**

HBlock can manage both local storage resources and OOS storage resources simultaneously, achieving unified management of storage space and meeting customers' needs for hybrid cloud storage. For customers who need to store massive amounts of data, HBlock can seamlessly connect local applications with cloud storage, synchronize

data to the cloud, and achieve on-demand use of storage space for elastic expansion. For scenarios where high data security is required and sensitive data is not suitable for cloud storage, HBlock can also help users achieve local data storage and improve data access speed. In addition, HBlock simplifies data management in hybrid cloud storage environments. Providing virtual targets and logical volumes for upper layer applications through the standard SCSI protocol, it can be deployed not only locally but also on private or public clouds.

- **Rapid Construction of a Production-like Environment**

Based on its distributed block storage-derived multi-LUNs consistent snapshot capability, HBlock can instantly capture the data states of multiple applications in the upper-layer business system. When data is accidentally deleted or logical errors occur, snapshots enable second-level recovery, significantly reducing the risk of data loss. This makes it an ideal choice for data backup, recovery, and version management. Additionally, leveraging the capability to create clone LUNs in seconds, HBlock can rapidly replicate multiple independent test environments that are consistent with the production environment. This greatly improves the efficiency of data replication and environment setup, facilitating the verification of rapid application recovery and disaster recovery capabilities. Meanwhile, it enables efficient business testing, upgrade validation, and data analysis without disrupting the production environment.

- **Geo-Diverse Multi-Active High-Reliability Storage**

HBlock optimizes the storage protocol stack and innovates in distributed algorithms to ensure cross-AZ strong data consistency (RPO=0) and achieve second-level fault recovery (RTO in seconds). The system can automatically detect faults, triggering cross-AZ data reconstruction. It uses intelligent data redundancy and redundancy overlap to ensure data security while reducing storage costs. In case of an AZ failure, services remain available with seamless business failover, requiring no manual intervention. This guarantees data security and business continuity.

- **Permanent Independent Backup of Critical Data**

HBlock supports full and incremental backups via snapshot technology, guaranteeing long-term retention and complete isolation from production systems. Backup files can be flexibly exported to various storage media (e.g., object storage, NAS), so data is fully decoupled from the platform. Even if the primary system is totally destroyed or hit by ransomware, these standalone backups remain a trusted source for precise LUN-level recovery, ensuring zero loss of business-critical data. The solution is ideal for finance,

healthcare, and other industries with strict long-term archival and compliance-audit requirements, providing rock-solid data protection.

- **Zero-Trust Storage Security Framework**

HBlock builds an end-to-end zero-trust storage access fabric through a triple lock: iSCSI target allowlist, CHAP authentication, and QoS flow control.

- At LUN granularity, only clients whose IQN is on the allowlist can discover or connect—unauthorized hosts remain completely invisible.
- CHAP handshakes block spoofed identities and man-in-the-middle attempts before a single I/O is issued.
- Coupled with QoS policies that dynamically throttle or block abnormal traffic in real time across both IOPS and bandwidth dimensions.

This delivers a triple guarantee: zero impact on storage performance, airtight isolation of data and traffic, and uninterrupted operation of mission-critical workloads.

1.3 Basic Concepts

1.3.1 iSCSI

iSCSI (Internet Small Computer System Interface) is a storage technology based on TCP/IP and SCSI-3 protocols, used to establish and manage IP storage devices, interconnections between hosts and clients, create a storage area network (SAN), etc.

1.3.2 LUN

LUN is the abbreviation of the term "Logical Unit Number," In the field of storage, LUN is a global identifier that uniquely identifies a storage volume or partition within a storage device (such as a hard disk or array). Through LUN, the operating system can recognize a specific storage volume and perform data read and write operations. You can create local LUNs (local mode) and cloud LUNs (cache mode and storage mode) according to your needs:

- Local mode: Stores all data locally.
- Cache mode: Stores part of hot data locally and stores all data in cloud asynchronously.
- Storage mode: Stores all data locally and stores them in cloud asynchronously.

`: For a LUN in replica mode, assuming that the number of LUN replicas is X and the minimum replica number is Y ($Y \leq X$), each time the data is written to LUN, at least Y replicas of data must be written successfully before this write request is considered successful. For a LUN in EC $N + M$ mode, assuming that the minimum replica number of the LUN is set to Y ($N \leq Y \leq N + M$), the data blocks and parity blocks that sum to at least Y blocks must be written successfully before this write request is considered successful.

1.3.3 Snapshot

A snapshot records an HBlock LUN's state and contents at a specific time point. Like a photo, it captures and saves the complete data status of the LUN at that moment. When needed, users can roll back via snapshots or create clone LUNs based on them to quickly restore to that state, effectively protecting data.

A consistency snapshot involves creating snapshots for all selected LUNs simultaneously at a specific moment, ensuring the snapshots reflect the data status of the same point in time.

1.3.4 Clone LUN

A clone LUN, based on a snapshot, mirrors the data and state at the snapshot's creation. It can be modified independently of the source LUN, tracking only new or changed data post - cloning, while relying on the source for original data. Once remove the reference to the parent snapshot from the clone LUN, it will copy the data stored in the snapshot to the clone. After the operation completes, the clone LUN will become an independent LUN.

Cloning depth refers to the number of layers in a chain of clone LUNs from a single snapshot. It indicates how many layers of clone LUNs can be created consecutively starting from the initial source LUN. For example:

- Create snapshot snap1 from source LUN lun1, then clone lun1-C1 from snap1. Cloning depth is 1.
- Create snapshot snap2 from lun1-C1, then clone lun1-C2 from snap2. Cloning depth becomes 2.

1.3.5 Backup

Captures the data state at a specific point-in-time and produces self-contained backup files that can be stored independently and imported on demand. Supports both full and incremental exports, and allows importing intact backups as well as those interrupted mid-process.

1.3.6 iSCSI Target

iSCSI target is a storage resource located on an iSCSI server. It is a protocol that connects data storage devices through an IP network infrastructure, enabling the mapping of remote storage devices to local hosts. This provides a network-based storage solution.

1.3.7 iSCSI Target Portal

The iSCSI target portal refers to the target portal of the HBlock server, used for communication with initiators that are not on the same LAN of HBlock.

If an Initiator outside the same LAN of HBlock wants to connect with HBlock server, network configuration (such as NAT) is required to ensure the Initiator can access the HBlock server through the IP address. Then, configure this address as the target portal of the HBlock server. Afterward, the Initiator can connect with the HBlock server through the configured target portal.

1.3.8 Storage Pool

Storage pool: A collection of storage resources provided by hardware. Physically, it refers to a collection of hard disks of the same medium across multiple servers. You can also form a storage pool as needed based on disk paths on servers in different server rooms and cabinets.

Base storage pool: The default storage pool created during initialization is the base storage pool. Physical resources upgraded from versions earlier than 3.7 belong to the base storage pool.

1.3.9 QoS Policy

Quality of Service (QoS) policies precisely regulate IOPS and throughput through a token-bucket algorithm, shaping traffic before congestion occurs and thereby preventing network bottlenecks at the source.

The policy can be applied to the following scenarios as required:

- Critical-business protection: By throttling the I/O of adjacent workloads, it indirectly safeguards the service quality of critical applications.
- Storage-tier optimization: Intelligently distributes I/O resources across performance tiers (e.g., SSD and HDD) to optimize performance in hybrid-storage environments.
- Burst-traffic handling: It allows short-term exceedance of baseline limits, smoothing sudden I/O bursts instead of rejecting requests outright.

1.3.10 Cluster Topology

Cluster topology: Displays the deployment of physical resources in a cluster logically and visually.

1.3.11 Fault Domain

You can set a fault domain for a storage pool. A fault domain is a collection of components that may fail simultaneously due to infrastructure sharing, such as rooms, racks, servers and disk paths. A fault of a component in the same fault domain affects only the data in the local fault domain.

HBlock allocates and stores data replicas in replica mode or data blocks in EC mode by fault domains to guarantee data security. Replicas of the same data and EC blocks of the same data are written into different fault domains.

1.3.12 Data Service

One disk path corresponds to one data service, and HBlock manages the user's file data blocks in the disk path through the data service.

1.3.13 Monitoring

Monitoring refers to monitoring and recording the performance indicators of the HBlock system, storage pool, servers, disk paths, and LUNs. Users can view real-time or historical performance data and pay attention to the performance of storage services. Please see the appendix **Monitoring Metrics**.

1.3.14 Event

Events refer to the user operations on HBlock or HBlock system behaviors recorded by the system, which reflect the comprehensive storage status and facilitate troubleshooting, auditing, and tracking.

Events are divided into user events and system events:

- User events: User operations on HBlock. See the appendix **User Event List** for a detailed list of user events.
- System events: HBlock system behaviors. See the appendix **System Event List** for a detailed list of system events.

1.3.15 Log Collection

Log collection means that users actively collect HBlock log data, generate log files, and download them locally, send to HBlock engineers to view the log details for troubleshooting. You can narrow the scope of log collection and speed up the collection progress by specifying the time period, log type, and server.

1.3.16 Alarm

Alarms refer to information generated when the system detects HBlock business or system abnormality.

Alarms are divided into three levels:

- Warning: Refers to a general situation where the system detects a potential or imminent fault that affects the business. There are currently no alarms that affect the business.

Maintenance personnel needs to find the cause of the alarm in time and eliminate potential faults.

- Major: Refers to an alarm within a local scope that affects system performance. Prompt action is required; otherwise, it will impact the operation of critical functions.
- Critical: Refers to a global alarm that has already caused business interruption or paralysis. It needs to be handled immediately, otherwise the system is in danger of collapsing.

For more information, see the appendix **Alarm List**.

1.4 Usage Restrictions

Item	Description
iSCSI target	HBlock supports a maximum of 32766 target IQNs. A target can be associated with up to 256 LUNs, but each LUN can only be associated with one target.
Storage Pool	A maximum of 32768 storage pools can be created in an HBlock cluster.
Snapshot	<ul style="list-style-type: none">● Maximum snapshots per LUN: 512.● Maximum snapshots per system: 100,000.● Maximum clone LUNs per snapshot: 512.● Maximum snapshot depth per system: 512.
Consistency Snapshot	Maximum number of LUNs for a consistency snapshot: 512.
Clone	<ul style="list-style-type: none">● Maximum number of clone LUNs supported by the system: 100,000.● Maximum number of clone LUNs per snapshot: 512.● Maximum clone depth supported by the system: 16.
Server	Each HBlock server can only add up to 100 disk paths.

1.5 Commercial Edition and Free Edition

During initialization, you may select either the Commercial Edition or the Free Edition based on your requirements. The features supported by each edition are as follows:

Features		Commercial Edition	Free Edition
Software License	import software license, list software licenses	Supported	Supported
LUN	create a LUN, expand a LUN, edit a LUN, set LUN primary-secondary priority or automatic failover, trigger Active/Standby switchover of the target corresponding to the LUN, set LUN extended attributes, delete specified extended attributes of LUN, delete all extended attributes of LUN, query LUN extended attributes, delete a LUN, query LUN information	Supported	Supported
	create a clone LUN, flatten the clone LUN, recover a LUN, wipe a LUN, resume LUN recovery	Supported	Not supported
Snapshot	create a snapshot, modify a snapshot, roll back a snapshot, delete a snapshot, query snapshot information	Supported	Not supported
Consistency Snapshot	create a consistency snapshot, modify a consistency snapshot, roll back a consistency snapshot, delete a consistency snapshot, query consistency snapshot information	Supported	Not supported
Backup	export backup file, import backup file	Supported	Not supported
iSCSI target	create an iSCSI target, delete an iSCSI target, set CHAP authentication of the iSCSI target, delete CHAP, migrate an iSCSI target, edit maximum sessions number per IQN under iSCSI target, modify the remain policy of the iSCSI target, set iSCSI target allowlist, delete an iSCSI target allowlist, query iSCSI target information, query iSCSI target connection, delete iSCSI target connection	Supported	Supported
Storage Pool	create a storage pool, add nodes to the storage pool, modify a storage	Supported	Supported

	pool, remove nodes from a storage pool, delete a non-base storage pool, query the storage pool information, query the QoS policy associated with the storage pool		
QoS Policy	create a QoS policy, modify a QoS policy, associate the QoS policy with LUNs, disassociate the QoS policy from LUNs, associate the QoS policy with storage pools, disassociate the QoS policy from storage pools, set the default QoS policy for LUNs in the storage pool, disassociate the default QoS policy for LUNs in the storage pool, delete a QoS policy, query QoS policy information, query information on objects associated or associable with the QoS policy	Supported	Not supported
Cluster Topology	create a topology node, modify topology node information, delete a topology node, query topology information	Supported	Supported
Server	add a server, set server properties (including: modify server port range, set server target portal IP, set the memory available for HBlock on the server, and set server default disk path), delete server properties (server target portal IPs), remove a server, query server information, add disk paths, modify capacity quota of a disk path, remove a disk path	Supported	Supported
	migrate the base services on the server	Supported	Not supported
Monitoring	query real-time performance data, export performance data (command line), obtain historical performance data (WEB/API)	Supported	Supported
Alarm	query HBlock alarm information, export HBlock alarms, manually resolve HBlock alarm, mute HBlock alarm, unmute HBlock alarm	Supported	Supported
Event	query HBlock events, export HBlock events	Supported	Supported

Log	initiate HBlock log collection, query collected logs, download the collected log file, delete collected logs	Supported	Supported
Administrator Password	change the administrator password	Supported	Supported
Email Settings	set the email, send a test email, delete email configuration, query email configuration	Supported	Supported
Remote Assistance Configuration	set remote assistance configuration, delete remote assistance configuration, query remote assistance configuration	Supported	Supported
Pushgateway Monitoring Configuration	add pushgateway monitoring configuration, modify pushgateway monitoring configuration, delete pushgateway monitoring configuration, query pushgateway monitoring configuration information	Supported	Supported
Authentication Mode	set the authentication mode, query the authentication mode	Supported	Supported
Pro Trial	enable the Pro Trial for the free edition	Not supported	Supported
Performance Parameters of HBlock	tune the performance parameters of HBlock, view the performance tuning configuration	Supported	Supported
Service Management	stop HBlock on the server, start HBlock on the server, restart HBlock on the server	Supported	Supported
Uninstall HBlock	uninstall HBlock	Supported	Supported
System Query	query HBlock information, query HBlock service status, query HBlock version	Supported	Supported
Upgrade	query the upgrade status of HBlock	Supported	Supported
	upgrade HBlock	Supported	Supported for 2 years

1.6 Features Available After Software License Expiration (Trial or Subscription Mode)

Module	Feature
Help Command	help command (command line only)
Software License	import software license, list software licenses
LUN	delete a LUN
Server	query server information
Service Management	stop HBlock on the server, start HBlock on the server, restart HBlock on the server
Alarm	query HBlock alarm information
Event	query HBlock events, export HBlock events
Log	initiate HBlock log collection, query collected logs, delete collected logs
Administrator Password	change administrator password
Remote Assistance Configuration	set remote assistance configuration, delete remote assistance configuration, query remote assistance configuration
System Query	query HBlock information, query HBlock service status, query HBlock version
Uninstall HBlock	uninstall HBlock
Other	WEB support overview page (license, data path, fault domain, LUN)

1.7 Terms and Abbreviations

Term/Abbreviation	Description
ALUA	Asymmetric Logical Unit Access
CHAP	Challenge Handshake Authentication Protocol
CIDR	Classless Inter-Domain Routing, a method of classifying IP addresses for assigning IP addresses to users and efficiently routing IP packets on the Internet.
DSM	Device Specific Module
EC	Erasur Coding
interface IP address	The interface IP address can be viewed through the command ifconfig.
IQN	iSCSI Qualified Name
iSCSI	Internet Small Computer System Interface
IOPS	Input/Output Operations Per Second, the number of reads and writes per second.
I/O	Input/Output
LUN	Logical Unit Number
MPIO	Multipath I/O
NFS	Network File System
NTP	Network Time Protocol
OOS	Object-Oriented Storage
RAID	Redundant Arrays of Independent Disks
SAN	Storage Area Network
SPC	Small Computer System Interface (SCSI) Primary Commands
SSD	Solid State Disk
SSL	Secure Sockets Layer
target	Storage target.
UUID	Universally Unique Identifier

WWID	World Wide Identifier, the unique identifier of the LUN. If there are multiple LUNs on the HBlock side when the client connects to a LUN, the WWID can be used to identify the LUN to be connected.
Disk Path	Path used to store HBlock data.

2 Server Side Deployment

2.1 Environmental Requirements

Item	Description
Supporting Linux OS	CentOS 7/8/9 64-bit, CTyunOS 2/4 64-bit.
Hardware	x86 server, ARM server, Loongson servers. Minimum configuration: Single-core CPU, 2 GB memory. Configurations can be added based on actual business needs.
Bandwidth	<ul style="list-style-type: none"> ● Bandwidth from client to HBlock: Read and write bandwidth capabilities are greater than business read and write bandwidth. ● The write bandwidth capacity of the disk partition corresponding to the disk path is greater than the bandwidth used by users to write data. ● Cloud bandwidth is greater than business write bandwidth.
Disk of installation directory	Above 10GB, RAID 1 or RAID 10 is recommended.
Disk path	<ul style="list-style-type: none"> ● Minimum configuration: 5GB, configuration can be increased according to actual business needs. ● Configure the capacity of the partition corresponding to the disk path based on the storage capacity and replica mode of LUN. <p>For the directories used by HBlock, it is recommended to set up automatic mounting at boot time or use a directory or subdirectory that has been set up for automatic mounting.</p>
Network settings	<p>Can connect to object storage network (not required for deploying local mode LUNs).</p> <p>The overall network architecture is as follows:</p> <ol style="list-style-type: none"> 1. The nodes within HBlock are interconnected through the intranet. 2. HBlock is interconnected with upper-layer applications through an intranet, dedicated line, or public network. 3. HBlock is interconnected with object storage through dedicated line or public network, local LUNs are not required.




Figure1. Standalone mode network topology diagram

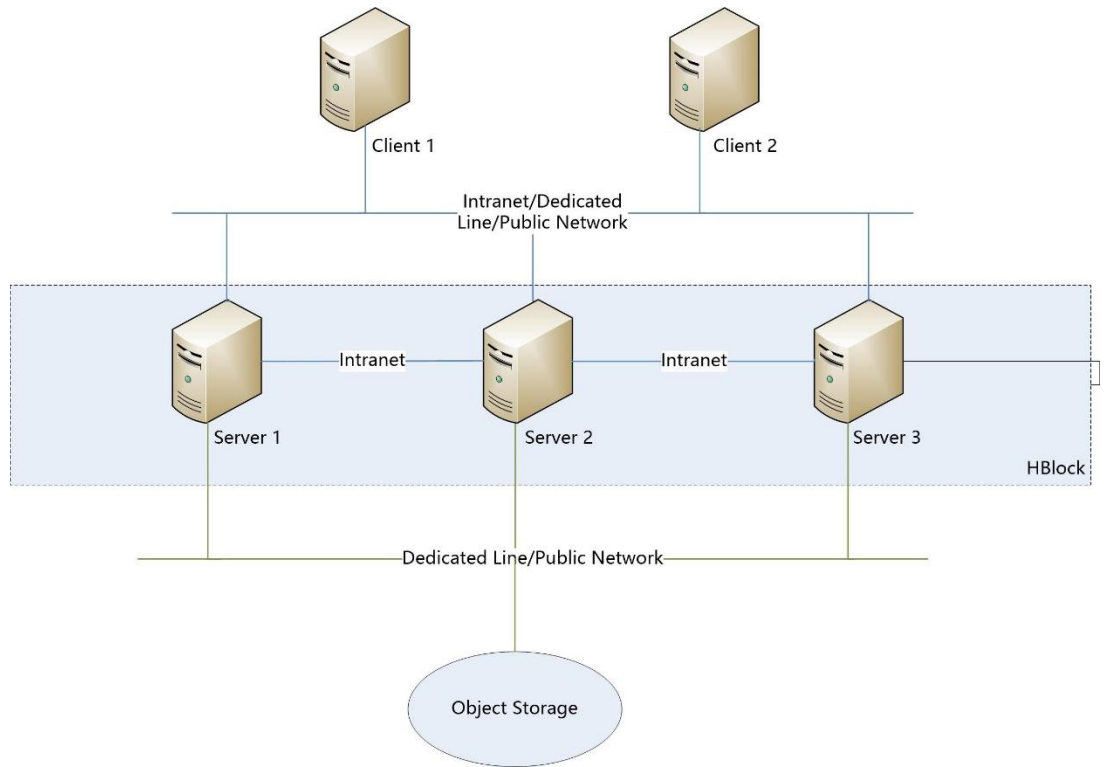


Figure2. HBlock cluster mode network topology diagram

Note: Confirm whether you are using the standalone mode or cluster mode before deploying HBlock. The mode cannot be changed once deployed.

2.2 HBlock Deployment – Standalone Mode

2.2.1 Configuration Environment

According to **Environmental Requirements**, prepare one server.

Note: Make sure that the **ping** command and **ps** command are available. For Debian/Ubuntu, you can use the following commands to install the **ping** command and **ps** command.

```
apt-get update      # Get the latest installation package
apt-get install iputils-ping    # Install ping
apt-get install procps        # Install ps
```

The server should be configured according to the following steps. The operations are illustrated using CentOS 7.x as an example:

Note: If the operating system is already installed, please ignore step I. If the disk is already mounted, please ignore step II. You can use the mounting path as a disk path for HBlock, or use the command `mkdir DIRECTORY` to create a directory under the mounting path and use it as an HBlock disk path.

(I) Install the CentOS 7.x operating system (Optional).

(II) Format the hard disk and mount it (Optional).

Please refer to the following examples to format the hard disk on your server for subsequent deployment.

```
lsblk      # View hard disk
mkfs.ext4 /dev/vdX    # Format the hard disk to ext4
mkdir DIRECTORY     # Create a mounting path, and DIRECTORY is the directory name.
mount /dev/vdX DIRECTORY    # Mount the hard disk
```

Note: The **mount** command is a temporary mounting command. After the server is restarted, it needs to be mounted again. For the directories used by HBlock, it is recommended to set up automatic mounting at boot time or use a directory or subdirectory that has been set up for automatic mounting.

Note: If the user installing HBlock is a non-root user, the read and write permissions on the directory used by HBlock are required. You can use the following command.

```
chown user:user-group DIRECTORY
# user is HBlock user, user-group is the user group to which HBlock users belong
```

(III) Close selinux and swap partitions (Recommended).

(IV) Configure the firewall.

If your server's firewall is not enabled, you can skip this step. If your server's firewall is

enabled, please open the iSCSI port to allow clients to connect to the server's target.

If the firewall software in use is firewalld, the example is as follows:

(a) Open the iSCSI port, for example, if the iSCSI port is 3260.

```
firewall-cmd --permanent --add-port=3260/tcp
```

(b) Reload the firewall to apply the configuration.

```
firewall-cmd --reload
```

If the firewall is iptables, the example is as follows:

(a) Open the iSCSI port, for example, if the iSCSI port is 3260.

```
iptables -I INPUT -p tcp --dport 3260 -j ACCEPT
```

(b) Save configuration.

```
iptables-save
```

(V) Set resource limits.

Edit file **/etc/security/limits.conf** to add the following contents, which set the maximum number of open files and processes in *domain*.

Note: Only non-root users need to manually modify **/etc/security/limits.conf**.

```
domain soft nfile 65536 # Set a value for the parameter domain as required
domain hard nfile 65536 # Set a value for the parameter domain as required
domain soft nproc 65535 # Set a value for the parameter domain as required
domain hard nproc 65535 # Set a value for the parameter domain as required
```

domain can be a *username*, *groupname*, *uid* or *wildcard*, and you can set it according to actual needs.

Note: If *domain* is set to *username*, *domain* must include the user who starts the HBlock service.

Example 1: If the value of the parameter *domain* is ***, it means all users can open up to 65,536 files and run up to 65,535 concurrent processes.

```
* soft nfile 65536 # * is the value of the parameter domain
* hard nfile 65536 # * is the value of the parameter domain
* soft nproc 65535 # * is the value of the parameter domain
* hard nproc 65535 # * is the value of the parameter domain
```

Example 2: If the value of the parameter *domain* is *root*, it means the root user can open up to 65,536 files and run up to 65,535 concurrent processes.

```
root soft nfile 65536 # root is the value of the parameter domain
root hard nfile 65536 # root is the value of the parameter domain
root soft nproc 65535 # root is the value of the parameter domain
root hard nproc 65535 # root is the value of the parameter domain
```

2.2.2 Deploy HBlock

The main steps to deploy HBlock are:

(I) Preparation before installation: Prepare one or more directories as HBlock disk paths to store HBlock data. Users installing HBlock must have read and write permissions on these paths.

Note: To avoid mutual influence, it is recommended not to share disks or file systems with the operating system.

(II) Unzip the installation package and go to the folder path of the unzipped files.

(III) Install and initialize HBlock:

(IV) Obtain and import the software license

(V) Create an iSCSI target and query its information

(VI) Create a LUN and query its information.

Note: The following shows how to install and deploy HBlock on an x86 server. The installation and deployment on an ARM server or a Loongnix server are the same as those on an x86 server.

The specific steps are as follows:

(I) Please complete the following preparations first: Prepare one or more directories as HBlock disk paths on the server for storing HBlock data. For example: /mnt/storage01, /mnt/storage02.

(II) Place the installation package in the directory where HBlock will be installed on the server, then unzip it and enter the unzipped folder.

```
unzip CTYUN_HBlock_Plus_3.10.0_x64.zip
cd CTYUN_HBlock_Plus_3.10.0_x64
```

(III) Follow these steps to install and initialize HBlock:

1. Install HBlock

Note: The installation of HBlock and the execution of HBlock management operations should belong to the same user.

Install HBlock on the server.

```
./stor install [ { -a | --api-port } API_PORT ] [ { -w | --web-port } WEB_PORT ]
```

API_PORT: Specifies the API port number, the default port number is 1443.

WEB_PORT: Specifies the WEB port number, the default port number is 2443.

2. Initialize HBlock

Note: HBlock can be initialized through the web, command line, and API.

For details about initialize HBlock command, see **Initialize**.

```
./stor setup { -n | --stor-name } STOR_NAME [ { -u | --user-name } USER_NAME ] { -p | --password } PASSWORD { -s | --server } { SERVER_IP[:PORT]:PATH <1-n> } [ { -P | --public-network } CIDR ] [ --iscsi-port ISCSI_PORT ] [--port-range PORT1-PORT2 ] [ --management-port1 MANAGEMENT_PORT1 ] [ --management-port2 MANAGEMENT_PORT2 ] [ --management-port3 MANAGEMENT_PORT3 ] [ --management-port4 MANAGEMENT_PORT4 ] [ --management-port6 MANAGEMENT_PORT6 ]
```

3. Query server information

```
./stor server ls [ { -n | --server } SERVER_ID ] [ --port ]
```

(IV) Obtain and import the Software License

HBlock software provides a 30-day trial period. After expiration, only partial features remain available. For details, see **Features Available After Software License Expiration (Trial or Subscription Mode)**. You can obtain a software license by following the steps below.

1. Get the HBlock serial number:

```
./stor info { -S | --serial-id }
```

2. Contact the HBlock software supplier to obtain the software license. You need to provide the HBlock serial number when obtaining it.
3. After obtaining the software license, import it.

```
./stor license add { -k | --key } KEY
```

(V) Create an iSCSI target and query it

1. For details about create an iSCSI target command, see **Create an iSCSI Target**.

```
./stor target add { -n | --name } TARGET_NAME [ --max-sessions MAX_SESSIONS ] [ { -c | --chap-name } CHAP_NAME { -p | --password } CHAP_PASSWORD { -s | --status } STATUS ]
```

Note: If multiple sessions are allowed to be established by IQN under the iSCSI target, you can be achieved by configuring the parameter **--max-sessions MAX_SESSIONS**.

2. Query the iSCSI target information.

```
./stor target ls [ -c | --connection ] [ { -n | --name } TARGET_NAME ]
```

(VI) Create a LUN and query it.

1. Create a LUN.

For details about create a LUN command, see **Create a LUN**.

Create a local mode LUN.

```
./stor lun add { -n | --name } LUN_NAME { -p | --capacity } CAPACITY { -t | --target }  
TARGET_NAME [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy }  
WRITE_POLICY ] [ { -P | --path } PATH ] [ { -m | --mode } STORAGE_MODE ]
```

Create a cache/storage mode LUN.

```
./stor lun add { -n | --name } LUN_NAME { -p | --capacity } CAPACITY { -t | --target }  
TARGET_NAME [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy }  
WRITE_POLICY ] [ { -P | --path } PATH ] { -m | --mode } STORAGE_MODE { -B | --bucket }  
BUCKET_NAME {-A | --ak } ACCESS_KEY { -S | --sk } SECRET_KEY [ { -C | --cloud-storage-  
class } CLOUD_STORAGE_CLASS ] {-E | --endpoint } ENDPOINT [ --sign-version VERSION ] [ --  
region REGION ] [ { -M | --cloud-compression } CLOUD_COMPRESSION ] [ { -O | --object-  
size } OBJECT_SIZE ] [ { -X | --prefix } PREFIX ]
```

2. Query the LUN information.

```
./stor lun ls [ { -n | --name } LUN_NAME ]
```

2.2.3 Example

Scenarios

- The server IP address is 192.168.0.32, API port number is 1443, WEB port number is 2443, the installation path is /mnt/stor (mounted by /dev/vdd), disk paths are /mnt/storage01 (mounted by /dev/vdc), /mnt/storage02 (mounted by /dev/vda).
- Create LUNs: LUN luna1 is local mode, corresponding to the iSCSI target targeta, with a LUN capacity of 100G; Cache mode LUN lunb1, corresponding to the iSCSI target targetb, with a LUN capacity of 200G; The storage mode LUN lunc1 corresponds to the iSCSI target targetc, with a LUN capacity of 300G.
- The Object-Oriented Storage (OOS) bucket corresponding to the cache mode and storage mode LUNs is hblocktest3.
- The HBlock system name is stor1, the HBlock username is storuser, and the password is hblock12@.

Steps

- (I) Please complete the following preparations first: the installation path which is /mnt/stor, disk paths which are /mnt/storage01 and /mnt/storage02, object storage OOS bucket hblocktest3, AK/SK and prefix stor1.
- (II) Place the installation package in the directory where HBlock will be installed on the server, then unzip it and enter the unzipped folder.

```
[root@hblockserver stor]# ls
CTYUN_HBlock_Plus_3.10.0_x64.zip
[root@hblockserver stor]# unzip CTYUN_HBlock_Plus_3.10.0_x64.zip
.....
[root@hblockserver stor]# cd CTYUN_HBlock_Plus_3.10.0_x64
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]#
```

- (III) Follow these steps to install and initialize HBlock:

1. Install HBlock.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor install -w 2443
Do you agree with HBlock User Agreement? [Yes/No]
Used in Chinese mainland, follow https://www.ctyun.cn/portal/protocol/10073150
Otherwise, follow https://www.esurfingcloud.com/portal/protocol/20692906
y
Installing HBlock...
Installed successfully.
```

When all servers are installed, please initialize HBlock in any of the following ways:

1. Use web portal to initialize HBlock. The https port is 2443.
2. Use management API (POST /rest/v1/system/setup) to initialize HBlock. The https port is 1443.
3. Use command line (stor setup) to initialize HBlock. Type 'stor --help setup' for more information.

2. Initialize HBlock

Passwords can be entered interactively, and when using interactive password input, they are displayed in ciphertext form.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor setup -n stor1 -u storuser -s 192.168.0.32:/mnt/storage01,/mnt/stor
Please enter password:
*****
Start to setup HBlock, please wait.
Processing...
Setup successfully and the HBlock services have been started.
Welcome to HBlock!
You are using a 30-day trial version. Please follow the steps to get a license.
1. Run "stor info --serial-id" to get the serial ID of the HBlock
2. Contact the software vendor to obtain a license
3. Run "stor license add -k KEY" to import the license

Type 'stor --help' to get more information, such as managing LUNs, targets, servers, etc.
```

3. Query server information

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls -n hblock_1
Server Name: hblockserver
Server ID:
hblock_1
Status: Connected
Public Address: 192.168.0.32:3260
Cluster Address: 192.168.0.32
Recent Start Time: 2025-12-26 10:25:05
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.95,"maxMemorySize":9223372036853727232}
Disk Path(s):
```

No.	Path	Used Capacity	Total Capacity	Used Capacity Quota	Capacity Quota	Health Status	Health Detail
1.	/mnt/storage01(*)	19.37 GiB	93.29 GiB	4 KiB	Unlimited	Healthy	-
2.	/mnt/stor	1.17 GiB	93.29 GiB	4 KiB	Unlimited	Healthy	-

(IV) Obtain and import the Software License

HBlock software provides a 30-day trial period, after which management operations cannot be performed. You can obtain a software license by following the steps below.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info --serial-id
HBlock serial ID: HBlock serial ID: 0A769770-1528-49B8-A82D-D1CB71333FFB-0201-040000
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor license add -k *****
Import license successfully.
Customer name: Customer ABC
As the operating entity utilizing the HBlock to run your business, you bear the responsibility for the network and information security of the business you are operating.
```

```
The current system time: 2025-12-26 11:31:29
+-----+-----+-----+-----+-----+
| LicenseId          | Account      | Type      | Status      | MaximumLocalCapacity |
+-----+-----+-----+-----+-----+
| qws2b6a9-f3fb-4098-a6b3-3652a5a76530 | test@ctyun.cn | Perpetual | Effective   | 2 PiB                |
+-----+-----+-----+-----+-----+

License qws2b6a9-f3fb-4098-a6b3-3652a5a76530 (Effective):
Usage:
+-----+-----+-----+-----+-----+
| MaximumLocalCapacity | MaintenanceEffectiveTime | MaintenanceExpireTime | Status      |
+-----+-----+-----+-----+-----+
| 2 PiB                | 2025-12-26 11:25:05      | 2026-12-26 11:25:05   | Effective   |
+-----+-----+-----+-----+-----+
```

(V) Create an iSCSI target and query it

1. Create iSCSI targets.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target add -n targeta --max-
sessions 3 -c chap-test -p ***** -s on
Created target targeta successfully.
iqn = iqn.2012-08.cn.ctyunapi.oos:targeta.1(192.168.0.32:3260)
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target add -n targetb --max-
sessions 2
Created target targetb successfully.
iqn = iqn.2012-08.cn.ctyunapi.oos:targetb.2(192.168.0.32:3260)
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target add -n targetc
Created target targetc successfully.
iqn = iqn.2012-08.cn.ctyunapi.oos:targetc.3(192.168.0.32:3260)
```

2. Query the iSCSI target information.

Query information of all iSCSI targets.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor target ls
+-----+-----+-----+-----+-----+
| No. | Target Name | Max Sessions | ISCSI Target | CHAP |
+-----+-----+-----+-----+-----+
| 1. | targeta     | 3           | iqn.2012-08.cn.ctyunapi.oos:targeta.1(192.168.0.32:3260) | chap-test,Enabled |
| 2. | targetb     | 2           | iqn.2012-08.cn.ctyunapi.oos:targetb.2(192.168.0.32:3260) | Disabled           |
| 3. | targetc     | 1           | iqn.2012-08.cn.ctyunapi.oos:targetc.3(192.168.0.32:3260) | Disabled           |
+-----+-----+-----+-----+-----+
```

Query the specific iSCSI target information, such as targeta.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor target ls -n targeta
Target Name: targeta
Max Sessions: 3
Create Time: 2025-12-26 12:28:05
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targeta.1(192.168.0.32:3260)
```

Reclaim Policy: Delete
 CHAP: chap-test,test1test2024,Enabled

(VI) Create a LUN and query it.

1. Create LUNs.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n luna1 -p 100 -t targeta
Created LUN luna1 successfully.
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lunb1 -t targetb -p 200 -m
Cache -B hblocktest3 -E oos-cn.ctyunapi.cn -A cb22b08b1f9229f85874 -S ***** --sign-
version v4 --region cn -X stor1
Created LUN lunb1 successfully.
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lunc1 -t targetc -p 300 -m
Storage -B hblocktest3 -E oos-cn.ctyunapi.cn -A cb22b08b1f9229f85874 -S ***** --sign-
version v4 --region cn -X stor1
Created LUN lunc1 successfully.
```

2. Query the LUN information.

Query the information of all LUNs.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun ls
```

No.	LUN Name	Storage Mode	Capacity	Status	Target	Snapshot Count	Snapshot Size	Is Clone
1.	luna1(LUN 0)	Local	100 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targeta.1(192.168.0.32:3260,Active)	0	0 B	
2.	lunb1(LUN 0)	Cache	200 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targetb.2(192.168.0.32:3260,Active)	0	0 B	
3.	lunc1(LUN 0)	Storage	300 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targetc.3(192.168.0.32:3260,Active)	0	0 B	

Query the specific LUN information, such as luna1.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun ls -n luna1
LUN Name: luna1 (LUN 0)
Storage Mode: Local
Capacity: 100 GiB
Status: Normal
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targeta.1(192.168.0.32:3260,Active)
Create Time: 2025-12-26 13:28:05
Local Sector Size: 4096 Bytes
Data Health: 100% normal, 0% low redundancy, 0% error
Write Policy: WriteBack
WWID: 33ffffffffce6bc2a6
UUID: lun-uuid-9885d5c6-23f7-4739-b580-fc5a233ea3bd
Path: /mnt/storage01
Snapshot Count: 0
Snapshot Size: 0 B (Note: Snapshot size may vary due to LUN issues or parent snapshot
deletion.)
```

2.3 HBlock Deployment – Cluster Mode

2.3.1 Configuration Environment

According to **Environmental Requirements**, prepare three or more servers.

Note: Make sure that the **ping** command and **ps** command are available. For Debian/Ubuntu, you can use the following commands to install the **ping** command and **ps** command.

```
apt-get update #Get the latest installation package
apt-get install iputils-ping # Install ping
apt-get install procps # Install ps
```

Each server should be configured according to following the steps. The following operations take CentOS 7.x version as an example:

Note: If the operating system is already installed, please ignore step I. If the disk is already mounted, please ignore step II. You can use the mounting path as the disk path for HBlock, or use the command `mkdir DIRECTORY` to create a directory under the mounting path and use it as the HBlock disk path.

(I) Install the CentOS 7.x operating system (Optional).

(II) Format the hard disk and mount it (Optional).

Please refer to the following examples to format the hard disk on the server to facilitate subsequent deployment.

```
lsblk # View hard disk
mkfs.ext4 /dev/vdX #Format the hard disk to ext4
mkdir DIRECTORY #Create a mounting path, and DIRECTORY is the directory name.
mount /dev/vdX DIRECTORY
#Mount the hard disk, and after mounting, you can use this path as the HBlock disk path
```

Note: The **mount** command is a temporary mounting command. After the server is restarted, it needs to be mounted again. For the directories used by HBlock, it is recommended to set up automatic mounting at boot time or use a directory or subdirectory that has been set up for automatic mounting.

Note: If the user installing HBlock is a non-root user, the read and write permissions on the directory used by HBlock are required. You can use the following command.

```
chown user:user-group DIRECTORY
# user is HBlock user, user-group is the user group to which HBlock users belong
```

(III) Close selinux and swap partitions (Recommended).

(IV) Configure the firewall.

Make sure that the cluster servers can access each other, and are mutually trustlisted. Also, open the iSCSI port to allow clients to connect to server's target. If it is installed on a cloud host, a trustlist also needs to be added to the security group.

If your server's firewall is not enabled, you can skip this step.

If the firewall software in use is firewalld, the example is as follows:

(a) Open the iSCSI port, for example, if the iSCSI port is 3260.

```
firewall-cmd --permanent --add-port=3260/tcp
```

(b) Add the IP of each server in the cluster to the trustlist:

Add IPv4 address.

```
firewall-cmd --permanent --add-rich-rule="rule family=ipv4 source address=your_IP accept"
# your_IP is IP address allowed to access
```

Add IPv6 address.

```
firewall-cmd --permanent --add-rich-rule="rule family=ipv6 source address=your_IP accept"
#your_IP is IP address allowed to access
```

(c) Reload the firewall to apply the configuration.

```
firewall-cmd --reload
```

If the firewall is iptables, the example is as follows:

(a) Open the iSCSI port, for example, if the iSCSI port is 3260.

```
iptables -I INPUT -p tcp --dport 3260 -j ACCEPT
```

(b) Configure mutual access between servers in the cluster:

Allow the loopback address.

```
iptables -I INPUT -i lo -j ACCEPT
```

Allow access from the internal network segment.

```
iptables -I INPUT -s your_IP -j ACCEPT
```

(c) Save configuration.

```
iptables-save
```

(V) Set resource limits.

Edit file **/etc/security/limits.conf** to add the following contents, which set the maximum number of open files and processes in domain.

Note: Only non-root users need to manually modify **/etc/security/limits.conf**.

```
domain soft nofile 65536
domain hard nofile 65536
domain soft nproc 65535
domain hard nproc 65535
```

domain can be a *username*, *groupname*, *uid* or *wildcard*, and you can set it according to actual needs.

Note: The *domain* must include the user who starts the HBlock service.

Example 1: If the value of the parameter domain is *, it means all users can open up to 65,536 files and run up to 65,535 concurrent processes.

```
* soft nofile 65536
* hard nofile 65536
* soft nproc 65535
* hard nproc 65535
```

Example 2: If the value of the parameter domain is root, it means the root user can open up to 65,536 files and run up to 65,535 concurrent processes.

```
root soft nofile 65536
root hard nofile 65536
root soft nproc 65535
root hard nproc 65535
```

2.3.2 Deploy HBlock

The main steps to deploy HBlock are:

(I) Preparation before installation: Prepare one or more directories as HBlock disk paths on each server to store HBlock data.

Note: To avoid mutual influence, it is recommended that directories do not share a disk or file system with the operating system.

(II) Unzip the installation package and enter the unzipped folder path.

(III) Install and initialize HBlock:

(IV) Obtain and import the software license

(V) Create an iSCSI target and query it.

(VI) Create a LUNs and query it.

Note: The following uses the installation and deployment of an x86 server as an example. The installation and deployment of an ARM server are the same as that of an x86 server.

The detailed steps are as follows:

(I) Please complete the following preparations first: Prepare one or more directories as HBlock disk paths on each server for storing HBlock data, the directories on each server can be different. For example: /mnt/storage01, /mnt/storage02.

(II) Place the installation package in the directory where HBlock will be installed on each server, then unzip it and enter the unzipped folder.

```
unzip CTYUN_HBlock_Plus_3.10.0_x64.zip
cd CTYUN_HBlock_Plus_3.10.0_x64
```

(III) Follow these steps to install and initialize HBlock:

1. Install HBlock

Note: The installation of HBlock and the execution of HBlock management operations should belong to the same user.

Install HBlock on each server.

```
./stor install [ { -a | --api-port } API_PORT ] [ { -w | --web-port } WEB_PORT ]
```

API_PORT: Specifies the API port number, the default port number is 1443.

WEB_PORT: Specifies the WEB port number, the default port number is 2443.

2. Initialize HBlock

Note: HBlock can be initialized through the web, command line, and API. Initialization can be performed on any server that has already installed HBlock.

For details about initialize HBlock command, see **Initialize**.

```
./stor setup { -n | --stor-name } STOR_NAME [ { -u | --user-name } USER_NAME ] { -p | --password } PASSWORD { -s | --server } { SERVER_IP[:PORT]:PATH <1-n> } <1-n> [ {-C | --cluster-network } CIDR ] [ { -P | --public-network } CIDR ] [ --fault-domain FAULT_DOMAIN ] [ --iscsi-port ISCSI_PORT ] [--port-range PORT1-PORT2 ] [ --data-port1 DATA_PORT1 ] [ --management-port1 MANAGEMENT_PORT1 ] [ --management-port2 MANAGEMENT_PORT2 ] [ --management-port3 MANAGEMENT_PORT3 ] [ --management-port4 MANAGEMENT_PORT4 ] [ --management-port5 MANAGEMENT_PORT5 ] [ --management-port6 MANAGEMENT_PORT6 ] [ --metadata-port1 METADATA_PORT1 ] [ --metadata-port2 METADATA_PORT2 ] [ --metadata-port3 METADATA_PORT3 ] [ --metadata-port4 METADATA_PORT4 ] [ --metadata-port5 METADATA_PORT5 ] [ --metadata-port6 METADATA_PORT6 ] [ --metadata-port7 METADATA_PORT7 ] [ --metadata-port8 METADATA_PORT8 ] [ --cs SERVER_IP,SERVER_IP,SERVER_IP ] [ --mdm SERVER_IP,SERVER_IP ] [ --ls SERVER_IP,SERVER_IP,SERVER_IP ]
```

3. Query server information

```
./stor server ls [ { -n | --server } SERVER_ID ] [ --port ]
```

(IV) Obtain and import the Software License

HBlock software provides a 30-day trial period. After expiration, only partial features remain available. For details, see **Features Available After Software License Expiration (Trial or Subscription Mode)**. You can obtain a software license by following the steps below.

1. Get the HBlock serial number:

```
./stor info { -S | --serial-id }
```

2. Contact the HBlock software supplier to obtain the software license. You need to provide the HBlock serial number when obtaining it.
3. After obtaining the software license, import it.

```
./stor license add { -k | --key } KEY
```

(V) Create an iSCSI target and query it

1. For details about create an iSCSI target command, see [Create an iSCSI Target](#).

```
./stor target add { -n | --name } TARGET_NAME [ --max-sessions MAX_SESSIONS ] [ { -c | --chap-name } CHAP_NAME { -p | --password } CHAP_PASSWORD { -s | --status } STATUS ]
```

Note: If multiple sessions are allowed to be established by IQN under the iSCSI target, you can be achieved by configuring the parameter **--max-sessions** *MAX_SESSIONS*.

2. Query the iSCSI target information.

```
./stor target ls [ -c | --connection ] [ { -n | --name } TARGET_NAME ]
```

(VI) Create a LUN and query it.

1. Create a LUN.

For details about create a LUN command, see **Create a LUN**.

Create a local mode LUN.

```
./stor lun add { -n | --name } LUN_NAME { -p | --capacity } CAPACITY { -t | --target } TARGET_NAME [ { -a | --ha } HIGH_AVAILABILITY ] [ { -c | --local-storage-class } LOCAL_STORAGE_CLASS ] [ --min-replica MIN_REPLICA ] [ --ec-fragment-size EC_FRAGMENT_SIZE ] [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy } WRITE_POLICY ] [ { -m | --mode } STORAGE_MODE ]
```

Create a cache/storage mode LUN.

```
./stor lun add { -n | --name } LUN_NAME { -p | --capacity } CAPACITY { -t | --target } TARGET_NAME [ { -a | --ha } HIGH_AVAILABILITY ] [ { -c | --local-storage-class } LOCAL_STORAGE_CLASS ] [ --min-replica MIN_REPLICA ] [ --ec-fragment-size EC_FRAGMENT_SIZE ] [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy } WRITE_POLICY ] { -m | --mode } STORAGE_MODE { -B | --bucket } BUCKET_NAME { -A | --ak } ACCESS_KEY { -S | --sk } SECRET_KEY [ { -C | --cloud-storage-class } CLOUD_STORAGE_CLASS ] { -E | --endpoint } ENDPOINT [ --sign-version VERSION ] [ --region REGION ] [ { -M | --cloud-compression } CLOUD_COMPRESSION ] [ { -O | --object-size } OBJECT_SIZE ] [ { -X | --prefix } PREFIX ]
```

2. Query the LUN information.

```
./stor lun ls [ { -n | --name } LUN_NAME ]
```

2.3.3 Examples

Scenarios

- Three servers:
 - The first server: The server IP address is 192.168.0.110, API port number is 1443, WEB port number is 2443, the installation path is /mnt/storage01 (mounted by /dev/vdc), disk paths are /mnt/stor (mounted by /dev/vdd), /mnt/storage02 (mounted by /dev/vdb).
 - The second server: The server IP address is 192.168.0.192, API port number is 1443, WEB port number is 2443, the installation path is /mnt/storage01 (mounted by /dev/vda), disk path is /mnt/stor (mounted by /dev/vdd).
 - The third server: The server IP address is 192.168.0.102, API port number is 1443, WEB port number is 2443, the installation path is /mnt/storage01 (mounted by /dev/vda), disk path is /mnt/stor (mounted by /dev/vdd).
- Create LUNs: LUN lun01a is local mode, corresponding to the iSCSI target target01, with a LUN capacity of 100G; Cache mode LUN lun02a, corresponding to the iSCSI target target02, with a LUN capacity of 200G; The storage mode LUN lun03a corresponds to the iSCSI target target03, with a LUN capacity of 300G.
- The OOS bucket corresponding to the cache mode and storage mode LUNs is hblocktest3.
- The HBlock system name is stor2, the HBlock username is storuser, and the password is hblock12@.

Steps

- (I) Please complete the following preparations:
 - The first server: the installation path is /mnt/storage01, disk paths are /mnt/stor and /mnt/storage02.
 - The second server: the installation path is /mnt/storage01, disk path is /mnt/stor.
 - The third server: the installation path is /mnt/storage01, disk path is /mnt/stor.
 - OOS bucket hblocktest3, AK/SK, and prefix stor2.
- (II) Place the installation package in the directory where HBlock will be installed on each server, then unzip it and enter the unzipped folder.

```
[root@hblockserver storage01]# ls
CTYUN_HBlock_Plus_3.10.0_x64.zip
[root@hblockserver storage01]# unzip CTYUN_HBlock_Plus_3.10.0_x64.zip
```

```
.....
[root@hblockserver storage01]# cd CTYUN_HBlock_Plus_3.10.0_x64
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]#
```

(III) Follow these steps to install and initialize HBlock:

1. Install HBlock on each server.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor install -w 2443
Do you agree with HBlock User Agreement? [Yes/No]
Used in Chinese mainland, follow https://www.ctyun.cn/portal/protocol/10073150
Otherwise, follow https://www.esurfingcloud.com/portal/protocol/20692906
y
Installing HBlock...
Installed successfully.
When all servers are installed, please initialize HBlock in any of the following ways:
1. Use web portal to initialize HBlock. The https port is 2443.
2. Use management API (POST /rest/v1/system/setup) to initialize HBlock. The https port is 1443.
3. Use command line (stor setup) to initialize HBlock. Type 'stor --help setup' for more information.
```

2. Initialize HBlock

Passwords can be entered interactively, and when using interactive password input, they are displayed in ciphertext form.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor setup -n stor2 -p hblock12@ -s 192.168.0.110:/mnt/stor,/mnt/storage02 192.168.0.192:/mnt/stor 192.168.0.102:/mnt/stor
Start to setup HBlock, please wait.
Processing...
Setup successfully and the HBlock services have been started.
Welcome to HBlock!
You are using a 30-day trial version. Please follow the steps to get a license.
1. Run "stor info --serial-id" to get the serial ID of the HBlock
2. Contact the software vendor to obtain a license
3. Run "stor license add -k KEY" to import the license

Type 'stor --help' to get more information, such as managing LUNs, targets, servers, etc.
```

3. Query server information

Query information of all servers.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor server ls
```

No.	Server ID	Server Name	Status	Public Address	Cluster Address	Recent Start Time
1.	hblock_1(M**)	ecs-9689-0915140	Connected	192.168.0.102:3260	192.168.0.102	2025-01-17 13:58:21
2.	hblock_2(**)	pm-006	Connected	192.168.0.192:3260	192.168.0.192	2025-01-17 14:11:52
3.	hblock_3(**)	hblockserver	Connected	192.168.0.110:3260	192.168.0.110	2025-01-17 14:12:01

Query information of the specific server, such as hblock_1.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor server ls -n hblock_1
Server Name: ecs-9689-0915140
Server ID: hblock_1
Node Name: default:server1
```

```
Parent Node: default
Status: Connected
Master Server: true
Base Server: true
Base Service: mdm(Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
               ls(Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
               cs(Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
Public Address: 192.168.0.102:3260
Cluster Address: 192.168.0.102
Recent Start Time: 2025-12-26 10:33:43
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.95,"maxMemorySize":9223372036853727232}
Disk Path(s):
```

No.	Path	Storage Pool	Used Capacity	Total Capacity	Used Capacity Quota	Capacity Quota	Health Status	Health Detail	Data Service
1.	/mnt/stor	default	10.22 GiB	93.29 GiB	138.49 KiB	Unlimited	Healthy	-	ds-1

(IV) Obtain and import the Software License

HBlock software provides a 30-day trial period, after which management operations cannot be performed. You can obtain a software license by following the steps below.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info --serial-id
HBlock serial ID: 1E46239F-BD8B-4CB3-A151-C3A58892AEF9-0202-040000
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor license add -k *****
Import license successfully.
Customer name: Customer ABC
As the operating entity utilizing the HBlock to run your business, you bear the responsibility for the network and
information security of the business you are operating.

The current system time: 2025-12-26 11:33:43
```

LicenseId	Account	Type	Status	MaximumLocalCapacity
ehc2b6a9-f3fb-4098-a6b3-3652a5d71232	testaccount1@ctyun.com	Subscription	Effective	1 PiB

```
License ehc2b6a9-f3fb-4098-a6b3-3652a5d71232 (Effective):
Usage:
```

MaximumLocalCapacity	EffectiveTime	ExpireTime	Status
1 PiB	2025-12-26 11:24:30	2026-12-26 11:24:30	Effective

(V) Create an iSCSI target and query it

1. Create iSCSI targets.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target add -n target01 --max-sessions 6 -c chap-test1 -p T12345678912 -s on --num 3 --server hblock_1,hblock_2,hblock_3
Created target target01 successfully.
iqn = iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.110:3260)
iqn = iqn.2012-08.cn.ctyunapi.oos:target01.2(192.168.0.192:3260)
iqn = iqn.2012-08.cn.ctyunapi.oos:target01.16(192.168.0.102:3260) [root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target add -n target02 --max-sessions 2
Created target target02 successfully.
iqn = iqn.2012-08.cn.ctyunapi.oos:target02.3(192.168.0.102:3260)
iqn = iqn.2012-08.cn.ctyunapi.oos:target02.4(192.168.0.110:3260)
```

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target add -n target03
Created target target03 successfully.
iqn = iqn.2012-08.cn.ctyunapi.oos:target03.5(192.168.0.102:3260)
iqn = iqn.2012-08.cn.ctyunapi.oos:target03.6(192.168.0.110:3260)
```

2. Query the iSCSI target information.

Query information of all iSCSI targets.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor target ls
```

No.	Target Name	Max Sessions	iSCSI Target	CHAP
1.	target01	6	iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.110:3260)	chap-test1,Enabled
			iqn.2012-08.cn.ctyunapi.oos:target01.2(192.168.0.192:3260)	
			iqn.2012-08.cn.ctyunapi.oos:target01.16(192.168.0.102:3260)	
2.	target02	2	iqn.2012-08.cn.ctyunapi.oos:target02.3(192.168.0.102:3260)	Disabled
			iqn.2012-08.cn.ctyunapi.oos:target02.4(192.168.0.110:3260)	
3.	target03	1	iqn.2012-08.cn.ctyunapi.oos:target03.5(192.168.0.102:3260)	Disabled
			iqn.2012-08.cn.ctyunapi.oos:target03.6(192.168.0.110:3260)	

Query the specific iSCSI target information, such as target01.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -n target01
Target Name: target01
Max Sessions: 6
Create Time: 2025-12-26 12:33:43
Number of Servers: 3
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.110:3260)
               iqn.2012-08.cn.ctyunapi.oos:target01.2(192.168.0.192:3260)
               iqn.2012-08.cn.ctyunapi.oos:target01.16(192.168.0.102:3260)
Reclaim Policy: Retain
CHAP: chap-test1,T12345678912,Enabled
ServerID: hblock_1,hblock_2,hblock_3
```

(VI) Create a LUN and query it.

1. Create LUNs.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lun01a -t target01 -c EC
2+1 --ec-fragment-size 32 -p 100 -m Local
Created LUN lun01a successfully.
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lun02a -t target02 -c EC
2+1 --ec-fragment-size 32 -p 200 -m Cache -B hblocktest3 -E oos-cn.ctyunapi.cn -A
cb22b08b1f9229f85874 -S ***** --sign-version v4 --region cn -X stor2
Created LUN lun02a successfully.
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lun03a -t target03 -c EC
2+1 --ec-fragment-size 32 -p 300 -m Storage -B hblocktest3 -E oos-cn.ctyunapi.cn -A
cb22b08b1f9229f85874 -S ***** --sign-version v4 --region cn -X stor2
Created LUN lun03a successfully.
```

2. Query the LUN information.

Query the information of all LUNs

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls
```

No.	LUN Name	Storage Mode	Capacity	Local Storage Class	Minimum Replicas	Status	Target	Snapshot Count	Snapshot Size	Is Clone
1.	lun01a(LUN 0)	Local	110 GiB	EC 2+1+32 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.102:3260,Active) iqn.2012-08.cn.ctyunapi.oos:target01.2(192.168.0.192:3260,Standby)	0	0 B	
2.	lun02a(LUN 0)	Cache	200 GiB	EC 2+1+32 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.110:3260,ColdStandby) iqn.2012-08.cn.ctyunapi.oos:target02.3(192.168.0.102:3260,Active) iqn.2012-08.cn.ctyunapi.oos:target02.4(192.168.0.110:3260,Standby)	0	0 B	
3.	lun03a(LUN 0)	Local	300 GiB	EC 2+1+32 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target03.5(192.168.0.102:3260,Active) iqn.2012-08.cn.ctyunapi.oos:target03.6(192.168.0.110:3260,Standby)	0	0 B	

Query the specific LUN information, such as lun02a.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun02a
LUN Name: lun02a (LUN 0)
Storage Mode: Cache
Capacity: 200 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target02.3(192.168.0.110:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target02.4(192.168.0.102:3260,Standby)
Create Time: 2025-12-26 12:36:43
Local Storage Class: EC 2+1+32 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
High Availability: ActiveStandby
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
Write Policy: WriteBack
WWID: 33ffffffff805dcfc1
UUID: lun-uuid-9457ebfc-8440-4e69-bc9c-c054938e5551
Object Storage Info:
+-----+-----+
| Provider       | OOS          |
| Bucket Name    | hblocktest3  |
| Prefix         | stor2        |
| Endpoint       | https://oos-cn.ctyunapi.cn |
| Signature Version | v4          |
| Region        | cn           |
| Storage Class  | STANDARD     |
| Access Key     | cb22b08b1f9229f85874 |
| Object Size    | 1024 KiB     |
| Compression    | Enabled      |
+-----+-----+
```

3 Client Operations

3.1 Windows Client – Standalone Mode

- (I) Prepare the client operating system.

It is recommended to use Windows 10, Windows Server 2012 R2, Windows Server 2016 R2 and other high versions of Windows operating systems. Because these systems have iSCSI initiator, it is unnecessary to install components separately.

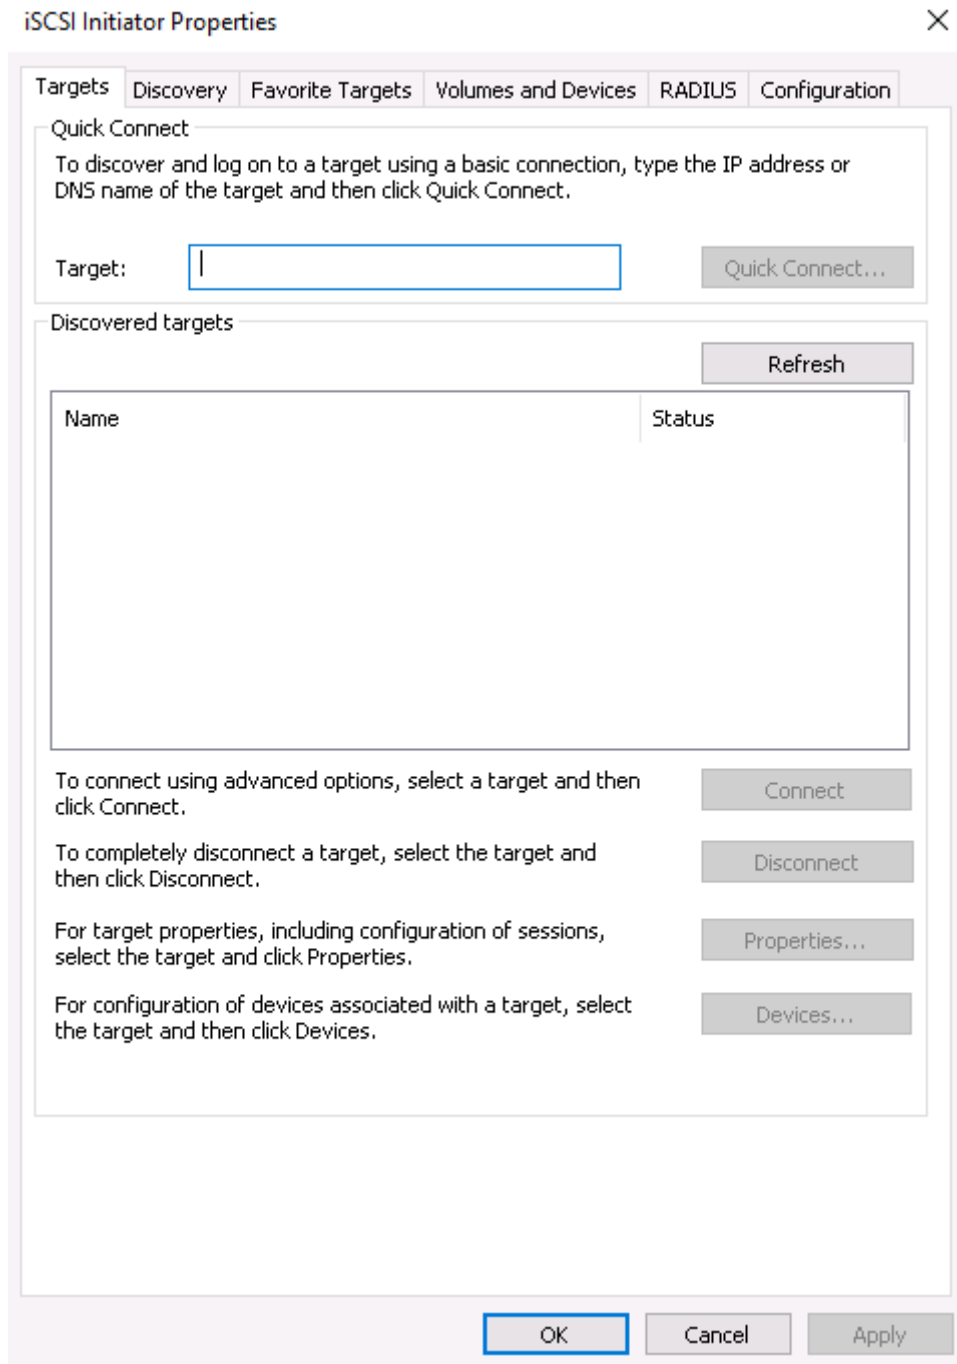
Different operating system versions of the client support different single-LUN capacities.

Please refer to the following table:

Windows Version	Block Size	Maximum Capacity of a Single LUN
Windows Server 2008R2	512 bytes / 4KiB	256 TiB
Windows Server 2012R2	512 bytes / 4KiB	256 TiB
Windows Server 2016	512 bytes / 4KiB	256 TiB
Windows 10	512 bytes / 4KiB	1 PiB

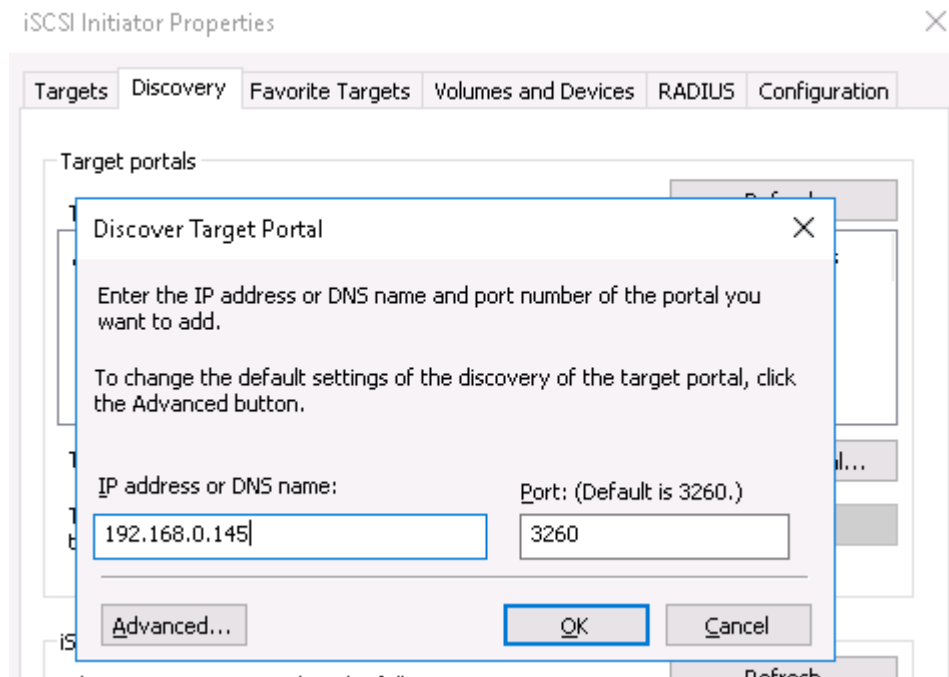
- (II) Run iSCSI initiator.

Run **iSCSI Initiator Properties** on the Windows client, as shown in the following figure:



1. Configure iSCSI initiator.

Enter the server IP and port in **Discovery Target Portal** of **Discovery**, as shown in the following figure:

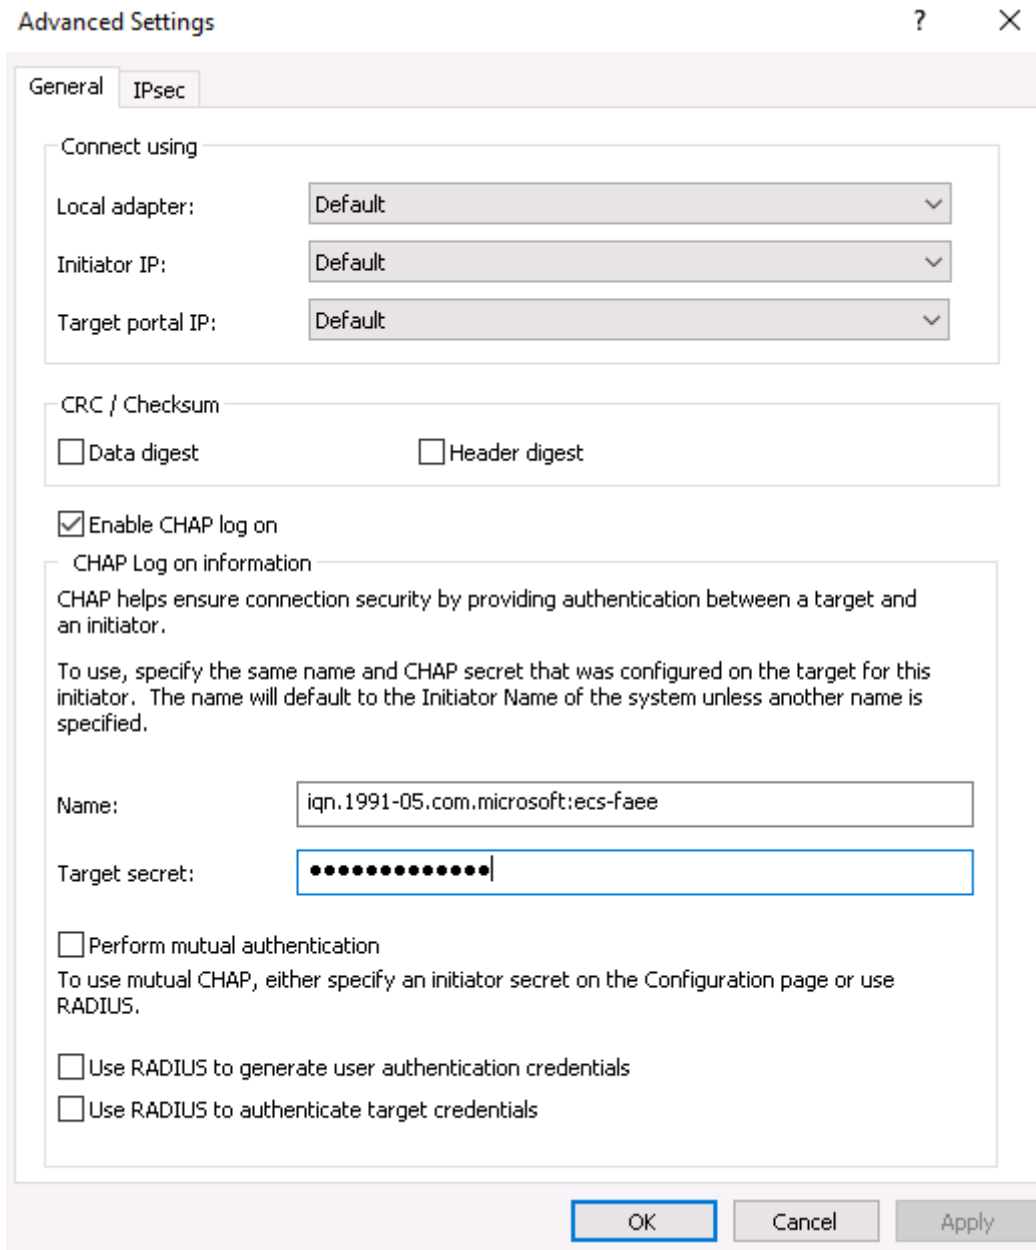


Search the iSCSI target issued by HBlock in **Discovered targets of Targets**, check that the status is **Inactive**, and click **Connect**.

2. Enable CHAP authentication.

If your iSCSI target has enabled CHAP authentication, click **Advanced** in the pop-up dialog box **Connect to Target**. Please ignore this step and direct connect if CHAP authentication is disabled.

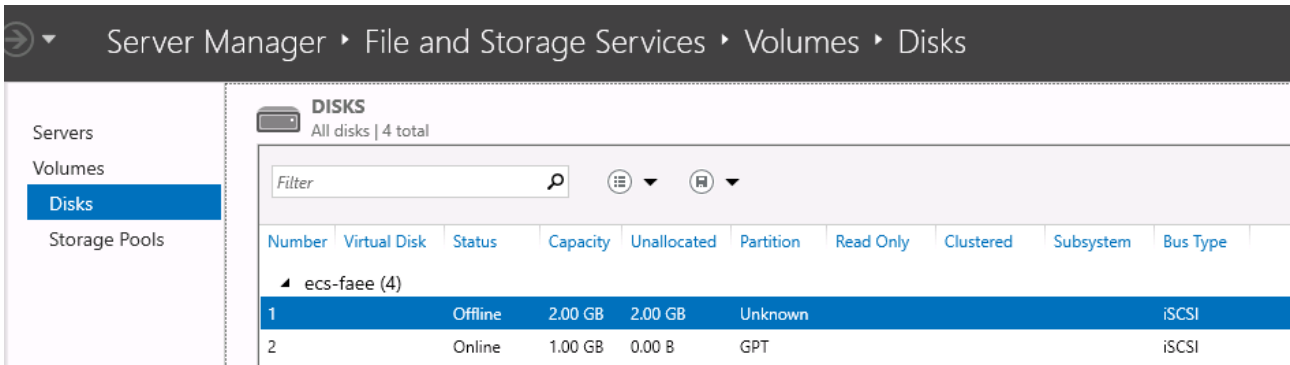
Select **Enable CHAP log on**, enter the same name and CHAP password that was configured on the target, and then click **OK**, as shown in the following figure.



3. Using iSCSI shared disks on the client.

Go to **Server Manager > File and Storage Services > Volumes > Disks**, and connect the disk with status off-line. Then **Initialize and New Volume**, specify the disk letter and format the disk. Start the **This PC**, you can see the new added disk and its capacity, and since then, you can use the iSCSI disk issued by as a local disk.

Note: During initialization of disk, if the storage capacity of LUN is less than or equal to 2TiB, either MBR or GPT can be used for partition; if the storage capacity of LUN is larger than 2TiB, only GPT can be used for partition.



Open the **This PC** and you can see the drive letter and capacity of the newly added disk. At this time, you can use the iSCSI disk released by HBlock according to the habit of using local disks.

Note:

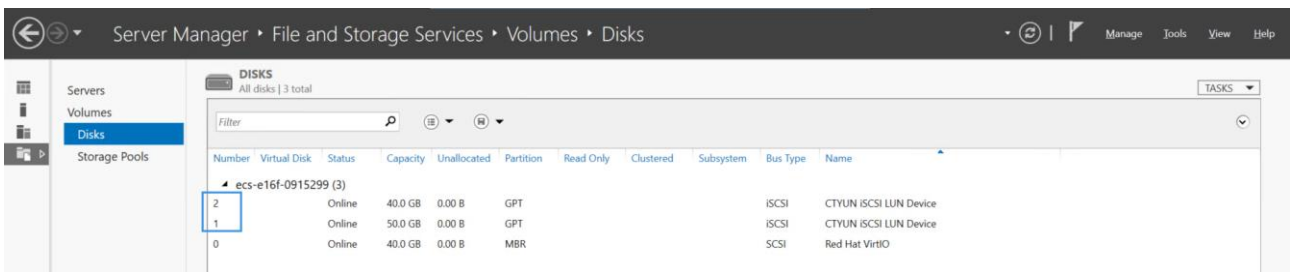
- If client needs to be disconnected or delete the disk, you need to go offline the disk first and then disconnect iSCSI target.
- If the client needs to disconnect and then re-connect, there is no need to initialize or create a new volume, and the disk can be seen after reconnecting.

Note: If you want to query the disk corresponding to HBlock LUN, you can enter the following command line on the client side to query.

```
wmic diskdrive get Name, Manufacturer, Model, InterfaceType, MediaType, SerialNumber
```

As shown in the following example, the query information "Name" column corresponds to the "Number" column on the "DISKS". SerialName corresponds to the name and uuid of HBlock LUN.

```
C:\Users\Administrator>wmic diskdrive get Name, Manufacturer, Model, InterfaceType, MediaType, SerialNumber
InterfaceType  Manufacturer      MediaType      Model          Name              SerialNumber
-----
SCSI           (Standard disk drives) Fixed hard disk media Red Hat VirtIO SCSI Disk Device \\.\PHYSICALDRIVE0 e3057108-d9d5-492d-9
SCSI           (Standard disk drives) Fixed hard disk media CTYUN          \\.\PHYSICALDRIVE1 lun01a-lun-uuid-1432bc36-3b5b-40bc-b74f-787722709399
SCSI           (Standard disk drives) Fixed hard disk media CTYUN          \\.\PHYSICALDRIVE2 lun02a-lun-uuid-5d2b6c71-2802-44ba-b6d3-3e1d1a39730d
```



3.2 Windows Client – Cluster Mode

Microsoft provides a common DSM (Device Specific Module) in Server 2008, 2012, and 2016, supports ALUA (Asymmetric Logical Unit Access), which can be used in conjunction with SPC (SCSI Primary Commands) compliant storage devices to configure MPIO (Multipath I/O) environments. MPIO ensures that the normal operation of the business will not be affected when switching between Active target and Standby target. Therefore, it is recommended to use Microsoft Server 2008, 2012, and 2016 as the client of HBlock and configure MPIO. Windows 7, 8, and 10 do not support MPIO, and it is not recommended to use this system as the HBlock client.

(I) Install Native MPIO software.

- Windows Server 2008 R2
 1. Open **Server Manager Management Console**.
 2. Go to **Features > Features Summary > Add Features**, and open **Add Features Wizard**.
 3. Click **Next**, select **Multipath I/O**, and install.
 4. Restart Windows.

- Windows Server 2012 or 2016

1. Open **Server Manager** and select **Add roles and features**.
2. Click **Next** and check **Multipath I/O** in the **Features**.
3. Click **Next** and select **Restart the destination server automatically** if required.
4. Install, click **Yes**.

(II) Open the MPIO tool to add a storage array.

Note: When using iSCSI MPIO in Windows, view and set the MPIO load-balancing policy via **Device Manager**. The iSCSI initiator doesn't use the MPIO_DSM_Path_V2 WMI class for status display. Thus, after setting the MPIO load-balancing policy through the iSCSI initiator, the displayed status may not meet expectations.

1. Click **MPIO** in **Administrative Tools**.
2. Click **Discover Multi-Paths**, select **Add support for iSCSI device**, and then click **Add**, select **Yes** in the dialog box.
3. Restart Windows.

(III) Adjust MPIO configuration

1. Open **Powershell** and enable path detection and custom path recovery functions.

```
Get-MPIOSetting      # View current configuration
Set-MPIOSetting -NewPathVerificationState Enabled      # Enable path verification
Set-MPIOSetting -CustomPathRecovery Enabled          # Enable custom path recovery function
```

2. Restart Windows.

(IV) Run iSCSI initiator

1. Run the **iSCSI Initiator Properties** on the Windows Client. Enter iSCSI in **Start > Search Programs and Files** to open the iSCSI initiator.
2. Configure iSCSI initiator **In Discovery > Discovery Portal**.
Enter the server IP and port in **Discovery Target Portal of Discovery**. You can use the command **./stor lun ls** on the server to query the ACTIVE target and STANDBY target of the LUN.

```
./stor lun ls #View the server IP and Port of the target corresponding to the LUN
```

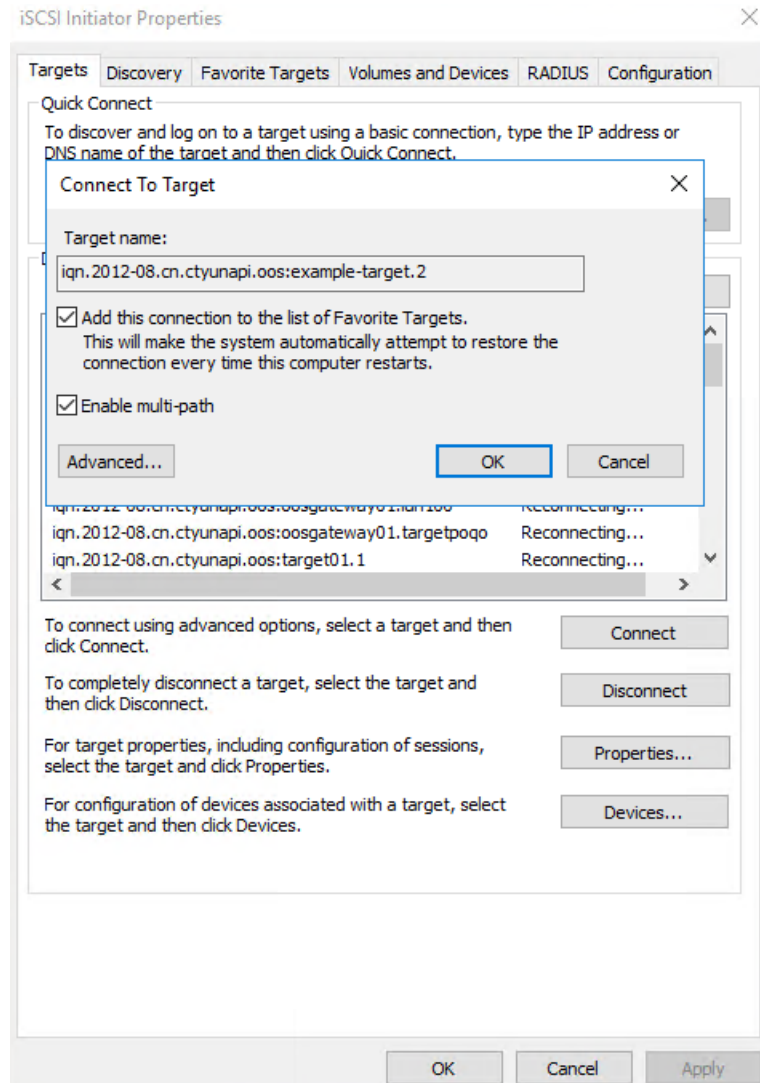
3. Search for the iSCSI target

Search iSCSI target issued by HBlock in **Discovered targets of Target**, check that the status is **Inactive**, click Connect, select **Enable multi-path**, and then click **OK**.

Note: You need to connect to the ACTIVE target first, and then connect to the STANDBY target.

Note:

- Windows Server 2012 or 2016: The same target can correspond to multiple LUNs. When one target can correspond to multiple LUNs, if the ACTIVE target and STANDBY target corresponding to different LUNs are different, it will take a while for the iSCSI connection to recognize all the LUNs. Therefore, it is recommended that each target corresponds to one LUN.
- Windows Server 2008: One target can only correspond to one LUN, and the iSCSI connection established first must be an ACTIVE target, and then the STANDBY target connection is established, otherwise the MPIO device cannot operate normally.



4. Enable CHAP authentication (Please ignore this step and direct connect if CHAP authentication is disabled)

If your iSCSI target has enabled CHAP authentication, click **Advanced** in the pop-up dialog box and check **Connect to Target**. Select **Enable CHAP log on**, enter the same name and CHAP password that was configured on the target for this initiator, and then click **OK**.

5. Using iSCSI shared disks on the client.

Go to **Server Manager > File and Storage Services > Volumes > Disks** and connect the disk with status **off-line**. Then **Initialize** and **New Volume**, specify the disk letter and format the disk. Start the **This PC**, you can see the new added disk and its capacity, and since then, you can use the iSCSI disk issued by HBlock as a local disk.

Note:

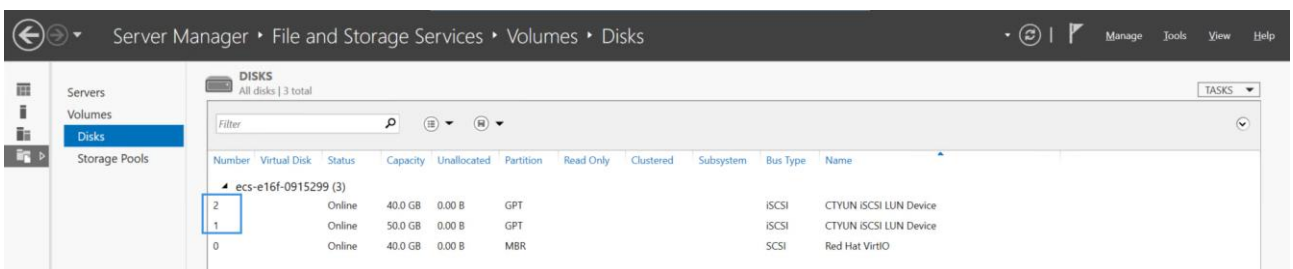
- During initialization of disk, if the storage capacity of LUN is less than or equal to 2TiB, either MBR or GPT can be used for partition; if the storage capacity of LUN is larger than 2TiB, only GPT can be used for partition.
- If client needs to be disconnected or delete the disk, you need to enter **Server Manager > File and Storage Services > Volumes > Disks**, right click on the disk and click **Take Offline** to go offline the disk. Then disconnect the iSCSI target on **iSCSI Initiator**.
- If the client needs to disconnect and then re-connect, there is no need to initialize or create a new volume, and the disk can be seen after reconnecting.

Note: If you want to query the disk corresponding to HBlock LUN, you can enter the following command line on the client side to query.

```
wmic diskdrive get Name, Manufacturer, Model, InterfaceType, MediaType, SerialNumber
```

As shown in the following example, the query information "Name" column corresponds to the "Number" column on the "DISKS". SerialName corresponds to the name and uuid of HBlock LUN.

```
C:\Users\Administrator>wmic diskdrive get Name, Manufacturer, Model, InterfaceType, MediaType, SerialNumber
InterfaceType  Manufacturer      MediaType          Model              Name              SerialNumber
SCSI           (Standard disk drives) Fixed hard disk media Red Hat VirtIO SCSI Disk Device \\.\PHYSICALDRIVE0 e3057108-d9d5-492d-9
SCSI           (Standard disk drives) Fixed hard disk media CTYUN              \\.\PHYSICALDRIVE1 lun01a-lun-uuid-1432bc36-3b5b-40bc-b74f-787722709399
SCSI           (Standard disk drives) Fixed hard disk media CTYUN              \\.\PHYSICALDRIVE2 lun02a-lun-uuid-5d2b6c71-2802-44ba-b6d3-3e1d1a39730d
```



3.3 Linux Client – Standalone Mode

3.3.1 Client Configuration

Prerequisites

- The LUNs have been successfully created on HBlock server side.
- Prepare the Linux Client.

Note: Root permission is required to configure the initiator.

If client is CentOS/RHEL, please install iscsi-initiator-utils. The installation command is as follows:

```
yum -y install iscsi-initiator-utils
```

Note: Install iSCSI initiator version 6.2.0-874-10 or above.

If client is Ubuntu/Debian, the installation command is as follows:

```
apt install open-iscsi
```

Steps

- **HBlock server side:** Query the detailed information of the LUN to be connected and its corresponding iSCSI target.

```
./stor lun ls { -n | --name } LUN_NAME  
./stor target ls { -n | --name } TARGET_NAME
```

- **Linux Client**

(I) Discover HBlock's target.

```
iscsiadm -m discovery -t st -p SERVER_IP
```

(II) Log in to iSCSI storage.

If your iSCSI target does not enable CHAP authentication, please directly proceed to step 4

Login to the target.

1. Enable authentication

```
iscsiadm -m node -T iSCSI_TARGET_IQN -o update --name node.session.auth.authmethod --value=CHAP
```

2. Enter CHAP username

```
iscsiadm -m node -T iSCSI_TARGET_IQN -o update --name node.session.auth.username --value=USER
```

3. Enter CHAP password

```
iscsiadm -m node -T iSCSI_TARGET_IQN -o update --name node.session.auth.password --value=PASSWORD
```

4. Login to the target

```
iscsiadm -m node -T iSCSI_TARGET_IQN -p SERVER_IP -l
```

(III) Display the session status and view the current iSCSI connection.

```
iscsiadm -m session
```

(IV) View the iSCSI disk and its corresponding LUN. You can view the wwid of the LUN by using the command `lsscsi -i`.

```
lsblk  
lsscsi -i
```

(V) Format the iSCSI disk.

Note: If the disk has been connected and formatted before, you can directly mount the iSCSI disk without formatting the disk after reconnection.

Use the following command to format the newly added iSCSI disk partition.

```
mkfs.ext4 /dev/sdX
```

or

```
mkfs.xfs /dev/sdX
```

Note: Common file systems include ext4 and XFS. It is determined based on user need.

(VI) Mount iSCSI disk

Mount the iSCSI disk partition to a local directory, and data can be written after mounting.

```
mount /dev/sdX PATH # PATH is the disk path
```

Note: If the user needs to disconnect or delete the disk, perform the following steps:

1. Ensure that no processes are using the folder of the file system before uninstalling it.
2. Use the command `sync` to ensure that all pending write operations have been written to disk.
3. Use the command `umount` to properly uninstall the file system and disconnect the iSCSI connection.

```
umount DIRECTORY_NAME_OR_PATH  
iscsiadm -m node -T iSCSI_TARGET_IQN -p SERVER_IP -u
```

Example:

```
[root@client ~]# sync  
[root@client ~]# umount /mnt/disk_sda  
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target1.1 -p  
192.168.0.32 -u  
Logging out of session [sid: 1, target: iqn.2012-08.cn.ctyunapi.oos:target1.1, portal:  
192.168.0.32,3260]  
Logout of [sid: 1, target: iqn.2012-08.cn.ctyunapi.oos:target1.1, portal:  
192.168.0.32,3260] successful.
```

3.3.2 Examples

Scenarios

- Linux clients need to connect to LUNs of HBlock standalone mode.
- The LUNs of HBlock standalone mode that need to be connected are lund1 and lunf1, with lund1 having CHAP authentication.

Prerequisites

- For client that need to connect to the HBlock LUNs, preparations have been made according to **Client Configuration**.
- The lund1 and f1 have been successfully created on HBlock server.

Steps

- **HBlock server side:** Query the detailed information of the LUN to be connected and its corresponding iSCSI target.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lund1
LUN Name: lund1 (LUN 0)
Storage Mode: Local
Capacity: 500 GiB
Status: Normal
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targetd.4(192.168.0.32:3260,Active)
Create Time: 2024-05-21 10:00:34
Local Sector Size: 4096 Bytes
Data Health: 100% normal, 0% low redundancy, 0% error
Write Policy: WriteBack
WWID: 33000000068f2f320
UUID: lun-uuid-3ddcc779-bf34-42b9-ac5e-0339dae28821
Path: /mnt/storage01
Snapshot Count: 0
Snapshot Size: 0 B (Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.)
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -n targetd
Target Name: targetd
Max Sessions: 2
Create Time: 2024-05-21 09:59:12
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targetd.4(192.168.0.32:3260)
LUN: lund1(LUN 0)
Reclaim Policy: Retain
CHAP: testd,T12345678912,Enabled
```

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lunf1
LUN Name: lunf1 (LUN 0)
Storage Mode: Local
Capacity: 600 GiB
Status: Normal
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targetf.5(192.168.0.32:3260,Active)
Create Time: 2024-05-21 10:00:56
Local Sector Size: 4096 Bytes
Data Health: 100% normal, 0% low redundancy, 0% error
Write Policy: WriteBack
WWID: 33000000030f798a5
UUID: lun-uuid-7b7f91d8-b75e-4de2-ac69-621e4be7a0cf
Path: /mnt/storage01
Snapshot Count: 0
Snapshot Size: 0 B (Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.)
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -n targetf
Target Name: targetf
Max Sessions: 2
Create Time: 2024-05-21 10:00:15
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targetf.5(192.168.0.32:3260)
LUN: lunf1(LUN 0)
Reclaim Policy: Retain
```

● Linux Client

(I) Discover targets of lund1 and lunf1.

```
[root@client ~]# iscsiadm -m discovery -t st -p 192.168.0.32
192.168.0.32:3260,1 iqn.2012-08.cn.ctyunapi.oos:targetf.5
192.168.0.32:3260,1 iqn.2012-08.cn.ctyunapi.oos:targetd.4
192.168.0.32:3260,1 iqn.2012-08.cn.ctyunapi.oos:targetc.3
```

(II) Log in to iSCSI storage.

- Log in to iSCSI storage of lund1, and CHAP authentication is required.

```
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:targetd.4 -o update --name node.session.auth.authmethod --value=CHAP
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:targetd.4 -o update --name node.session.auth.username --value=testd
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:targetd.4 -o update --name node.session.auth.password --value=*****
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:targetd.4 -p 192.168.0.32:3260 -l
Logging in to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:targetd.4, portal: 192.168.0.32,3260] (multiple)
Login to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:targetd.4, portal: 192.168.0.32,3260] successful.
```

- Log in to iSCSI storage of lunf1.

```
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:targetf.5 -p
192.168.0.32:3260 -l
Logging in to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:targetf.5, portal:
192.168.0.32,3260] (multiple)
Login to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:targetf.5, portal:
192.168.0.32,3260] successful.
```

(III) Display the session status and view the current iSCSI connection.

```
[root@client ~]# iscsiadm -m session
tcp: [1] 192.168.0.32:3260,1 iqn.2012-08.cn.ctyunapi.oos:targetd.4 (non-flash)
tcp: [2] 192.168.0.32:3260,1 iqn.2012-08.cn.ctyunapi.oos:targetf.5 (non-flash)
```

(IV) View the iSCSI disk and its corresponding LUN. You can view the wwid of the LUN by using the command `lsscsi -i`.

```
[root@client ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 500G 0 disk
sdb 8:16 0 600G 0 disk
vda 253:0 0 40G 0 disk
├─vda1 253:1 0 4G 0 part
└─vda2 253:2 0 36G 0 part /
vdb 253:16 0 100G 0 disk
└─vdb1 253:17 0 100G 0 part /mnt/storage01
vdc 253:32 0 100G 0 disk
vdd 253:48 0 100G 0 disk
[root@client ~]# lsscsi -i
[2:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sda 33000000068f2f320
[3:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sdb 33000000030f798a5
```

Note: It can be seen that `/dev/sda` corresponds to HBlock LUN lund1 (LUN WWID 33000000068f2F320), and `/dev/sdb` corresponds to HBlock LUN lunf1 (LUN WWID 33000000030f798a5).

(V) Format the iSCSI disk.

```
[root@client ~]# mkfs.ext4 /dev/sda
mke2fs 1.42.9 (28-Dec-2013)
/dev/sda is entire device, not just one partition!
Proceed anyway? (y,n) y
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
32768000 inodes, 131072000 blocks
6553600 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2279604224
```

```
4000 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, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

[root@client ~]# mkfs.ext4 /dev/sdb
mke2fs 1.42.9 (28-Dec-2013)
/dev/sdb is entire device, not just one partition!
Proceed anyway? (y,n) y
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
39321600 inodes, 157286400 blocks
7864320 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2304770048
4800 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, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

(VI) Mount iSCSI disk

Mount the iSCSI disk partition to a local directory, and data can be written after mounting.

```
[root@client ~]# mount /dev/sda /mnt/disk_sda
[root@client ~]# mount /dev/sdb /mnt/disk_sdb
[root@client ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda   8:0    0 500G 0 disk /mnt/disk_sda
sdb   8:16   0 600G 0 disk /mnt/disk_sdb
vda   253:0   0  40G 0 disk
```

```
|_vda1 253:1    0    4G  0 part
|_vda2 253:2    0   36G 0 part /
vdb   253:16   0  100G 0 disk
|_vdb1 253:17  0  100G 0 part /mnt/storage01
vdc   253:32   0  100G 0 disk
vdd   253:48   0  100G 0 disk
```

3.4 Linux Client – Cluster Mode

3.4.1 Client Configuration

- The LUNs have been successfully created on HBlock server side.
- Prepare the Linux Client

(I) Install the Linux Client.

Note: Root permission is required to configure the initiator.

If client is CentOS/RHEL, please install iscsi-initiator-utils. The installation command is as follows:

```
yum -y install iscsi-initiator-utils
```

Note: Please install iSCSI initiator 6.2.0-874-10 or above.

If client is Ubuntu/Debian, the installation command is as follows:

```
apt install open-iscsi
```

(II) Install MPIO.

- For CentOS

```
yum install device-mapper-multipath device-mapper-multipath-libs
```

- For Ubuntu

```
apt install multipath-tools
```

(III) Configure MPIO.

1. Copy **/usr/share/doc/device-mapper-multipath-X.Y.Z/multipath.conf** to **/etc/multipath.conf**, where X.Y.Z is the actual version number of multipath.
2. Add the following configuration in **/etc/multipath.conf**:

Note: In multipath.conf, if the multipath section has the same parameters as the devices section, the parameter values in multipath override those in devices. In order to use HBlock volumes properly, you must delete the same parameters as the following fields in multipath.

```
defaults {
    user_friendly_names yes
    find_multipaths yes
    uid_attribute "ID_WWN"
}
devices {
    device {
        vendor "CTYUN"
        product "iSCSI LUN Device"
```

```

    path_grouping_policy failover
    path_checker tur
    path_selector "round-robin 0"
    hardware_handler "1 alua"
    rr_weight priorities
    no_path_retry queue
    prio alua
}
}

```

Note: user_friendly_names can be set to yes or no.

- user_friendly_names yes: The system assigns aliases to multipath devices using settings in /etc/multipath/bindings, in the format mpath*n* (e.g., mpatha, mpathb).
- user_friendly_names no: The system uses the WWID as the alias for multipath devices.

(IV) Restart the multipathd service.

- For CentOS

```

systemctl restart multipathd
systemctl enable multipathd

```

- For Ubuntu

```

systemctl restart multipath-tools.service
systemctl enable multipath-tools.service

```

Steps

- **HBlock server side:** Query the detailed information of the LUN to be connected and its corresponding iSCSI target.

```

./stor lun ls { -n | --name } LUN_NAME
./stor target ls { -n | --name } TARGET_NAME

```

- **Linux Client**

(I) Discover HBlock's target.

Note: If the LUN corresponds to multiple target IQNs, it is recommended to connect all of them.

```

iscsiadm -m discovery -t st -p ACTIVE_IP
iscsiadm -m discovery -t st -p STANDBY_IP
iscsiadm -m discovery -t st -p ColdStandby_IP

```

(II) Log in to iSCSI storage: establish multiple iSCSI connections (including target IQN with status Active, Standby, and ColdStandby).

Note: If your iSCSI target does not enable CHAP authentication, please directly proceed to step 4 Login to the target.

1. Enable authentication.

```
iscsiadm -m node -T iSCSI_TARGET_IQN -o update --name node.session.auth.authmethod --value=CHAP
```

2. Enter the CHAP username.

```
iscsiadm -m node -T iSCSI_TARGET_IQN -o update --name node.session.auth.username --value=USER
```

3. Enter the CHAP password.

```
iscsiadm -m node -T iSCSI_TARGET_IQN -o update --name node.session.auth.password --value=PASSWORD
```

4. Login to the target

```
iscsiadm -m node -T iSCSI_TARGET_IQN -p SERVER_IP -l
```

(III) Display the session status and view the current iSCSI connection.

```
iscsiadm -m session  
lsscsi
```

(IV) View MPIO devices, the iSCSI disk's corresponding LUN WWID.

```
multipath -ll # You can add parameter -v 3 to display more detailed information  
ll /dev/mapper/mpathX  
/lib/udev/scsi_id --whitelisted --device=/dev/sdX # View the iSCSI disk's corresponding LUN WWID
```

(V) Operate MPIO devices.

Mount the iSCSI disk partition to a local directory, and data can be written after mounting.

Note: If the disk has been connected and formatted before, you can directly mount the iSCSI disk without formatting the disk after reconnection.

```
lsblk  
mkfs -t ext4 /dev/mapper/mpathX # Formatted into ext4  
mkdir DIRECTORY_NAME_OR_PATH #Create directory  
mount /dev/mapper/mpath X DIRECTORY_NAME_OR_PATH #Mount mpath X to the directory  
lsblk
```

Note: Common file systems include ext4 and XFS. It is determined based on user need.

Note: If the user needs to disconnect or delete the disk, perform the following steps:

1. Ensure that no processes are using the folder of the file system before uninstalling it.
2. Use the command `sync` to ensure that all pending write operations have been written to disk.

3. Use the command `umount` to properly uninstall the file system and disconnect the iSCSI connection.

```
umount DIRECTORY_NAME_OR_PATH  
iscsiadm -m node -T iSCSI_TARGET_IQN -p SERVER_IP -u
```

Example:

```
[root@client ~]# sync  
[root@client ~]# umount /mnt/disk_mpatha  
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target01.1 -p  
192.168.0.102 -u  
Logging out of session [sid: 2, target: iqn.2012-08.cn.ctyunapi.oos:target01.1, portal:  
192.168.0.102,3260]  
Logout of [sid: 2, target: iqn.2012-08.cn.ctyunapi.oos:target01.1, portal:  
192.168.0.102,3260] successful.  
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target01.2 -p  
192.168.0.110 -u  
Logging out of session [sid: 3, target: iqn.2012-08.cn.ctyunapi.oos:target01.2, portal:  
192.168.0.110,3260]  
Logout of [sid: 3, target: iqn.2012-08.cn.ctyunapi.oos:target01.2, portal:  
192.168.0.110,3260] successful.  
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target01.3 -p  
192.168.0.192 -u  
Logging out of session [sid: 4, target: iqn.2012-08.cn.ctyunapi.oos:target01.3, portal:  
192.168.0.192,3260]  
Logout of [sid: 4, target: iqn.2012-08.cn.ctyunapi.oos:target01.3, portal:  
192.168.0.192,3260] successful.
```

3.4.2 Examples

Scenarios

- Linux clients need to connect to LUNs of HBlock cluster mode.
- The LUNs of HBlock cluster mode that need to be connected are lun6a and lun7a, with lun7a having CHAP authentication.

Prerequisites

- For client that need to connect to the HBlock LUNs, preparations have been made according to **Client Configuration**.
- The lun6a and 7a have been successfully created on HBlock server.

Steps

- **HBlock server side:** Query the detailed information of the LUN to be connected and its corresponding iSCSI target.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun6a
LUN Name: lun6a (LUN 0)
Storage Mode: Cache
Capacity: 500 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target6.12(192.168.0.192:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target6.11(192.168.0.110:3260,Standby)
               iqn.2012-08.cn.ctyunapi.oos:target6.13(192.168.0.102:3260,ColdStandby)
Create Time: 2024-05-21 14:14:48
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 33ffffffffc69cbabb
UUID: lun-uuid-40731bfd-d0e5-49fb-9784-1d825635daf8
Object Storage Info:
+-----+
| Provider      | OOS                |
| Bucket Name  | hblocktest3       |
```

```

| Prefix          | stor2          |
| Endpoint        | https://oos-cn.ctyunapi.cn |
| Signature Version | v2            |
| Region          |                |
| Storage Class    | STANDARD      |
| Access Key       | cb22b08b1f9229f85874 |
| Object Size      | 1024 KiB      |
| Compression      | Enabled       |
+-----+-----+
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -n target6
Target Name: target6
Max Sessions: 2
Create Time: 2024-05-21 14:12:44
Number of Servers: 3
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target6.11(192.168.0.110:3260)
               iqn.2012-08.cn.ctyunapi.oos:target6.12(192.168.0.192:3260)
               iqn.2012-08.cn.ctyunapi.oos:target6.13(192.168.0.102:3260)
LUN: lun6a(LUN 0)
Reclaim Policy: Retain
ServerID: hblock_1,hblock_2,hblock_3
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun7a
LUN Name: lun7a (LUN 0)
Storage Mode: Local
Capacity: 500 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target7.14(192.168.0.110:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target7.15(192.168.0.192:3260,Standby)
Create Time: 2024-05-21 14:15:22
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 330000000727497eb
UUID: lun-uuid-3429b79f-cd7d-47cb-9fb6-c79136deb237
Snapshot Count: 0
Snapshot Size: 0 B (Note: Snapshot size may vary due to LUN issues or parent snapshot
deletion.)
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -n target7
Target Name: target7
Max Sessions: 1
Create Time: 2024-05-21 14:13:27
Number of Servers: 2
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target7.14(192.168.0.110:3260)
               iqn.2012-08.cn.ctyunapi.oos:target7.15(192.168.0.192:3260)
LUN: lun7a(LUN 0)
Reclaim Policy: Retain

```

```
CHAP: test2,T12345678912,Enabled
ServerID: hblock_1,hblock_2
```

● Linux Client

(I) Discover targets of lun6a and lun7a.

```
[root@client ~]# iscsiadm -m discovery -t st -p 192.168.0.110
192.168.0.110:3260,1 iqn.2012-08.cn.ctyunapi.oos:target7.14
192.168.0.110:3260,1 iqn.2012-08.cn.ctyunapi.oos:target02.3
192.168.0.110:3260,1 iqn.2012-08.cn.ctyunapi.oos:target04.7
192.168.0.110:3260,1 iqn.2012-08.cn.ctyunapi.oos:target6.11
[root@client ~]# iscsiadm -m discovery -t st -p 192.168.0.192
192.168.0.192:3260,1 iqn.2012-08.cn.ctyunapi.oos:target7.15
192.168.0.192:3260,1 iqn.2012-08.cn.ctyunapi.oos:target6.12
192.168.0.192:3260,1 iqn.2012-08.cn.ctyunapi.oos:test.10
192.168.0.192:3260,1 iqn.2012-08.cn.ctyunapi.oos:target04.8
[root@client ~]# iscsiadm -m discovery -t st -p 192.168.0.102
192.168.0.102:3260,1 iqn.2012-08.cn.ctyunapi.oos:target02.4
192.168.0.102:3260,1 iqn.2012-08.cn.ctyunapi.oos:target6.13
192.168.0.102:3260,1 iqn.2012-08.cn.ctyunapi.oos:test.9
```

(II) Log in to iSCSI storage.

- Log in to iSCSI storage of lun6a, and establish multiple iSCSI connections in the order of Active target, Standby target, and ColdStandby target.

```
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target6.12 -p
192.168.0.192:3260 -l
Logging in to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target6.12, portal:
192.168.0.192,3260] (multiple)
Login to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target6.12, portal:
192.168.0.192,3260] successful.
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target6.11 -p
192.168.0.110:3260 -l
Logging in to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target6.11, portal:
192.168.0.110,3260] (multiple)
Login to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target6.11, portal:
192.168.0.110,3260] successful.
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target6.13 -p
192.168.0.102:3260 -l
Logging in to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target6.13, portal:
192.168.0.102,3260] (multiple)
Login to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target6.13, portal:
192.168.0.102,3260] successful.
```

- Log in to iSCSI storage of lun7a, and CHAP authentication is required.

```
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target7.14 -o update --
name node.session.auth.authmethod --value=CHAP
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target7.14 -o update --
name node.session.auth.username --value=test2
```

```
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target7.14 -o update --
name node.session.auth.password --value=*****
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target7.14 -p
192.168.0.110:3260 -l
Logging in to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target7.14, portal:
192.168.0.110,3260] (multiple)
Login to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target7.14, portal:
192.168.0.110,3260] successful.
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target7.15 -o update --
name node.session.auth.authmethod --value=CHAP
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target7.15 -o update --
name node.session.auth.username --value=test2
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target7.15 -o update --
name node.session.auth.password --value=*****
[root@client ~]# iscsiadm -m node -T iqn.2012-08.cn.ctyunapi.oos:target7.15 -p
192.168.0.192:3260 -l
Logging in to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target7.15, portal:
192.168.0.192,3260] (multiple)
Login to [iface: default, target: iqn.2012-08.cn.ctyunapi.oos:target7.15, portal:
192.168.0.192,3260] successful.
```

(III) Display the session status and view the current iSCSI connection.

```
[root@client ~]# iscsiadm -m session
tcp: [3] 192.168.0.192:3260,1 iqn.2012-08.cn.ctyunapi.oos:target6.12 (non-flash)
tcp: [4] 192.168.0.110:3260,1 iqn.2012-08.cn.ctyunapi.oos:target6.11 (non-flash)
tcp: [5] 192.168.0.102:3260,1 iqn.2012-08.cn.ctyunapi.oos:target6.13 (non-flash)
tcp: [6] 192.168.0.110:3260,1 iqn.2012-08.cn.ctyunapi.oos:target7.14 (non-flash)
tcp: [7] 192.168.0.192:3260,1 iqn.2012-08.cn.ctyunapi.oos:target7.15 (non-flash)
[root@client ~]# lsscsi
[4:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sdc
[5:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sdd
[6:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sde
[7:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sdf
[8:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sdg
```

(IV) View MPIO devices, the iSCSI disk's corresponding LUN WWID.

```
[root@client ~]# multipath -ll
mpathc (0x30000000727497eb) dm-1 CTYUN ,iSCSI LUN Device
size=500G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active
| ` 7:0:0:0 sdf 8:80 active ready running
`+- policy='round-robin 0' prio=1 status=enabled
  ` 8:0:0:0 sdg 8:96 active ghost running
mpathb (0x3ffffffffffc69cbabb) dm-0 CTYUN ,iSCSI LUN Device
size=500G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active
| ` 4:0:0:0 sdc 8:32 active ready running
|+- policy='round-robin 0' prio=1 status=enabled
| ` 5:0:0:0 sdd 8:48 active ghost running
```

```

`-- policy='round-robin 0' prio=0 status=enabled
  `- 6:0:0:0 sde 8:64 failed faulty running
[root@client ~]# ll /dev/mapper/mpathc
lrwxrwxrwx 1 root root 7 May 21 15:03 /dev/mapper/mpathc -> ../dm-1
[root@client ~]# ll /dev/mapper/mpathb
lrwxrwxrwx 1 root root 7 May 21 14:57 /dev/mapper/mpathb -> ../dm-0
[root@client ~]# # /lib/udev/scsi_id --whitelisted --device=/dev/sdc
33ffffffffc69cbabb
[root@client ~]# # /lib/udev/scsi_id --whitelisted --device=/dev/sdd
33ffffffffc69cbabb
[root@client ~]# # /lib/udev/scsi_id --whitelisted --device=/dev/sde
33ffffffffc69cbabb
[root@client ~]# # /lib/udev/scsi_id --whitelisted --device=/dev/sdf
330000000727497eb
[root@client ~]# # /lib/udev/scsi_id --whitelisted --device=/dev/sdg
330000000727497eb
    
```

Note: It can be seen that /dev/mapper/mpathb (/dev/sdc, /dev/sdd, /dev/sde) corresponds to HBlock LUN lun6a (LUN WWID 33ffffffffc69cbabb), and /dev/mapper/mpathc (/dev/sdf, /dev/sdg) corresponds to HBlock LUN lun7a (LUN WWID 330000000727497eb).

(V) Operate MPIO devices.

Mount the iSCSI disk partition to a local directory, and data can be written after mounting.

- Mount iSCSI disk /dev/mapper/mpathb.

```

[root@client ~]# lsblk
sdc          8:32    0 500G  0 disk
└─mpathb    252:0    0 500G  0 mpath
sdd          8:48    0 500G  0 disk
└─mpathb    252:0    0 500G  0 mpath
sde          8:64    0 500G  0 disk
└─mpathb    252:0    0 500G  0 mpath
sdf          8:80    0 500G  0 disk
└─mpathc    252:1    0 500G  0 mpath
sdg          8:96    0 500G  0 disk
└─mpathc    252:1    0 500G  0 mpath
vda         253:0    0   40G  0 disk
├─vda1     253:1    0    4G  0 part
└─vda2     253:2    0   36G  0 part  /
vdb         253:16   0  100G  0 disk
└─vdb1     253:17   0  100G  0 part  /mnt/storage01
vdc         253:32   0  100G  0 disk
vdd         253:48   0  100G  0 disk
[root@client ~]# mkfs -t ext4 /dev/mapper/mpathb
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
32768000 inodes, 131072000 blocks
6553600 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2279604224
4000 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, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

[root@client ~]# mkdir /mnt/disk_mpathb
[root@client ~]# mount /dev/mapper/mpathb /mnt/disk_mpathb
[root@client ~]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sdc          8:32   0 500G  0 disk
└─mpathb 252:0   0 500G  0 mpath /mnt/disk_mpathb
sdd          8:48   0 500G  0 disk
└─mpathb 252:0   0 500G  0 mpath /mnt/disk_mpathb
sde          8:64   0 500G  0 disk
└─mpathb 252:0   0 500G  0 mpath /mnt/disk_mpathb
sdf          8:80   0 500G  0 disk
└─mpathc 252:1   0 500G  0 mpath
sdg          8:96   0 500G  0 disk
└─mpathc 252:1   0 500G  0 mpath
vda         253:0   0   4G  0 disk
├─vda1    253:1   0    4G  0 part
└─vda2    253:2   0   36G  0 part /
vdb         253:16  0 100G  0 disk
└─vdb1    253:17  0 100G  0 part /mnt/storage01
vdc         253:32  0 100G  0 disk
vdd         253:48  0 100G  0 disk

```

- Mount iSCSI disk /dev/mapper/mpathc.

```

[root@client ~]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sdc          8:32   0 500G  0 disk
└─mpathb 252:0   0 500G  0 mpath /mnt/disk_mpathb
sdd          8:48   0 500G  0 disk
└─mpathb 252:0   0 500G  0 mpath /mnt/disk_mpathb
sde          8:64   0 500G  0 disk
└─mpathb 252:0   0 500G  0 mpath /mnt/disk_mpathb
sdf          8:80   0 500G  0 disk

```

```

└─mpathc 252:1    0 500G 0 mpath
sdg      8:96    0 500G 0 disk
└─mpathc 252:1    0 500G 0 mpath
vda     253:0    0  40G 0 disk
├─vda1  253:1    0   4G 0 part
└─vda2  253:2    0  36G 0 part /
vdb     253:16   0 100G 0 disk
└─vdb1  253:17   0 100G 0 part /mnt/storage01
vdc     253:32   0 100G 0 disk
vdd     253:48   0 100G 0 disk
[root@client ~]# mkfs -t ext4 /dev/mapper/mpathc
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
32768000 inodes, 131072000 blocks
6553600 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2279604224
4000 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, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

[root@client ~]# mkdir /mnt/disk_mpathc
[root@client ~]# mount /dev/mapper/mpathc /mnt/disk_mpathc
[root@client ~]# lsblk
NAME      MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sdc       8:32    0 500G  0 disk
└─mpathb 252:0    0 500G  0 mpath /mnt/disk_mpathb
sdd       8:48    0 500G  0 disk
└─mpathb 252:0    0 500G  0 mpath /mnt/disk_mpathb
sde       8:64    0 500G  0 disk
└─mpathb 252:0    0 500G  0 mpath /mnt/disk_mpathb
sdf       8:80    0 500G  0 disk
└─mpathc 252:1    0 500G  0 mpath /mnt/disk_mpathc
sdg       8:96    0 500G  0 disk
└─mpathc 252:1    0 500G  0 mpath /mnt/disk_mpathc
vda       253:0    0  40G  0 disk
├─vda1   253:1    0   4G  0 part
└─vda2   253:2    0  36G  0 part /

```

vdb	253:16	0	100G	0	disk	
└vdb1	253:17	0	100G	0	part	/mnt/storage01
vdc	253:32	0	100G	0	disk	
vdd	253:48	0	100G	0	disk	

4 Management Operations

4.1 Management Operation Command Line Format

The following command line formats may appear in this manual, and their meanings are as follows:

Format	Description
Bold	Command line keywords.
<i>Italic</i>	Command line variable parameters which need to be replaced with actual values.
[]	Indicates that the [] part is optional during configuration.
{A B ...}	Indicates that one must be selected from keyword or parameter in {}.
[A B ...]	Indicates that one could be selected from keyword or parameter in [], or not selected.
&<1-n>	Indicates that the parameter before the symbol & can be repeated 1-n times, separated by commas (,).

Note: The subsequent command line uses the x86 server as an example. The commands of the ARM server are the same as those of the x86 server.

4.2 Help Command

```
./stor { --help | -h }
```

This command is used to query the command help information of HBlock.

Examples

- Query all commands help information of HBlock.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor -help
Usage:
  stor <command> [options]
  stor <command> <subcommand> [options]

Type 'stor --help <command>' or 'stor --help <command> <subcommand>' to print help on a
specific subcommand.
Type 'stor --version' to print HBlock version.

Available commands:
  alarm          View and manage alarms.
  backup         Manage backup(s).
  config        Configure HBlock.
  conssnap      Manage consistency snapshot(s).
  event         Event information.
  info          Get HBlock information.
  install       Install HBlock.
  license       Manage HBlock license(s).
  logcollect   Log collection information.
  lun           Manage LUN(s).
  monitor       Get performance data.
  qos           Manage QoS policy.
  restart      Restart HBlock.
  server       Manage HBlock server(s).
  setup        Initialize HBlock.
  snapshot     Manage snapshot(s).
  start        Start HBlock.
  status       Get HBlock status.
  stop         Stop HBlock.
  storagepool  Manage storage pools.
  target       Manage target(s).
  topology     Manage topology.
  tuning       Adjusting parameters to fit different scenarios.
  uninstall    Uninstall HBlock on all servers.
  upgrade      Upgrade HBlock.

General subcommands:
  add          Add an entity, such as adding a LUN or server.
  ls           List entities or the specified entity, such as listing
servers or the specified server.
  rm           Remove an entity, such as removing a LUN or server.
```

```
set                Set entity properties, such as setting LUN or server
properties.

Available subcommands for alarm:
  E[export]        Export alarm records to a file.
  M[mute]          Set the mute status of a specified alarm record to Muted
status, stop receiving its alarm emails and not list it by default.
  UM[unmute]       Set the mute status of a specified alarm record to Normal
status.
  R[resolve]       Set a specified alarm record to Resolved status.

Available subcommands for backup:
  E[export]        Export backup file.
  I[import]        Import backup file to a LUN.

Available subcommands for conssnap:
  R[rollback]      Rollback consistency snapshot.

Available subcommands for lun:
  C[clone]         Create clone LUN.
  F[flatten]       Flatten clone LUN.
  RC[recover]      Recover the specified LUN from the cloud.
  RS[resume]       Resume the specified LUN.
  SUS[suspend]     Suspend the specified LUN.
  S[switch]        Switch between active and standby iSCSI targets.
  W[wipe]          Wipe the specified LUN.
  X[expand]        Expand LUN.

Available subcommands for monitor:
  E[export]        Export historical performance data to a file.
  V[view]          Get real-time performance data.

Available subcommands for qos:
  A[assoc]         Associate the QoS policy with the specified objects.
  D[disass]        Disassociate the QoS policy with the specified objects.

Available subcommands for server:
  A[addpath]       Add path(s) to store data.
  R[rmpath]        Remove path(s).
  S[setpath]       Set path(s) properties.

Available subcommands for snapshot:
  R[rollback]      Rollback snapshot or a snapshot in consistency snapshot.

Available subcommands for storagepool:
  A[addnode]       Add node(s) to storage pool.
  R[rmnode]        Remove node(s) from storage pool.

Available subcommands for target:
  R[rmallow]       Remove iSCSI target allowlist.
  S[setallow]      Set iSCSI target allowlist.
```

```
Available subcommands for upgrade:  
T[status]           Get upgrade status.
```

- Query help information of LUN operations.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun --help  
lun: Manage LUN(s).  
  
Usage: lun <subcommand> [options]  
Type 'stor --help lun <subcommand>' to print help on a specific subcommand.  
  
Available subcommands:  
add                Create LUN.  
C[clone]           Create clone LUN.  
F[flatten]         Flatten clone LUN.  
ls                 List all iSCSI LUNs or the specified LUN.  
RC[recover]        Recover the specified LUN from the cloud.  
RS[resume]         Resume the specified LUN.  
prefer             Set server affinity of active and standby iSCSI  
targets for LUN.  
rm                 Remove LUN.  
SUS[suspend]       Suspend the specified LUN.  
S[switch]          Switch between active and standby iSCSI targets.  
set                Set LUN properties.  
W[wipe]            Wipe the specified LUN.  
X[expand]          Expand LUN.
```

4.3 Install

```
./stor install [ { -a | --api-port } API_PORT ] [ { -w | --web-port } WEB_PORT ]
```

This command is used to install HBlock.

Note:

- Please ensure that the Linux user has permission for the required ports. By default, Linux systems do not open ports less than 1024 to ordinary users without root privileges.
- The installation of HBlock and the execution of HBlock management operations should belong to the same user.
- If the server does not have the font library installed, a corresponding prompt message will be displayed after executing the installation command. In addition, when the server lacks the font library, logging in via the WEB will result in the verification code not being displayed, thus preventing the login operation from being completed.

Parameters

Parameter	Description
<code>-a API_PORT</code> or <code>--api-port API_PORT</code>	Specifies the API port. The value is an integer that ranges from 1 to 65535. The default value is 1443.
<code>-w WEB_PORT</code> or <code>--web-port WEB_PORT</code>	Specifies the web port. The value is an integer that ranges from 1 to 65535. The default value is 2443.

Examples:

Install HBlock

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor install
Do you agree with HBlock User Agreement? [Yes/No]
Used in Chinese mainland, follow https://www.ctyun.cn/portal/protocol/10073150
Otherwise, follow https://www.esurfingcloud.com/portal/protocol/20692906
y
Installing HBlock...
Installed successfully.
When all servers are installed, please initialize HBlock in any of the following ways:
1. Use web portal to initialize HBlock. The https port is 2443.
2. Use management API (POST /rest/v1/system/setup) to initialize HBlock. The https port is 1443.
3. Use command line (stor setup) to initialize HBlock. Type 'stor --help setup' for more information.
```

4.4 Initialize

Standalone Mode:

```
./stor setup { -n | --stor-name } STOR_NAME [ { -u | --user-name } USER_NAME ] { -p | --password } PASSWORD { -s | --server } { SERVER_IP[:PORT]:PATH &<1-n> } [ --edition EDITION ] [ { -P | --public-network } CIDR ] [ --iscsi-port ISCSI_PORT ] [--port-range PORT1-PORT2 ] [ --management-port1 MANAGEMENT_PORT1 ] [ --management-port2 MANAGEMENT_PORT2 ] [ --management-port3 MANAGEMENT_PORT3 ] [ --management-port4 MANAGEMENT_PORT4 ] [ --management-port6 MANAGEMENT_PORT6 ]
```

Cluster Mode:

```
./stor setup { -n | --stor-name } STOR_NAME [ { -u | --user-name } USER_NAME ] { -p | --password } PASSWORD { { -s | --server } { SERVER_IP[:PORT][:PATH] &<1-n> } &<1-n> | --topology-file TOPOLOGY_FILE } [ --edition EDITION ] [ { -C | --cluster-network } CIDR ] [ { -P | --public-network } CIDR ] [ --fault-domain FAULT_DOMAIN ] [ --iscsi-port ISCSI_PORT ] [--port-range PORT1-PORT2 ] [ --data-port1 DATA_PORT1 ] [ --management-port1 MANAGEMENT_PORT1 ] [ --management-port2 MANAGEMENT_PORT2 ] [ --management-port3 MANAGEMENT_PORT3 ] [ --management-port4 MANAGEMENT_PORT4 ] [ --management-port5 MANAGEMENT_PORT5 ] [ --management-port6 MANAGEMENT_PORT6 ] [ --metadata-port1 METADATA_PORT1 ] [ --metadata-port2 METADATA_PORT2 ] [ --metadata-port3 METADATA_PORT3 ] [ --metadata-port4 METADATA_PORT4 ] [ --metadata-port5 METADATA_PORT5 ] [ --metadata-port6 METADATA_PORT6 ] [ --metadata-port7 METADATA_PORT7 ] [ --metadata-port8 METADATA_PORT8 ] [ --cs SERVER_IP[:DIR], SERVER_IP[:DIR], SERVER_IP[:DIR] ] [ --mdm SERVER_IP[:DIR], SERVER_IP[:DIR] ] [ --ls SERVER_IP[:DIR], SERVER_IP[:DIR], SERVER_IP[:DIR] ]
```

This command is used to initialize HBlock.

Note: You can perform this operation on any server in the cluster after installing HBlock on each server.

Note:

- Please ensure that the Linux user has permission for the required ports. By default, Linux systems do not open ports less than 1024 to ordinary users without root privileges.
- When setting the port range (`--port-range PORT1-PORT2`), please avoid overlapping with the local temporary port (`ip_local_port_range`) range of the Linux system, otherwise, the port used by the HBlock service may be occupied. Run the command `cat`

/proc/sys/net/ipv4/ip_local_port_range to query the local temporary port range.

Parameters

Parameter	Description
-n <i>STOR_NAME</i> or --stor-name <i>STOR_NAME</i>	Specifies the HBlock name. The value is a string of 1 to 64 case-sensitive characters. It can contain letters, digits, underscores (_), or hyphens (-). Only supports starting with a letter or a digit.
-u <i>USER_NAME</i> or --user-name <i>USER_NAME</i>	Specifies HBlock Administrator Username. The value is a string of 5 to 16 case-sensitive characters. It can only contain letters or digits. The default value is storuser.
-p <i>PASSWORD</i> or --password <i>PASSWORD</i>	Specifies the password for admin user. The password must be set when initializing HBlock. The value is a string of 8 to 16 case-sensitive characters: <ul style="list-style-type: none"> ● The password must contain at least three of the following: lowercase letters, uppercase letters, digits, or special characters. The special characters only include ~ ! @ # \$ % ^ & * () _ + [] { } ; : , . / < > ?. ● The password cannot contain t 3 consecutive repeating characters, 3 consecutive or in-reverse order of numbers or letters (case-insensitive), 3 consecutive or in-reverse order of keyboard sequences (case-insensitive).
-s { <i>SERVER_IP[:PORT] [:PATH]</i> &<1-n> } &<1-n> or --server { <i>SERVER_IP[:PORT] [:PATH]</i> &<1-n> } &<1-n>	Specifies the server IP, API port and disk path when initializing HBlock. <ul style="list-style-type: none"> ● server IP: <ul style="list-style-type: none"> ■ For standalone mode, only one sever is required. ■ For cluster mode, at least three servers are required. ● API port: The value is an integer that ranges from 1 to 65535, the default value is 1443. It needs to be consistent with the API port number set when installing HBlock on the server. ● disk path: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; : . : Multiple disk paths can be set, separated by comma (,). The disk path is used to

	<p>store data. It is not recommended to share a disk or file system with the operating system.</p> <p>Note: For standalone mode, the first disk path is the default disk path. For the cluster mode, you must have at least one disk path.</p>
<p>--topology-file <i>TOPOLOGY_FILE</i></p>	<p>Import the cluster topology file (only supported by cluster mode).</p> <p>The topology file is a JSON file in UTF-8 encoding format. For details, see Cluster Topology File.</p>
<p>--edition <i>EDITION</i></p>	<p>Specifies the edition type.</p> <p>Value:</p> <ul style="list-style-type: none"> ● Free: Free edition. ● Commercial (Comm): Commercial edition. <p>The default value is Commercial.</p>
<p>-C <i>CIDR</i> or --cluster-network <i>CIDR</i></p>	<p>Cluster network for data communication between clusters, only supported by cluster mode. The format is IP CIDR format.</p> <ul style="list-style-type: none"> ● If a cluster network is specified, please ensure that each server has an IP that matches the specified network segment, and the system will automatically use this IP for communication. ● If the cluster network is not specified, the <i>SERVER_IP</i> will be used by default. In this case, <i>SERVER_IP</i> cannot be specified as localhost, 127.0.0.1 or 0:0:0:0:0:0:1.
<p>-P <i>CIDR</i> or --public-network <i>CIDR</i></p>	<p>Public network for data transmission between client and server. The format is IP CIDR format.</p> <ul style="list-style-type: none"> ● If the public network is specified, please ensure that each server has an IP that matches the specified network segment, and the system will automatically use this IP to communicate with the client. ● If the public network is not specified or the specified network segment does not match any IP of the server, the <i>SERVER_IP</i> is used by default. In this case, <i>SERVER_IP</i> cannot be specified as localhost, 127.0.0.1 or 0:0:0:0:0:0:1.
<p>--fault-domain <i>FAULT_DOMAIN</i></p>	<p>Sets the fault domain level of the base storage pool (only supported by cluster mode). All disk path nodes in the cluster topology configured during initialization are added to the base storage pool.</p> <p>Value:</p> <ul style="list-style-type: none"> ● path: disk path level. ● server: server level. ● rack: rack level. ● room: room level.

	<p>The default value is server.</p> <p>Note: If the fault domain level is room or rack, you must import the topology file for initialization.</p>
<code>--iscsi-port</code> <i>ISCSI_PORT</i>	<p>Specifies the iSCSI port.</p> <p>The value is an integer that ranges from 1 to 65535. The default value is 3260.</p>
<code>--port-range</code> <i>PORT1-PORT2</i>	<p>Specifies the port range. Storage services and services without specified ports will automatically be assigned ports from this range.</p> <p>The value is an integer that ranges from 1 to 65535. <i>PORT1</i> is the minimum value of the port range, <i>PORT2</i> is the maximum value of the port range, and <i>PORT1</i><<i>PORT2</i>. The default value of <i>PORT1</i> is 20000, and the default value of <i>PORT2</i> is 20500.</p> <p>Note: It is recommended that the specified port range contains at least 500 ports.</p>
<code>--data-port1</code> <i>DATA_PORT1</i>	<p>Specifies data port 1, only supported by cluster mode.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--management-port1</code> <i>MANAGEMENT_PORT1</i>	<p>Specifies management service port 1.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--management-port2</code> <i>MANAGEMENT_PORT2</i>	<p>Specifies management service port 2.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--management-port3</code> <i>MANAGEMENT_PORT3</i>	<p>Specifies management service port 3.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--management-port4</code> <i>MANAGEMENT_PORT4</i>	<p>Specifies management service port 4.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--management-port5</code> <i>MANAGEMENT_PORT5</i>	<p>Specifies management service port 5, only supported by cluster mode.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--management-port6</code> <i>MANAGEMENT_PORT6</i>	<p>Specifies management service port 6.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--metadata-port1</code> <i>METADATA_PORT1</i>	<p>Specifies metadata port 1, only supported by cluster mode.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--metadata-port2</code> <i>METADATA_PORT2</i>	<p>Specifies metadata port 2, only supported by cluster mode.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--metadata-port3</code> <i>METADATA_PORT3</i>	<p>Specifies metadata port 3, only supported by cluster mode.</p> <p>The value is an integer that ranges from 1 to 65535.</p>
<code>--metadata-port4</code> <i>METADATA_PORT4</i>	<p>Specifies Metadata port 4, only supported by cluster mode.</p> <p>The value is an integer that ranges from 1 to 65535.</p>

--metadata-port5 <i>METADATA_PORT5</i>	Specifies metadata port 5, only supported by cluster mode. The value is an integer that ranges from 1 to 65535.
--metadata-port6 <i>METADATA_PORT6</i>	Specifies metadata port 6, only supported by cluster mode. The value is an integer that ranges from 1 to 65535.
--metadata-port7 <i>METADATA_PORT7</i>	Specifies metadata port 7, only supported by cluster mode. The value is an integer that ranges from 1 to 65535.
--metadata-port8 <i>METADATA_PORT8</i>	Specifies metadata port 8, only supported by cluster mode. The value is an integer that ranges from 1 to 65535.
--cs <i>SERVER_IP[:DIR] , SERVER_IP[:DIR] , SERVER_IP[:DIR]</i>	Specifies the address of the coordination service and the disk path for storing coordination service data, only supported by cluster mode. If specifying the address of the coordination service, you must also specify all three IP addresses in the cluster. Note: In order to improve read-write performance, it is recommended that the disk path for storing coordination service data be independent of the installation disk path and the disk path where data is stored. <ul style="list-style-type: none"> ● <i>SERVER_IP</i>: The value is IPv4 or [IPv6] address. ● <i>DIR</i>: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; : . The default disk path is the installation disk path.
--mdm <i>SERVER_IP[:DIR] , SERVER_IP[:DIR]</i>	Specifies the address of the metadata management service and the disk path for storing metadata management data, only supported by cluster mode. If specifying the address of the metadata management service, you must also specify all two IP addresses in the cluster. Note: In order to improve read-write performance, it is recommended that the disk path for storing metadata management service data be independent of the installation disk path and the disk path where data is stored. <ul style="list-style-type: none"> ● <i>SERVER_IP</i>: The value is IPv4 or [IPv6] address. ● <i>DIR</i>: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; : . The default disk path is the installation disk path.
--ls <i>SERVER_IP[:DIR] , SERVER_IP[:DIR] , SERVER_IP[:DIR]</i>	Specifies the address of the log service and the disk path for storing log service data, only supported by cluster mode. If specifying the address of the log service, you must also specify all three IP addresses in the cluster.

	<p>Note: In order to improve read-write performance, it is recommended that the disk path for storing log service data be independent of the installation disk path and the disk path where data is stored.</p> <ul style="list-style-type: none"> ● <i>SERVER_IP</i>: The value is IPv4 or [IPv6] address. ● <i>DIR</i>: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; : . The default disk path is the installation disk path.
--	--

Examples:

- Initialize HBlock (standalone mode - commercial edition)

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor setup -n test -s
192.168.0.69:/mnt/storage01,/mnt/storage02
Please enter password:
*****
Start to setup HBlock, please wait.
Processing...
Setup successfully and the HBlock services have been started.
Welcome to HBlock!
You are using a 30-day trial version. Please follow the steps to get a license.
1. Run "stor info --serial-id" to get the serial ID of the HBlock
2. Contact the software vendor to obtain a license
3. Run "stor license add -k KEY" to import the license

Type 'stor --help' to get more information, such as managing LUNs, targets, servers, etc.
```

- Initialize HBlock (standalone mode - free edition)

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor setup -n stor1 -s
192.168.0.66:/mnt/stor01 --edition Free
Please enter password:
*****
Start to setup HBlock, please wait.
Processing...
Setup successfully and the HBlock services have been started.
Welcome to HBlock!
You are using a free edition with basic features to get started. The deadline for the
upgrade support is 2028-01-12 10:39:27. To ensure business continuity and access the full
product experience, please follow the steps to get a license.
1. Run "stor info --serial-id" to get the serial ID of the HBlock
2. Contact the software vendor to obtain a license
3. Run "stor license add -k KEY" to import the license

Type 'stor --help' to get more information, such as managing LUNs, targets, servers, etc.
```

- Initialize HBlock (cluster mode - commercial edition): The cluster topology uses the server mode.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor setup -n test -s
192.168.0.209:/mnt/storage01,/mnt/storage02 192.168.0.121:/mnt/storage01
192.168.0.72:/mnt/storage01
Please enter password:
*****
Start to setup HBlock, please wait.
Processing...
Setup successfully and the HBlock services have been started.
Welcome to HBlock!
You are using a 30-day trial version. Please follow the steps to get a license.
1. Run "stor info --serial-id" to get the serial ID of the HBlock
2. Contact the software vendor to obtain a license
3. Run "stor license add -k KEY" to import the license

Type 'stor --help' to get more information, such as managing LUNs, targets, servers, etc.
```

- Initialize HBlock (cluster mode - commercial edition): The cluster topology uses the topology file import mode.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor setup -n stor -p ***** --
topology-file /mnt/storage01/topology.txt --fault-domain path
Start to setup HBlock, please wait.
Processing...
Setup successfully and the HBlock services have been started.
Welcome to HBlock!
You are using a 30-day trial version. Please follow the steps to get a license.
1. Run "stor info --serial-id" to get the serial ID of the HBlock
2. Contact the software vendor to obtain a license
3. Run "stor license add -k KEY" to import the license

Type 'stor --help' to get more information, such as managing LUNs, targets, servers, etc.
```

The source code of the topology file is as follows:

```
{
  "name": "default",
  "childNodes": [
    {
      "name": "room1",
      "type": "room",
      "childNodes": [
        {
          "type": "server",
          "name": "server1",
          "ip": "192.168.0.192",
          "apiPort": 1443,
```

```

        "childNodes": [
            {
                "name": "/mnt/stor",
                "type": "path"
            },
            {
                "name": "/mnt/storage01",
                "type": "path"
            }
        ]
    },
    {
        "type": "server",
        "name": "server2",
        "ip": "192.168.0.110",
        "apiPort": 1443,
        "childNodes": [
            {
                "name": "/mnt/stor",
                "type": "path"
            }
        ]
    },
    {
        "type": "server",
        "name": "server3",
        "ip": "192.168.0.102",
        "apiPort": 1443,
        "childNodes": [
            {
                "name": "/mnt/stor",
                "type": "path"
            }
        ]
    }
]
}
]
}
}

```

- Initialize HBlock (cluster mode - free edition): The cluster topology uses the server mode.

```

[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor setup -n stor2 -s
192.168.0.64:/mnt/stor02,/mnt/data1 192.168.0.65:/mnt/stor01 192.168.0.67:/mnt/stor01 --
edition Free
Please enter password:
*****
Start to setup HBlock, please wait.
Processing...
Setup successfully and the HBlock services have been started.
Welcome to HBlock!

```

You are using a 30-day trial version. Please follow the steps to get a license.

1. Run "stor info --serial-id" to get the serial ID of the HBlock
2. Contact the software vendor to obtain a license
3. Run "stor license add -k KEY" to import the license

Type 'stor --help' to get more information, such as managing LUNs, targets, servers, etc.

4.5 Software License

4.5.1 Import Software License

```
./stor license add { -k | --key } KEY
```

This command is used to import software license.

Note:

- After the HBlock software is initialized, a 30-day trial period is provided. If you decide to use it, it is recommended to contact the software vendor as soon as possible to obtain the software license and import it.
- You need to execute command `./stor info --serial-id` to get serial number of HBlock and contact the software vendor to get the software license.
- Trial period expired, or software license expired: only partial features are available. For details, see **Features Available After Software License Expiration (Trial or Subscription Mode)**.
- For perpetual license, after the maintenance service expires, the upgrade function becomes unavailable, while all other features remain fully functional.

Parameters

Parameter	Description
<code>-k KEY</code> or <code>--key KEY</code>	Import software license.

Examples

Import Software License.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor license add -k *****
Import license successfully.
Customer name: Customer ABC
As the operating entity utilizing the HBlock to run your business, you bear the responsibility for the
network and information security of the business you are operating.

The current system time: 2025-08-05 09:38:45
+-----+-----+-----+-----+-----+
| LicenseId                | Account          | Type           | Status         | MaximumLocalCapacity |
+-----+-----+-----+-----+-----+
| qws2b6a9-f3fb-4098-a6b3-3652a5a76530 | test@ctyun.cn   | Perpetual      | Effective      | 2 PiB                 |
+-----+-----+-----+-----+-----+

License qws2b6a9-f3fb-4098-a6b3-3652a5a76530 (Effective):
Usage:
+-----+-----+-----+-----+-----+
| MaximumLocalCapacity | MaintenanceEffectiveTime | MaintenanceExpireTime | Status         |
+-----+-----+-----+-----+-----+
```

2 PiB	2025-08-05 09:34:27	2026-08-05 09:34:27	Effective
+-----+	+-----+	+-----+	+-----+

4.5.2 List Software Licenses

```
./stor license ls [ -r | --record ] [ -u | --usage ] [ { -n | --license } LICENSE_ID ]
```

This command is used to list all software licenses or the specified license.

Parameters

Parameter	Description
-r or --record	List purchase information for the software license.
-u or --usage	List usage information for the software license.
-n LICENSE_ID or --license LICENSE_ID	List specified software license information.

Examples

- List software licenses (Perpetual mode).

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor license ls
Customer name: Customer ABC
As the operating entity utilizing the HBlock to run your business, you bear the responsibility for the network and information security of the business you are operating.

License List:
+-----+-----+-----+-----+-----+
| LicenseId          | Account          | Type      | Status   | MaximumLocalCapacity |
+-----+-----+-----+-----+-----+
| qws2b6a9-f3fb-4098-a6b3-3652a5a76530 | test@ctyun.cn | Perpetual | Effective | 2 PiB                 |
+-----+-----+-----+-----+-----+

License qws2b6a9-f3fb-4098-a6b3-3652a5a76530 (Effective):
Record:
+-----+-----+-----+-----+-----+-----+
| PurchaseTime      | Operation        | LocalCapacity | MaintenanceEffectiveTime | MaintenanceExpireTime | Status   |
+-----+-----+-----+-----+-----+-----+
| 2025-08-05 09:34:27 | New              | 2 PiB         | 2025-08-05 09:34:27     | 2026-08-05 09:34:27 | Effective |
+-----+-----+-----+-----+-----+-----+

Usage:
+-----+-----+-----+-----+-----+
| MaximumLocalCapacity | MaintenanceEffectiveTime | MaintenanceExpireTime | Status   |
+-----+-----+-----+-----+-----+
| 2 PiB                 | 2025-08-05 09:34:27     | 2026-08-05 09:34:27 | Effective |
+-----+-----+-----+-----+-----+
```

- List software licenses (Subscription mode).

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor license ls
As the operating entity utilizing the HBlock to run your business, you bear the responsibility for the network and information security of the business you are operating.

License List:
+-----+-----+-----+-----+-----+
| LicenseId          | Account          | Type      | Status   | MaximumLocalCapacity |
+-----+-----+-----+-----+-----+
```

```

| ehc2b6a9-f3fb-4098-a6b3-3652a5d71232 | testaccount1@ctyun.com | Subscription | Effective | 1 PiB |
+-----+-----+-----+-----+-----+
License ehc2b6a9-f3fb-4098-a6b3-3652a5d71232 (Effective):
Record:
+-----+-----+-----+-----+-----+
| PurchaseTime | Operation | LocalCapacity | EffectiveTime | ExpireTime | Status |
+-----+-----+-----+-----+-----+
| 2025-07-30 14:40:36 | New | 1 PiB | 2025-07-30 14:40:36 | 2026-07-30 14:40:36 | Effective |
+-----+-----+-----+-----+-----+

Usage:
+-----+-----+-----+-----+-----+
| MaximumLocalCapacity | EffectiveTime | ExpireTime | Status |
+-----+-----+-----+-----+-----+
| 1 PiB | 2025-07-30 14:40:36 | 2026-07-30 14:40:36 | Effective |
+-----+-----+-----+-----+-----+
    
```

Table 1 Software license description

Item	Description
Customer name	Customer name. Note: This field will not be displayed if it was left blank when applying for the software license.
License List: Software license list	
Licenseld	Software license ID.
Type	Types of software license subscriptions: <ul style="list-style-type: none"> ● Subscription ● Perpetual
Status	Software license status: <ul style="list-style-type: none"> ● Effective: The software license is effective. ● Expired: The software license has expired. ● Invalid: The software license is invalid.
MaximumLocalCapacity	The total capacity allowed for all local LUNs within the specified period. "-" Indicates that there is no limit to the total capacity of local LUNs.
Record: software license purchase information	
PurchaseTime	Software license purchase time.
Operation	Software license purchase record: <ul style="list-style-type: none"> ● New. ● Expand. ● Renew.
LocalCapacity	All local LUNs' capacity limit of the license within the specified period.
EffectiveTime	The effective time of the subscription license in this purchase.

ExpireTime	The expiration time of the subscription license in this purchase.
MaintenanceEffectiveTime	The maintenance effective time of the perpetual license in this purchase.
MaintenanceExpireTime	The maintenance expiration time of the perpetual license in this purchase.
Status	Software license status corresponding to the purchase record: <ul style="list-style-type: none"> ● Effective: The software license is effective. ● Expired: The software license has expired. ● NotStart: The software license has not taken effect.
Usage: Software license usage.	
MaximumLocalCapacity	The total capacity allowed for all local LUNs within the specified period. "-" indicates that there is no limit to the total capacity of local LUNs.
EffectiveTime	The effective time of the subscription license.
ExpireTime	The expiration time of the subscription license.
MaintenanceEffectiveTime	The maintenance effective time of the perpetual license.
MaintenanceExpireTime	The maintenance expiration time of the perpetual license.
Status	Software license status of the usage: <ul style="list-style-type: none"> ● Effective: The software license is effective. ● Expired: The software license has expired. ● NotStart: The software license has not taken effect.

4.6 LUN Operations

4.6.1 Create a LUN

Local Mode LUN:

- Standalone Mode

```
./stor lun add { -n | --name } LUN_NAME { -p | --capacity } CAPACITY { -t | --target } TARGET_NAME [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy } WRITE_POLICY ] [ { -P | --path } PATH ] [ { -m | --mode } STORAGE_MODE ] [ [ --qos-name QOS_NAME ] [ --iops IOPS ] [ --read-iops READ_IOPS ] [ --write-iops WRITE_IOPS ] [ --bps BPS ] [ --read-bps READ_BPS ] [ --write-bps WRITE_BPS ] [ --iops-burst IOPS_BURST ] [ --read-iops-burst READ_IOPS_BURST ] [ --write-iops-burst WRITE_IOPS_BURST ] [ --bps-burst BPS_BURST ] [ --read-bps-burst READ_BPS_BURST ] [ --write-bps-burst WRITE_BPS_BURST ] [ --iops-burst-secs IOPS_BURST_SECS ] [ --read-iops-burst-secs READ_IOPS_BURST_SECS ] [ --write-iops-burst-secs WRITE_IOPS_BURST_SECS ] [ --bps-burst-secs BPS_BURST_SECS ] [ --read-bps-burst-secs READ_BPS_BURST_SECS ] [ --write-bps-burst-secs WRITE_BPS_BURST_SECS ] ] ]
```

- Cluster Mode

```
./stor lun add { -n | --name } LUN_NAME { -p | --capacity } CAPACITY { -t | --target } TARGET_NAME [ --priority SERVER_ID <1-n> [ --auto-failback AUTO_FAILBACK ] ] [ --pool POOL ] [ --cache-pool CACHE_POOL ] [ { -a | --ha } HIGH_AVAILABILITY ] [ { -c | --local-storage-class } LOCAL_STORAGE_CLASS ] [ --min-replica MIN_REPLICA ] [ --redundancy-overlap REDUNDANCY_OVERLAP ] [ --ec-fragment-size EC_FRAGEMENT_SIZE ] [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy } WRITE_POLICY ] [ { -m | --mode } STORAGE_MODE ] [ [ --qos-name QOS_NAME ] [ --iops IOPS ] [ --read-iops READ_IOPS ] [ --write-iops WRITE_IOPS ] [ --bps BPS ] [ --read-bps READ_BPS ] [ --write-bps WRITE_BPS ] [ --iops-burst IOPS_BURST ] [ --read-iops-burst READ_IOPS_BURST ] [ --write-iops-burst WRITE_IOPS_BURST ] [ --bps-burst BPS_BURST ] [ --read-bps-burst READ_BPS_BURST ] [ --write-bps-burst WRITE_BPS_BURST ] [ --iops-burst-secs IOPS_BURST_SECS ] [ --read-iops-burst-secs READ_IOPS_BURST_SECS ] [ --write-iops-burst-secs WRITE_IOPS_BURST_SECS ] [ --bps-burst-secs BPS_BURST_SECS ] [ --read-bps-burst-secs READ_BPS_BURST_SECS ] [ --write-bps-burst-secs WRITE_BPS_BURST_SECS ] ] ]
```

Cache/Storage Mode LUN:

● Standalone Mode

```
./stor lun add { -n | --name } LUN_NAME { -p | --capacity } CAPACITY { -t | --target }  
TARGET_NAME [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy }  
WRITE_POLICY ] [ { -P | --path } PATH ] { -m | --mode } STORAGE_MODE [ --cloud-provider  
CLOUD_PROVIDER ] { -B | --bucket } BUCKET_NAME { -A | --ak } ACCESS_KEY { -S | --sk }  
SECRET_KEY [ { -C | --cloud-storage-class } CLOUD_STORAGE_CLASS ] { -E | --endpoint }  
ENDPOINT [ --sign-version VERSION ] [ --region REGION ] [ { -M | --cloud-compression }  
CLOUD_COMPRESSION ] [ { -O | --object-size } OBJECT_SIZE ] [ { -X | --prefix } PREFIX ]  
[ [ --qos-name QOS_NAME ] [ --iops IOPS ] [ --read-iops READ_IOPS ] [ --write-iops  
WRITE_IOPS ] [ --bps BPS ] [ --read-bps READ_BPS ] [ --write-bps WRITE_BPS ] [ --iops-  
burst IOPS_BURST ] [ --read-iops-burst READ_IOPS_BURST ] [ --write-iops-burst  
WRITE_IOPS_BURST ] [ --bps-burst BPS_BURST ] [ --read-bps-burst READ_BPS_BURST ] [ --  
write-bps-burst WRITE_BPS_BURST ] [ --iops-burst-secs IOPS_BURST_SECS ] [ --read-iops-  
burst-secs READ_IOPS_BURST_SECS ] [ --write-iops-burst-secs WRITE_IOPS_BURST_SECS ] [ --  
bps-burst-secs BPS_BURST_SECS ] [ --read-bps-burst-secs READ_BPS_BURST_SECS ] [ --write-  
bps-burst-secs WRITE_BPS_BURST_SECS ] ] ]
```

● Cluster Mode

```
./stor lun add { -n | --name } LUN_NAME { -p | --capacity } CAPACITY { -t | --target }  
TARGET_NAME [ --priority SERVER_ID &<1-n> [ --auto-failback AUTO_FAILBACK ] ] [ --pool  
POOL ] [ --cache-pool CACHE_POOL ] [ { -a | --ha } HIGH_AVAILABILITY ] [ { -c | --local-  
storage-class } LOCAL_STORAGE_CLASS ] [ --min-replica MIN_REPLICA ] [ --redundancy-  
overlap REDUNDANCY_OVERLAP ] [ --ec-fragment-size EC_FRAGMENT_SIZE ] [ { -o | --sector-  
size } SECTOR_SIZE ] [ { -w | --write-policy } WRITE_POLICY ] { -m | --mode }  
STORAGE_MODE [ --cloud-provider CLOUD_PROVIDER ] { -B | --bucket } BUCKET_NAME { -A | --  
ak } ACCESS_KEY { -S | --sk } SECRET_KEY [ { -C | --cloud-storage-class }  
CLOUD_STORAGE_CLASS ] { -E | --endpoint } ENDPOINT [ --sign-version VERSION ] [ --region  
REGION ] [ { -M | --cloud-compression } CLOUD_COMPRESSION ] [ { -O | --object-size }  
OBJECT_SIZE ] [ { -X | --prefix } PREFIX ] [ [ --qos-name QOS_NAME ] [ --iops IOPS ] [ --  
read-iops READ_IOPS ] [ --write-iops WRITE_IOPS ] [ --bps BPS ] [ --read-bps READ_BPS ]  
[ --write-bps WRITE_BPS ] [ --iops-burst IOPS_BURST ] [ --read-iops-burst  
READ_IOPS_BURST ] [ --write-iops-burst WRITE_IOPS_BURST ] [ --bps-burst BPS_BURST ] [ --  
read-bps-burst READ_BPS_BURST ] [ --write-bps-burst WRITE_BPS_BURST ] [ --iops-burst-secs  
IOPS_BURST_SECS ] [ --read-iops-burst-secs READ_IOPS_BURST_SECS ] [ --write-iops-burst-
```

```
secs WRITE_IOPS_BURST_SECS ] [ --bps-burst-secs BPS_BURST_SECS ] [ --read-bps-burst-secs
READ_BPS_BURST_SECS ] [ --write-bps-burst-secs WRITE_BPS_BURST_SECS ] ]
```

This command is used to create a LUN.

Note:

- HBlock supports a maximum of 32766 target IQNs. A target can be associated with up to 256 LUNs, but each LUN can only be associated with one target.
- After a LUN is created, only the following parameters can be modified: LUN capacity, write policy, minimum number of replicas (only supported by cluster mode), redundancy overlap (only supported by cluster mode), priority (only supported by cluster mode), Endpoint, cloud signature authentication version, region, AK/SK, and whether to compress data.

Parameters

Parameter	Description
<code>-n LUN_NAME</code> or <code>--name LUN_NAME</code>	Specifies the LUN name. The value is a string of 1 to 16 case-sensitive characters. It can contain letters, digits, or hyphens (-). Only supports starting with a letter or a digit.
<code>-p CAPACITY</code> or <code>--capacity CAPACITY</code>	Specifies the storage capacity of the LUN. The value is an integer. The unit abbreviation G/g, T/t or P/p can be entered after the value, representing GiB, TiB, and PiB respectively. The default unit is GiB. <ul style="list-style-type: none"> ● If the unit is GiB, the value is an integer that ranges from 1 to 1048576. ● If the unit is TiB, the value is an integer that ranges from 1 to 1024. ● If the unit is PiB, the value is 1.
<code>-t TARGET_NAME</code> or <code>--target TARGET_NAME</code>	Specifies the iSCSI target name. The value is a string of 1 to 16 case-sensitive characters. It can contain lowercase letters, dots (.), digits, or hyphens (-). Only supports starting with a letter or a digit. Note: If the specified iSCSI target name does not exist when creating a LUN, the iSCSI target will be created at the same time, the reclaim policy of the iSCSI target is Delete.
<code>--priority SERVER_ID <1-n></code>	Specifies the sequential list of server ID used to identify the affinity priority of active and standby

	<p>iSCSI targets, only supported by cluster mode. The system will select the active and standby IQN of the LUN based on the sequence of the specified server IDs. Multiple server IDs can be added at one time, separated by comma (,).</p> <p>Prerequisite: The name of the iSCSI target already exists, and the specified server must be the same as the one where the iSCSI target is located.</p>
<p>--auto-failback <i>AUTO_FAILBACK</i></p>	<p>Indicates whether to automatically switch active and standby iSCSI targets of the LUN based on affinity priority, only supported by cluster mode. That is, when the high-priority server returns to normal, whether to automatically switch active and standby iSCSI targets of the LUN.</p> <p>Value:</p> <ul style="list-style-type: none"> ● Enabled: Automatic switch active and standby iSCSI targets. ● Disabled: No automatic switch active and standby iSCSI targets. <p>The default value is Enabled.</p>
<p>--pool <i>POOL</i></p>	<p>Specifies a storage pool, only supported by cluster mode. The storage pool is the final storage pool in which LUN data is stored. By default, the base storage pool in the cluster is used.</p>
<p>--cache-pool <i>CACHE_POOL</i></p>	<p>Specifies a cache storage pool, only supported by cluster mode. If a cache storage pool is specified, LUN data is first written to the cache storage pool and then stored to the storage pool.</p> <p>Note: The storage pool and cache storage pool must not be the same.</p>
<p>-a <i>HIGH_AVAILABILITY</i> or --ha <i>HIGH_AVAILABILITY</i></p>	<p>Sets the high availability type of the LUN, only supported by cluster mode.</p> <p>Value:</p> <ul style="list-style-type: none"> ● ActiveStandby (as): Enable active and standby. The LUN is associated with all IQNs under the corresponding target. ● Disabled (off): Disable active and standby LUNs. The LUN is associated with 1 target IQN under the corresponding target. For the client's method of connecting this type of LUN, see Windows Client – Standalone , Linux Client – Standalone . <p>The default value is ActiveStandby(as).</p>
<p>-c <i>LOCAL_STORAGE_CLASS</i> or</p>	<p>Specifies LUN data storage redundancy mode, only supported by cluster mode.</p>

<p><code>--local-storage-class</code> <code>LOCAL_STORAGE_CLASS</code></p>	<p>Value:</p> <ul style="list-style-type: none">● single-copy.● 2-copy.● 3-copy.● EC $N+M$: Erasure code mode. N and M are positive integers, $N \geq M$, and $N + M \leq 128$. This indicates that the data is divided into N fragments and M pieces of verification data are generated. <p>The default value is EC 2+1.</p> <p>Note: (All scenarios are based on the premise of cluster availability):</p> <ul style="list-style-type: none">● After an EC $N+M$ LUN is created:<ul style="list-style-type: none">■ Data can be written to the LUN if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to the minimum number of replicas of the LUN. Data cannot be written to the LUN and an alarm is generated if the number of available fault domains in the storage pool where the LUN resides is smaller than the minimum number of replicas.■ Data in the LUN is normal and will not degrade if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to $N+M$. Data in the LUN is being degraded if the number of available fault domains in the storage pool where the LUN resides is between $[N, N+M]$. We recommend that you add or repair the fault domains as soon as possible. Data written to the storage pool is corrupted if the number of available fault domains in the storage pool where the LUN resides is less than N.● After a LUN in replica mode is created:<ul style="list-style-type: none">■ Data can be written to the LUN if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to the minimum number of replicas of the LUN. Data cannot be written to the LUN and an alarm is generated if the number of available fault
--	--

	<p>domains in the storage pool where the LUN resides is smaller than the minimum number of replicas.</p> <ul style="list-style-type: none"> ■ Data in the LUN is normal and will not degrade if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to the number of replicas. For a two-replica or three-replica LUN, data in the LUN data is being degraded if the fault domain in the storage pool where the LUN resides is greater than or equal to 1, but less than the number of replicas. We recommend that you add or repair fault domains as soon as possible. Data written to the storage pool is corrupted if no fault domain is available in the storage pool where the LUN resides.
<p>--min-replica <i>MIN_REPLICA</i></p>	<p>Specifies the minimum replica number, only supported by cluster mode.</p> <ul style="list-style-type: none"> ● For a LUN in replica mode, assuming that the number of LUN replicas is X and the minimum replica number is Y (must satisfy $Y \leq X$), each time the data is written to LUN, at least Y replicas of data are written successfully before this write request is considered successful. ● For a LUN in EC $N + M$ mode, assuming that the minimum replica number of the LUN is set to Y (must satisfy $N \leq Y \leq N + M$), the data blocks and parity blocks that sum to at least Y blocks are written successfully before this write request is considered successful. <p>Value: For a LUN in replica mode, the value is an integer that ranges from 1 to N, where N is the number of replicas, the default value is 1. For a LUN in EC $N + M$ mode, the value is an integer that ranges from N to $N + M$, the default value is N.</p>
<p>--redundancy-overlap <i>REDUNDANCY_OVERLAP</i></p>	<p>Specifies the number of copies/fragments from the same data which are allowed to be distributed in the same fault domain, only supported by cluster mode. Different copies/fragments of the same data are distributed in different fault domains generally. When the fault domain is damaged, it is allowed to place multiple copies/fragments in the same fault domain</p>

	<p>but different paths according to the redundancy overlap principle.</p> <p>Note: If fault domain level of the storage pool is path, this parameter does not take effect.</p> <p>Value: For a LUN in replica mode, the value is an integer that ranges from 1 to N, where N is the number of replicas, the default value is 1. For a LUN in EC $N + M$ mode, the value is an integer that ranges from 1 to $N + M$, the default value is 1.</p>
<p>--ec-fragment-size <i>EC_FRAGEMENT_SIZE</i></p>	<p>The fragment size of erasure code. This parameter takes effect only when LUN data storage redundancy mode is in EC mode, otherwise it is ignored.</p> <p>Value: The value is 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048 or 4096, and the unit is KiB. The default value is 16.</p>
<p>-o <i>SECTOR_SIZE</i> or --sector-size <i>SECTOR_SIZE</i></p>	<p>Specifies the sector size.</p> <p>The value is 512 or 4096, the unit is bytes. The default value is 4096.</p> <p>Note: Selection of sector size: According to your business scenario, under normal circumstances, if the data size of a single I/O operation is greater than or close to 4 KiB, it is recommended to choose 4096; if the data size of a single I/O operation is close to 512 Bytes, it is recommended to choose 512. If connecting to virtualization platforms such as VMware, it is recommended to choose 512 Bytes.</p>
<p>-w <i>WRITE_POLICY</i> or --write-policy <i>WRITE_POLICY</i></p>	<p>Sets the write policy for the LUN:</p> <ul style="list-style-type: none"> ● WriteBack (wb): After the data is written to the memory, it returns to the client successfully, and then the data is written to the disk asynchronously. It is suitable for scenarios with high performance requirements and low stability requirements. ● WriteThrough (wt): The data is written to both memory and disk at the same time, and then returns to the client after successful writing. It is suitable for scenarios that the stability requirements are high, the write performance requirements are not high, and the recently written data will be read in a short time. ● WriteAround (wa): Write data directly to the disk without writing to memory. It is suitable for scenarios with high stability requirements, low

	<p>performance requirements, and more writing and less reading.</p> <p>The default value is WriteBack (wb).</p>
<p>-p <i>PATH</i> or --path <i>PATH</i></p>	<p>Specifies the disk path to store LUN data, only supported by standalone mode.</p> <p>Value: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; . :</p> <p>If not specified, the default disk path of the server will be used for the LUN.</p>
<p>-m <i>STORAGE_MODE</i> or --mode <i>STORAGE_MODE</i></p>	<p>Specifies the storage mode for the LUN:</p> <ul style="list-style-type: none"> ● Local: Store all data at local only. ● Cache: Store part of hot data at local and store all data in cloud asynchronously. ● Storage: Store all data at local and asynchronously store it in cloud. <p>The default value is Local.</p>
<p>--cloud-provider <i>CLOUD_PROVIDER</i></p>	<p>Specifies the type of object storage:</p> <ul style="list-style-type: none"> ● OOS: eSurfing Cloud Object-Oriented Storage. ● S3: Other object storage compatible with S3. <p>The default value is OOS.</p>
<p>-B <i>BUCKET_NAME</i> or --bucket <i>BUCKET_NAME</i></p>	<p>Specifies Bucket name of object storage.</p> <p>Note: Do not enable bucket lifecycle settings and compliance retention.</p>
<p>-A <i>ACCESS_KEY</i> or --ak <i>ACCESS_KEY</i></p>	<p>Specifies access key of object storage.</p>
<p>-S <i>SECRET_KEY</i> or --sk <i>SECRET_KEY</i></p>	<p>Specifies secret key of object storage.</p>
<p>-C <i>CLOUD_STORAGE_CLASS</i> or --cloud-storage-class <i>CLOUD_STORAGE_CLASS</i></p>	<p>Specifies storage class of object storage:</p> <ul style="list-style-type: none"> ● STANDARD. ● STANDARD_IA. <p>The default value is STANDARD.</p>
<p>-E <i>ENDPOINT</i> or --endpoint <i>ENDPOINT</i></p>	<p>Specifies endpoint of object storage.</p> <p>Note:</p> <ul style="list-style-type: none"> ● The Endpoint must match the object storage type you're using. If you use OOS, enter its Endpoint. If you use other object storage compatible with S3, enter its Endpoint. ● Access will be via HTTPS if only the domain name is entered. <p>For OOS endpoint details, see OOS Endpoint and Region.</p>
<p>--sign-version <i>VERSION</i></p>	<p>Specifies signature authentication version:</p>

	<ul style="list-style-type: none"> ● v2 ● v4 The default value is v2.
<code>--region REGION</code>	Specifies the region of endpoint. For the specific OOS region, see OOS Endpoint and Region . If the signature authentication version is v4, this item is required.
<code>-M CLOUD_COMPRESSION</code> or <code>--cloud-compression CLOUD_COMPRESSION</code>	Whether to compress data and upload it to object storage: <ul style="list-style-type: none"> ● Enabled (on): Compress data and upload it to object storage. ● Disabled (off): Do not compress data and upload it to object storage. The default value is Enabled (on).
<code>-O OBJECT_SIZE</code> or <code>--object-size OBJECT_SIZE</code>	Specifies the size of data stored in object storage. Values: 128, 256, 512, 1024, 2048, 4096, 8192, in KiB. The default value is 1024.
<code>-X PREFIX</code> or <code>--prefix PREFIX</code>	Specifies the prefix name for the LUN in object storage. After specifying the prefix name, the LUN data will be stored in the folder named with the prefix in the Bucket. If no prefix is specified, it is stored directly in the folder named after the LUN name. The value is a string of 1 to 256 characters.
<code>--qos-name QOS_NAME</code>	Specifies the name of a QoS policy. Note: If the specified QoS policy name does not exist when creating a LUN, the QoS policy will be created at the same time, the reclaim policy of the QoS policy is Delete. For cluster mode, when no QoS policy is explicitly specified: <ul style="list-style-type: none"> ● If the LUN has both a cache storage pool and a storage pool, its QoS policy is the default QoS policy for LUNs in the cache storage pool. If the cache storage pool has not defined such a default policy, the LUN has no QoS policy. ● If the LUN has only a storage pool, its QoS policy is the default QoS policy for LUNs in that storage pool. If the storage pool has not defined such a default policy, the LUN has no QoS policy. Value: The value is a string of 1 to 64 case-sensitive characters. It can contain letters, digits, and hyphens (-). Only supports starting with a letter or a digit. Note:

	<ul style="list-style-type: none"> ● If QoS policy name is not specified but any IOPS/Bps parameter is set, a new QoS policy is automatically created based on the specified IOPS/Bps. The system assigns the policy name in the format: <code>lun- lunname-qos-timestamp</code>. ● If QoS policy name is specified and any IOPS/Bps parameter is set, QoS policy name must be unique and cannot match any existing QoS policy. ● If QoS policy name is specified and no IOPS/Bps parameter is set, QoS policy name must be the name of an existing QoS policy.
<code>--iops IOPS</code>	Specifies the limit of I/O operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
<code>--read-iops READ_IOPS</code>	Specifies the limit of read operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
<code>--write-iops WRITE_IOPS</code>	Specifies the limit of write operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
<code>--bps BPS</code>	<p>Specifies the limit of I/O throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
<code>--read-bps READ_BPS</code>	Specifies the limit of read throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value,

	<p>representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
--write-bps <i>WRITE_BPS</i>	<p>Specifies the limit of write throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
--iops-burst <i>IOPS_BURST</i>	<p>Specifies the burst limit of I/O operations per second. This setting only takes effect when --iops <i>IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>IOPS</i>, 999,999,999]. The default value is -1, indicating no limit.</p>
--read-iops-burst <i>READ_IOPS_BURST</i>	<p>Specifies the burst limit of read operations per second. This setting only takes effect when --read-iops <i>READ_IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>READ_IOPS</i> 999,999,999]. The default value is -1, indicating no limit.</p>
--write-iops-burst <i>WRITE_IOPS_BURST</i>	<p>Specifies the burst limit of write operations per second.</p>

	<p>This setting only takes effect when --write-iops <i>WRITE_IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>WRITE_IOPS</i>, 999,999,999]. The default value is -1, indicating no limit.</p>
<p>--bps-burst <i>BPS_BURST</i></p>	<p>Specifies the burst limit of I/O throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --bps <i>BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3].
<p>--read-bps-burst <i>READ_BPS_BURST</i></p>	<p>Specifies the burst limit of read throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --read-bps <i>READ_BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,096,000,000,000].

	<ul style="list-style-type: none"> ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3].
<p>--write-bps-burst <i>WRITE_BPS_BURST</i></p>	<p>Specifies the burst limit of write throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --write-bps <i>WRITE_BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3].
<p>--iops-burst-secs <i>IOPS_BURST_SECS</i></p>	<p>Specifies the duration of I/O operations with burst limit.</p> <p>Note: This setting is only effective when the IOPS Burst feature is enabled.</p>

	Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--read-iops-burst-secs <i>READ_IOPS_BURST_SECS</i>	Specifies the duration of read operations with burst limit. Note: This setting is only effective when the read IOPS Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--write-iops-burst-secs <i>WRITE_IOPS_BURST_SECS</i>	Specifies the duration of write operations with burst limit. Note: This setting is only effective when the write IOPS Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--bps-burst-secs <i>BPS_BURST_SECS</i>	Specifies the duration of I/O throughput with burst limit. Note: This setting is only effective when the Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--read-bps-burst-secs <i>READ_BPS_BURST_SECS</i>	Specifies the duration of read throughput with burst limit. Note: This setting is only effective when the read Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--write-bps-burst-secs <i>WRITE_BPS_BURST_SECS</i>	Specifies the duration of write throughput with burst limit. Note: This setting is only effective when the write Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.

Examples:

- Standalone mode: Create local mode LUN lun1 with a capacity of 10 GiB, target as target1, and the LUN data is stored in the disk path /mnt/storage01.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lun1 -t target1 -p 10 -P /mnt/storage01
```

```
Created LUN lun1 successfully.
```

- Cluster mode: Create local mode LUN lun1 with a capacity of 10 GiB, target as target1, and LUN redundancy mode as 3-copy.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lun1 -t target1 -p 10 -c 3-copy  
Created LUN lun1 successfully.
```

- Cluster mode: Create cache mode LUN lun01a with a capacity of 100 GiB, target as target01, LUN redundancy mode as EC 2+1, and fragment size as 32 KiB.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lun01a -t target01 -c EC 2+1 --ec-fragment-size 32 -p 100  
Created LUN lun01a successfully.
```

- Standalone mode: Create storage mode LUN luna2 with a capacity of 110 GiB, target as targeta, use v4 signature authentication, and prefix as hblock.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n luna2 -t targeta -p 110 -m Storage -B hblocktest3 -E oos-cn.ctyunapi.cn -A cb22b08b1f9229f85874 -S ***** --sign-version v4 --region cn -X hblock  
Created LUN luna2 successfully.
```

- Cluster mode: Create cache mode LUN lun03a with a capacity of 100 GiB, target as target03, LUN redundancy mode as EC 2+1, fragment size as 32 KiB, and use v4 signature authentication.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun add -n lun03a -t target03 -c EC 2+1 --ec-fragment-size 32 -p 1000 -m Cache -B hblocktest1 -E oos-cn.ctyunapi.cn -A 9fe711b8b89ad73c2a46 -S ***** --sign-version v4 --region cn  
Created LUN lun03a successfully.
```

4.6.2 Create a Clone LUN

Standalone Mode:

```
./stor lun { C | clone } { -n | --name } LUN_NAME { -s | --snapshot } SNAPSHOT_NAME { -t | --target } TARGET_NAME [ { -p | --capacity } CAPACITY ] [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy } WRITE_POLICY ] [ { -P | --path } PATH ]
```

Cluster Mode:

```
./stor lun { C | clone } { -n | --name } LUN_NAME { -s | --snapshot } SNAPSHOT_NAME { -t | --target } TARGET_NAME [ { -p | --capacity } CAPACITY ] [ --priority SERVER_ID <1-n> [ --auto-failback AUTO_FAILBACK ] ] [ --pool POOL ] [ --cache-pool CACHE_POOL ] [ { -a | -ha } HIGH_AVAILABILITY ] [ { -c | --local-storage-class } LOCAL_STORAGE_CLASS ] [ --min-replica MIN_REPLICA ] [ --redundancy-overlap REDUNDANCY_OVERLAP ] [ --ec-fragment-size EC_FRAGMENT_SIZE ] [ { -o | --sector-size } SECTOR_SIZE ] [ { -w | --write-policy } WRITE_POLICY ]
```

This command is used to create a clone LUN.

Prerequisite: The snapshot used for creating a clone LUN must be in "Normal" status.

Notes:

- Maximum clone LUNs supported by the system: 100,000.
- Maximum clone LUNs creatable from a single snapshot: 512.
- Maximum cloning depth supported by the system: 16.

Parameters

Parameter	Description
-n LUN_NAME or --name LUN_NAME	Specifies the clone LUN name. The value is a string of 1 to 16 case-sensitive characters. It can contain letters, digits, or hyphens (-). Only supports starting with a letter or a digit.
-s SNAPSHOT_NAME or --snapshot SNAPSHOT_NAME	Specifies the name of the snapshot associated with the clone LUN. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit.
-t TARGET_NAME or --target TARGET_NAME	Specifies the name of the iSCSI target associated with the clone LUN. The value is a string of 1 to 16 case-sensitive characters. It can contain lowercase letters, dots (.),

	<p>digits, or hyphens (-). Only supports starting with a letter or a digit.</p> <p>Note: If the specified iSCSI target name does not exist when creating a clone LUN, the iSCSI target will be created at the same time, the reclaim policy of the iSCSI target is Delete.</p>
<p>-p <i>CAPACITY</i> or --capacity <i>CAPACITY</i></p>	<p>Specifies the storage capacity of the clone LUN. The value is an integer. The unit abbreviation G/g, T/t or P/p can be entered after the value, representing GiB, TiB, and PiB respectively. The default unit is GiB.</p> <ul style="list-style-type: none"> ● If the unit is GiB, the value is an integer that ranges from 1 to 1048576. ● If the unit is TiB, the value is an integer that ranges from 1 to 1024. ● If the unit is PiB, the value is 1. <p>The default is the source LUN's capacity at the snapshot time. If reconfigured, it must be no less than the source LUN's capacity at the snapshot time.</p>
<p>--priority <i>SERVER_ID &<1-n></i></p>	<p>Specifies the sequential list of server ID used to identify the affinity priority of active and standby iSCSI targets, only supported by cluster mode. The system will select the active and standby IQN of the clone LUN based on the sequence of the specified server IDs. Multiple server IDs can be added at one time, separated by comma (,).</p> <p>Prerequisite: The name of the iSCSI target already exists, and the specified server must be the same as the one where the iSCSI target is located.</p>
<p>--auto-failback <i>AUTO_FAILBACK</i></p>	<p>Indicates whether to automatically switch active and standby iSCSI targets of the clone LUN based on affinity priority, only supported by cluster mode. That is, when the high-priority server returns to normal, whether to automatically switch active and standby iSCSI targets of the clone LUN.</p> <p>Value:</p> <ul style="list-style-type: none"> ● Enabled: Automatic switch active and standby iSCSI targets. ● Disabled: No automatic switch active and standby iSCSI targets. <p>The default value is Enabled.</p>

<p><code>--pool POOL</code></p>	<p>Specifies the storage pool, only supported by cluster mode. The default value matches the source LUN's configuration.</p> <p>Note:</p> <ul style="list-style-type: none"> ● The <i>POOL</i> and <i>CACHE_POOL</i> must not be the same. ● If the clone LUN's <i>POOL</i> and <i>CACHE_POOL</i> settings match the source LUN's, no extra configuration is needed. If either the <i>POOL</i> or <i>CACHE_POOL</i> for the clone LUN is set separately, the corresponding settings of the source LUN are overridden, and the configured parameter values are used. For example, if the source LUN has both <i>CACHE_POOL</i> and <i>POOL</i> and the clone LUN only sets <i>POOL</i> without setting <i>CACHE_POOL</i>, then the clone LUN will only use the newly set <i>POOL</i> and have no <i>CACHE_POOL</i>. Conversely, if the clone LUN sets <i>CACHE_POOL</i>, you must also set <i>POOL</i> or use the base storage pool.
<p><code>--cache-pool CACHE_POOL</code></p>	<p>Specifies a cache storage pool, only supported by cluster mode. The default value matches the source LUN's configuration.</p> <p>Note:</p> <ul style="list-style-type: none"> ● The <i>POOL</i> and <i>CACHE_POOL</i> must not be the same. ● If the clone LUN's <i>POOL</i> and <i>CACHE_POOL</i> settings match the source LUN's, no extra configuration is needed. If either the <i>POOL</i> or <i>CACHE_POOL</i> for the clone LUN is set separately, the corresponding settings of the source LUN are overridden, and the configured parameter values are used. For example, if the source LUN has both <i>CACHE_POOL</i> and <i>POOL</i> and the clone LUN only sets <i>POOL</i> without setting <i>CACHE_POOL</i>, then the clone LUN will only use the newly set <i>POOL</i> and have no <i>CACHE_POOL</i>. Conversely, if the clone LUN

	sets <i>CACHE_POOL</i> , you must also set <i>POOL</i> or use the base storage pool.
<p>-a <i>HIGH_AVAILABILITY</i> or --ha <i>HIGH_AVAILABILITY</i></p>	<p>Sets the high availability type of the clone LUN, only supported by cluster mode.</p> <p>Value:</p> <ul style="list-style-type: none"> ● ActiveStandby (as): Enable active and standby. The LUN is associated with all IQNs under the corresponding target. ● Disabled (off): Disable active and standby LUNs. The LUN is associated with 1 target IQN under the corresponding target. For the client's method of connecting this type of LUN, see Windows Client – Standalone , Linux Client – Standalone . <p>The default value matches the source LUN's configuration.</p>
<p>-c <i>LOCAL_STORAGE_CLASS</i> or --local-storage-class <i>LOCAL_STORAGE_CLASS</i></p>	<p>Specifies the clone LUN data storage redundancy mode, only supported by cluster mode.</p> <p>Note: If --min-replica <i>MIN_REPLICA</i> or --redundancy-overlap <i>REDUNDANCY_OVERLAP</i> is set for the cloned LUN, this parameter must also be specified.</p> <p>Value:</p> <ul style="list-style-type: none"> ● single-copy. ● 2-copy. ● 3-copy. ● EC <i>N+M</i>: Erasure code mode. N and M are positive integers, $N \geq M$, and $N + M \leq 128$. This indicates that the data is divided into N fragments and M pieces of verification data are generated. <p>The default value matches the source LUN's configuration.</p> <p>Note: (All scenarios are based on the premise of cluster availability):</p> <ul style="list-style-type: none"> ● After an EC <i>N+M</i> LUN is created: <ul style="list-style-type: none"> ■ Data can be written to the LUN if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to the minimum number of replicas of the LUN. Data cannot be written to the LUN and an alarm is generated if the number of available fault

	<p>domains in the storage pool where the LUN resides is smaller than the minimum number of replicas.</p> <ul style="list-style-type: none"> ■ Data in the LUN is normal and will not degrade if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to $N+M$. Data in the LUN is being degraded if the number of available fault domains in the storage pool where the LUN resides is between $[N, N+M]$. We recommend that you add or repair the fault domains as soon as possible. Data written to the storage pool is corrupted if the number of available fault domains in the storage pool where the LUN resides is less than N. ● After a LUN in replica mode is created: <ul style="list-style-type: none"> ■ Data can be written to the LUN if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to the minimum number of replicas of the LUN. Data cannot be written to the LUN and an alarm is generated if the number of available fault domains in the storage pool where the LUN resides is smaller than the minimum number of replicas. ■ Data in the LUN is normal and will not degrade if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to the number of replicas. For a two-replica or three-replica LUN, data in the LUN data is being degraded if the fault domain in the storage pool where the LUN resides is greater than or equal to 1, but less than the number of replicas. We recommend that you add or repair fault domains as soon as possible. Data written to the storage pool is corrupted if no fault domain is available in the storage pool where the LUN resides.
<p>--min-replica <i>MIN_REPLICA</i></p>	<p>Specifies the minimum replica number, only supported by cluster mode.</p>

	<ul style="list-style-type: none"> ● For a LUN in replica mode, assuming that the number of LUN replicas is X and the minimum replica number is Y (must satisfy $Y \leq X$), each time the data is written to LUN, at least Y replicas of data are written successfully before this write request is considered successful. ● For a LUN in EC $N + M$ mode, assuming that the minimum replica number of the LUN is set to Y (must satisfy $N \leq Y \leq N + M$), the data blocks and parity blocks that sum to at least Y blocks are written successfully before this write request is considered successful. <p>Note: If this parameter is specified, you must also specify --local-storage-class <i>LOCAL_STORAGE_CLASS</i> for the cloned LUN.</p> <p>Value: For a LUN in replica mode, the value is an integer that ranges from 1 to N, where N is the number of replicas. For a LUN in EC $N + M$ mode, the value is an integer that ranges from N to $N + M$. The default value matches the source LUN's configuration.</p>
<p>--redundancy-overlap <i>REDUNDANCY_OVERLAP</i></p>	<p>Specifies the number of copies/fragments from the same data which are allowed to be distributed in the same fault domain, only supported by cluster mode. Different copies/fragments of the same data are distributed in different fault domains generally. When the fault domain is damaged, it is allowed to place multiple copies/fragments in the same fault domain but different paths according to the redundancy overlap principle.</p> <p>Note: If fault domain level of the storage pool is path, this parameter does not take effect. If this parameter is specified, you must also specify --local-storage-class <i>LOCAL_STORAGE_CLASS</i> for the cloned LUN.</p> <p>Value: For a LUN in replica mode, the value is an integer that ranges from 1 to N, where N is the number of replicas. For a LUN in EC $N + M$ mode, the value is an integer that ranges from 1 to $N + M$. The default value matches the source LUN's configuration.</p>
<p>--ec-fragment-size <i>EC_FRAGEMENT_SIZE</i></p>	<p>The fragment size of erasure code. This parameter takes effect only when LUN data storage</p>

	<p>redundancy mode is in EC mode, otherwise it is ignored.</p> <p>Value: The value is 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048 or 4096, and the unit is KiB. The default value matches the source LUN's configuration.</p>
<p>-o <i>SECTOR_SIZE</i> or --sector-size <i>SECTOR_SIZE</i></p>	<p>Specifies the sector size.</p> <p>The value is 512 or 4096, the unit is bytes. The default value matches the source LUN's configuration.</p> <p>Note: Selection of sector size: According to your business scenario, under normal circumstances, if the data size of a single I/O operation is greater than or close to 4 KiB, it is recommended to choose 4096; if the data size of a single I/O operation is close to 512 Bytes, it is recommended to choose 512. If connecting to virtualization platforms such as VMware, it is recommended to choose 512 Bytes.</p>
<p>-w <i>WRITE_POLICY</i> or --write-policy <i>WRITE_POLICY</i></p>	<p>Sets the write policy for the clone LUN:</p> <ul style="list-style-type: none"> ● WriteBack (wb): After the data is written to the memory, it returns to the client successfully, and then the data is written to the disk asynchronously. It is suitable for scenarios with high performance requirements and low stability requirements. ● WriteThrough (wt): The data is written to both memory and disk at the same time, and then returns to the client after successful writing. It is suitable for scenarios that the stability requirements are high, the write performance requirements are not high, and the recently written data will be read in a short time. ● WriteAround (wa): Write data directly to the disk without writing to memory. It is suitable for scenarios with high stability requirements, low performance requirements, and more writing and less reading. <p>The default value matches the source LUN's configuration.</p>
<p>-P <i>PATH</i> or --path <i>PATH</i></p>	<p>Specifies the disk path to store clone LUN data, only supported by standalone mode.</p> <p>Value: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; . :</p>

	The default disk path matches the source LUN's configuration.
--	---

Examples

- Standalone mode: Create clone LUN luna1-C1, associated with snapshot luna1-snapshot, iSCSI target targeta1, with a capacity of 200 GiB and other settings matching the source LUN.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun C -n luna1-C1 -s luna1-snapshot -t targeta1 -p 200
Confirm the information you set. If correct, enter Yes(Y). If not, enter No(N).
LUN Name: luna1-C1
Source Snapshot: luna1-snapshot
Capacity: 200 GiB
iSCSI Target: targeta1
Local Sector Size: 4096 Bytes
Write Policy: WriteBack
Path: /mnt/stor01
yes
Created clone LUN luna1-C1 successfully.
```

- Cluster mode: Create clone LUN lun01a-C2, associated with snapshot lun01a-s1 and iSCSI target target01, and other settings matching the source LUN.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun C -n lun01a-C2 -s lun01a-s1 -t target01
Confirm the information you set. If correct, enter Yes(Y). If not, enter No(N).
LUN Name: lun01a-C2
Source Snapshot: lun01a-s1
Capacity: 110 GiB
iSCSI target: target01
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
High Availability: ActiveStandby
Write Policy: WriteBack
yes
Created clone LUN lun01a-C2 successfully.
```

4.6.3 Flatten the Clone LUN

```
./stor lun { F | flatten } { -n | --name } LUN_NAME
```

This command is used to flatten the clone LUN.

Note:

- If the source LUN's data is damaged or unreadable due to other reasons, the system will keep retrying until it can read and copy the data to the clone LUN.
- This operation cannot be performed if the clone LUN is being rolled back.

Parameters

Parameter	Description
<code>-n LUN_NAME</code> or <code>--name LUN_NAME</code>	Specifies the name of the clone LUN to be flattened.

Examples

Flatten the clone LUN clone-lun02-2.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun F -n clone-lun02-2
Start flattening the clone LUN clone-lun02-2.
```

4.6.4 Delete a LUN

Local Mode LUN:

```
./stor lun rm { -n | --name } LUN_NAME [ -f | --force ]
```

Cache/Storage Mode LUN:

```
./stor lun rm { -n | --name } LUN_NAME [ { -c | --del-cloud } ] [ -f | --force ]
```

This command is used to delete the specified LUN.

For the cache or storage mode LUN, if the data on the cloud is not deleted when deleting the LUN, the LUN restore function can be used later for LUN data restoration. When deleting a LUN, ensure it has no snapshots or consistency snapshots.

Note:

- For the cache or storage mode LUN, if forcibly deleting the LUN while also deleting cloud data may result in residual cloud data, which needs to be manually deleted.
- If a LUN has associated clone LUNs that are not in "Deleting" or "Flattening" status, deleting the source LUN is prohibited.
- If a LUN has associated snapshots or consistency snapshots, you can only force-delete the LUN. Forcing deletion will also delete the snapshots and LUN snapshots in the consistency snapshots. This may cause data residue, please proceed with caution.

Parameters

Parameter	Description
<code>-n LUN_NAME</code> or <code>--name LUN_NAME</code>	Specifies the name of the deleted LUN.
<code>-c</code> or <code>--del-cloud</code>	Delete the data on the cloud when the cloud LUN is deleted. Note: If the LUN status is Suspended, Suspending, or SuspendFailed, the cloud data cannot be deleted even if the LUN is forcibly deleted.
<code>-f</code> or <code>--force</code>	Forcibly deleting the LUN. Note: Forcibly deleting the LUN may result in data residue. Please proceed with caution.

Examples

Delete LUN lun1.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun rm -n lun1
Start deleting LUN lun1. You can list LUNs to check whether the deletion is completed.
```

4.6.5 Expand a LUN

```
./stor lun { X | expand } { -n | --name } LUN_NAME { -p | --capacity } CAPACITY
```

The command is used to expand the storage capacity of the specified LUN.

Parameters

Parameter	Description
-n <i>LUN_NAME</i> or --name <i>LUN_NAME</i>	Specifies the name of the LUN to be expanded.
-p <i>CAPACITY</i> or --capacity <i>CAPACITY</i>	<p>Specifies the storage capacity after the LUN is expanded. The storage capacity of the LUN after expansion must be greater than or equal to the storage capacity of the LUN before expansion.</p> <p>The value is an integer. The unit abbreviation G/g, T/t or P/p can be entered after the value, representing GiB, TiB, and PiB respectively. The default unit is GiB.</p> <ul style="list-style-type: none"> ● If the unit is GiB, the value is an integer that ranges from 1 to 1048576. ● If the unit is TiB, the value is an integer that ranges from 1 to 1024. ● If the unit is PiB, the value is 1.

Examples

Note: If the LUN is not attached to the client, after the storage capacity of the LUN is expanded, the client will use the expanded storage capacity of the LUN when connecting. If the LUN is already attached to the client as local disk, after the storage capacity of the LUN is expanded, the client also needs to perform the LUN expansion operations.

- When LUN is not attached to the client, just perform LUN expansion directly on the server:

```
[root@server CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun X -n lun3 -p 11
Expanded LUN lun3 successfully.
```

- When the server is connected to the client:
 - The client is a Windows system:

Server:

```
[root@server CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun X -n lun3 -p 11
Expanded LUN lun3 successfully.
```

Client:

Enter **Server Manager > File and Storage services > Volumes > Disks**, right-click in blank, and then click **Rescan Storage**. After the disk capacity is updated, select the volume to be expanded in the **VOLUMES** area, right-click, and then click **Extend Volume** to expand.

➤ The client is a Linux system:

Server: Execute LUN expansion command

```
./stor lun { X | expand } { -n | --name } LUN_NAME { -p | --capacity } CAPACITY
```

Client:

1. Scan disk:

- If HBlock is standalone mode, you need to find the drive letter corresponding to the iSCSI LUN on the client and scan the disk.
- If HBlock is cluster mode, you need to find multiple drive letters corresponding to the iSCSI LUN on the client and scan the disk.

```
echo 1 > /sys/class/block/sdX/device/rescan
# Where sdX is the drive letter of the iSCSI LUN on the Client.
```

2. Resize the map (the client of HBlock cluster mode)

```
multipathd resize map mpathX
```

3. Resize the file system

```
resize2fs /dev/mapper/mpathX
# ext4 expansion. If it is XFS expansion, use xfs_growfs /dev/mapper/mpathX
```

Note: The device name of the standalone mode is `/dev/sdX`, and the device name of the cluster mode is `/dev/mapper/mpathX`. You can execute relevant commands according to the situation.

4. View the expanded disk

```
lsblk
```

Example: HBlock is cluster mode, the client is Linux system, the storage capacity of LUN lun01 is expanded from 30 GiB to 40 GiB.

■ Server

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun01
LUN Name: lun01 (LUN 0)
Storage Mode: Local
Capacity: 30 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.102:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target01.2(192.168.0.110:3260,Standby)
               iqn.2012-08.cn.ctyunapi.oos:target01.3(192.168.0.192:3260,ColdStandby)
```

```

Create Time: 2026-01-02 13:47:11
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 333fffffffffaae6e920
UUID: lun-uuid-1d5ec154-7fb8-412a-b14c-b490eb33dad3
Snapshot Count: 0
Snapshot Size: 0 B (Note: Snapshot size may vary due to LUN issues or parent snapshot
deletion.)
[root@hblockserver CTYUN_HBblock_Plus_4.0.0_x64]# ./stor lun X -n lun01 -p 40
Expanded LUN lun01 successfully.
[root@hblockserver CTYUN_HBblock_Plus_4.0.0_x64]# ./stor lun ls -n lun01
LUN Name: lun01 (LUN 0)
Storage Mode: Local
Capacity: 40 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.102:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target01.2(192.168.0.110:3260,Standby)
               iqn.2012-08.cn.ctyunapi.oos:target01.3(192.168.0.192:3260,ColdStandby)
Create Time: 2026-01-02 13:47:11
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 333fffffffffaae6e920
UUID: lun-uuid-1d5ec154-7fb8-412a-b14c-b490eb33dad3
Snapshot Count: 0
Snapshot Size: 0 B (Note: Snapshot size may vary due to LUN issues or parent snapshot
deletion.)
    
```

■ Client: Before the lun01 expansion, the client has connected to lun01

```

[root@client ~]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda          8:0    0   30G  0 disk
└─mpatha    252:3    0   30G  0 mpath /mnt/disk_mpatha
sdb          8:16    0  100G  0 disk
└─mpathb    252:2    0  100G  0 mpath
sdc          8:32    0   30G  0 disk
└─mpatha    252:3    0   30G  0 mpath /mnt/disk_mpatha
    
```

```
sdd            8:48    0 100G  0 disk
└─mpathb      252:2    0 100G  0 mpath
sde            8:64    0  30G  0 disk
└─mpatha      252:3    0  30G  0 mpath /mnt/disk_mpatha
sdf            8:80    0 100G  0 disk
└─mpathb      252:2    0 100G  0 mpath
vda            253:0    0  40G  0 disk
├─vda1        253:1    0   4G  0 part [SWAP]
└─vda2        253:2    0  36G  0 part /
[root@client ~]# echo 1 > /sys/class/block/sda/device/rescan
[root@client ~]# echo 1 > /sys/class/block/sdc/device/rescan
[root@client ~]# echo 1 > /sys/class/block/sde/device/rescan
[root@client ~]# multipathd resize map mpatha
ok
[root@client ~]# resize2fs /dev/mapper/mpatha
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/mapper/mpatha is mounted on /mnt/disk_mpatha; on-line resizing required
old_desc_blocks = 4, new_desc_blocks = 5
The filesystem on /dev/mapper/mpatha is now 10485760 blocks long.

[root@client ~]# lsblk
NAME            MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda              8:0     0   40G  0 disk
└─mpatha         252:3    0   40G  0 mpath /mnt/disk_mpatha
sdb              8:16    0  100G  0 disk
└─mpathb        252:2    0  100G  0 mpath
sdc              8:32    0   40G  0 disk
└─mpatha         252:3    0   40G  0 mpath /mnt/disk_mpatha
sdd              8:48    0  100G  0 disk
└─mpathb        252:2    0  100G  0 mpath
sde              8:64    0   40G  0 disk
└─mpatha         252:3    0   40G  0 mpath /mnt/disk_mpatha
sdf              8:80    0  100G  0 disk
└─mpathb        252:2    0  100G  0 mpath
vda             253:0    0   40G  0 disk
├─vda1          253:1    0    4G  0 part [SWAP]
└─vda2          253:2    0  36G  0 part /
```

4.6.6 Edit a LUN

Local Mode LUN:

```
./stor lun set { -n | --name } LUN_NAME [ { -w | --write-policy } WRITE_POLICY ] [ --min-replica MIN_REPLICA ] [ --redundancy-overlap REDUNDANCY_OVERLAP ] [ --priority { SERVER_ID &<1-n> } | RANDOM ] [ --auto-failback AUTO_FAILBACK ]
```

Cache/Storage Mode LUN:

```
./stor lun set { -n | --name } LUN_NAME [ { -w | --write-policy } WRITE_POLICY ] [ --min-replica MIN_REPLICA ] [ --redundancy-overlap REDUNDANCY_OVERLAP ] [ --priority { SERVER_ID &<1-n> } | RANDOM ] [ --auto-failback AUTO_FAILBACK ] [ { -A | --ak } ACCESS_KEY { -S | --sk } SECRET_KEY ] [ { -E | --endpoint } ENDPOINT ] [ --sign-version VERSION ] [ --region REGION ] [ { -M | --cloud-compression } CLOUD_COMPRESSION ]
```

This command is used to edit the specified LUN.

Note: LUN configurations cannot be modified if the LUN is in "Deleting" or "Rollbacking" status.

Parameters

Parameter	Description
<code>-n LUN_NAME</code> or <code>--name LUN_NAME</code>	Specifies LUN name to be edited.
<code>-w WRITE_POLICY</code> or <code>--write-policy WRITE_POLICY</code>	Sets the write policy of the LUN: <ul style="list-style-type: none"> ● WriteBack (wb): After the data is written to the memory, it returns to the client successfully, and then the data is written to the disk asynchronously. It is suitable for scenarios with high performance requirements and low stability requirements. ● WriteThrough (wt): The data is written to both memory and disk at the same time, and then returns to the client after successful writing. It is suitable for scenarios that the stability requirements are high, the write performance requirements are not high, and the recently written data will be read in a short time. ● WriteAround (wa): Write data directly to the disk without writing to memory. It is suitable for scenarios with high stability requirements, low performance requirements, and more writing and less reading.
<code>--min-replica MIN_REPLICA</code>	Specifies the minimum replica number, only supported by cluster mode. For a LUN in replica mode, assuming that the number of LUN replicas is X and the minimum replica number is Y (must satisfy

	<p>Y <= X), each time the data is written to LUN, at least Y replicas of data are written successfully before this write request is considered successful. For a LUN in EC N + M mode, assuming that the minimum replica number of the LUN is set to Y (must satisfy N <= Y <= N + M), the data blocks and parity blocks that sum to at least Y blocks are written successfully before this write request is considered successful.</p> <p>Value: For a LUN in replica mode, the value is an integer that ranges from 1 to N, where N is the number of replica. For a LUN in EC N + M mode, the value is an integer that ranges from N to N + M.</p>
<p>--redundancy-overlap <i>REDUNDANCY_OVERLAP</i></p>	<p>Specifies the number of copies/fragments from the same data which are allowed to be distributed in the same fault domain, only supported by cluster mode. Different copies/fragments of the same data are distributed in different fault domains generally. When the fault domain is damaged, it is allowed to place multiple copies/fragments in the same fault domain but different paths according to the redundancy overlap principle.</p> <p>Note: If fault domain level of the storage pool is path, this parameter does not take effect.</p> <p>Value: For a LUN in replica mode, the value is an integer that ranges from 1 to N, where N is the number of replicas. For a LUN in EC N + M mode, the value is an integer that ranges from 1 to N + M.</p>
<p>--priority <i>SERVER_ID</i> &<1-n></p>	<p>Specifies the server IDs for LUN active-standby distribution priority. The system selects the active and standby IQNs for the LUN based on the specified server ID order. One or multiple server IDs can be specified; when multiple servers are specified, separate them with by comma (,).</p> <p>Example: lun1 corresponds to target1, and its IQN is distributed across hblock_1, hblock_2, and hblock_4. To specify the server IDs for lun1 LUN primary-backup distribution priority, such as --priority hblock_1,hblock_3,hblock_4, indicates that the target1 IQN corresponding to hblock_1 is prioritized as the active IQN for lun1, and the target1 IQN corresponding to hblock_3 is used as the standby IQN. If hblock_1 fails, the system automatically switches to the target1 IQN corresponding to hblock_3 as the active IQN, and the target1 IQN corresponding to hblock_4 as the standby IQN. When hblock_1 recovers from the failure:</p> <ul style="list-style-type: none"> ● If --auto-failback is set to Enabled, the system automatically restores to the configured priority order, with Active target IQN, Standby target IQN, and Cold target IQN switching to hblock_1, hblock_3, and hblock_4 respectively.

	<ul style="list-style-type: none"> ● If <code>--auto-failback</code> is set to Disabled, the system maintains the current state unchanged; to restore to the configured priority order, execute the command <code>./stor lun S -n LUN_NAME</code> to switch manually. <p>Note: If you specify the server IDs for LUN primary-backup distribution priority, you must and can only specify one LUN name.</p>
<code>--priority RANDOM</code>	The system randomly selects the active and standby IQNs for the LUN. That is, when the server hosting the target IQN corresponding to the LUN fails, the system randomly switches the active and standby IQNs for the LUN.
<code>--auto-failback AUTO_FAILBACK</code>	Whether to automatically switch active and standby iSCSI targets of the LUN based on affinity priority. That is, when the high-priority server returns to normal, whether to automatically switch active and standby iSCSI targets of the LUN. Value: <ul style="list-style-type: none"> ● Enabled: Automatic switch active and standby iSCSI targets. ● Disabled: No automatic switch active and standby iSCSI targets.
<code>-A ACCESS_KEY</code> or <code>--ak ACCESS_KEY</code>	Access key of object storage. Note: Access key and Secret key must be modified simultaneously.
<code>-S SECRET_KEY</code> or <code>--sk SECRET_KEY</code>	Secret key of object storage. Note: Access key and Secret key must be modified simultaneously.
<code>-E ENDPOINT</code> or <code>--endpoint ENDPOINT</code>	Specifies endpoint of object storage. Note: <ul style="list-style-type: none"> ● The Endpoint must match the object storage type you're using. If you use OOS, enter its Endpoint. If you use other object storage compatible with S3, enter its Endpoint. ● Access will be via HTTPS if only the domain name is entered. For OOS endpoint details, see OOS Endpoint and Region .
<code>--sign-version VERSION</code>	Specifies signature authentication version: <ul style="list-style-type: none"> ● v2 ● v4
<code>--region REGION</code>	Specifies the region of endpoint. If the signature authentication version is v4, this item is required.
<code>-M CLOUD_COMPRESSION</code> or	Whether to compress data and upload it to object storage: <ul style="list-style-type: none"> ● Enabled (on): Compress data and upload it to object storage.

--cloud-compression
CLOUD_COMPRESSION

- Disabled (off): Do not compress data and upload it to object storage.

Examples

- Modify the write policy of LUN lun01a to WriteThrough (wt).

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun set -n lun01a -w wt
Set LUN lun01a successfully.
```

- Modify the minimum replica number of LUN lun01a to 3.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun set -n lun01a --min-replica 3
Set LUN lun01a successfully.
```

- Modify sequential list of server ID for identifying the affinity priority of active and standby iSCSI targets.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun set -n lun01a --priority
hblock_1
Set LUN lun01a successfully.
```

- Modify the signature authentication version of the LUN to v4 and upload the data to object storage without compressing it.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun set -n lun03a --sign-version v4
--region cn -M off
Set LUN lun03a successfully.
```

4.6.7 Set LUN Primary-Secondary Priority or Automatic Failover (Cluster Mode)

```
./stor lun prefer [ { -n | --name } LUN_NAME &<1-n> ] [--priority { SERVER_ID &<1-n> |
RANDOM | CUSTOM } ] [ --auto-failback AUTO_FAILBACK ]
```

This command is used to set the primary-secondary priority or automatic failover of the LUN.

Parameters

Parameter	Description
<code>-n LUN_NAME &<1-n></code> or <code>--name LUN_NAME &<1-n></code>	Specifies the LUN name. You can specify one or more LUN, separated by comma (.). Note: If no LUN name is specified, it indicates that the setting applies to all LUNs.
<code>--priority SERVER_ID &<1-n></code>	Specifies the server IDs for LUN active-standby distribution priority. The system selects the active and standby IQNs for the LUN based on the specified server ID order. One or multiple server IDs can be specified; when multiple servers are specified, separate them with by comma (.). Example: lun1 corresponds to target1, and its IQN is distributed across hblock_1, hblock_2, and hblock_4. To specify the server IDs for lun1 LUN primary-backup distribution priority, such as --priority hblock_1,hblock_3,hblock_4 , indicates that the target1 IQN corresponding to hblock_1 is prioritized as the active IQN for lun1, and the target1 IQN corresponding to hblock_3 is used as the standby IQN. If hblock_1 fails, the system automatically switches to the target1 IQN corresponding to hblock_3 as the active IQN, and the target1 IQN corresponding to hblock_4 as the standby IQN. When hblock_1 recovers from the failure: <ul style="list-style-type: none"> ● If <code>--auto-failback</code> is set to <code>Enabled</code>, the system automatically restores to the configured priority order, with Active target IQN, Standby target IQN, and Cold target IQN switching to hblock_1, hblock_3, and hblock_4 respectively. ● If <code>--auto-failback</code> is set to <code>Disabled</code>, the system maintains the current state unchanged; to restore to the configured priority order, execute the command <code>./stor lun S -n LUN_NAME</code> to switch manually.

	<p>Note: If you specify the server IDs for LUN primary-backup distribution priority, you must and can only specify one LUN name.</p>
--priority RANDOM	<p>The system randomly selects the active and standby IQNs for the LUN. That is, when the server hosting the target IQN corresponding to the LUN fails, the system randomly switches the active and standby IQNs for the LUN.</p>
--priority CUSTOM	<p>Configures the priority based on the servers where the LUN's current active and standby target IQNs (Active target IQN, Standby target IQN, Cold target IQN) are located. After configuration, whether automatic adjustment is required depends on the <code>auto-failback</code> setting.</p> <p>Example: lun1 corresponds to target1, with its Active target, Standby target, and Cold target located on hblock_1, hblock_2, and hblock_4 respectively. If hblock_1 fails, the system automatically switches to the target1 IQN corresponding to hblock_2 as the active IQN, and the target1 IQN corresponding to hblock_4 as the standby IQN. When hblock_1 recovers from the failure:</p> <ul style="list-style-type: none"> ● If <code>--auto-failback</code> is set to Enabled, the system automatically restores to the configured priority order, with Active target IQN, Standby target IQN, and Cold target IQN switching to hblock_1, hblock_2, and hblock_4 respectively. ● If <code>--auto-failback</code> is set to Disabled, the system maintains the current state unchanged; to restore to the configured priority order, execute the command <code>./stor lun S -n LUN_NAME</code> to switch manually.
--auto-failback <i>AUTO_FAILBACK</i>	<p>Indicate whether to automatically switch active and standby iSCSI targets of the LUN based on affinity priority. That is, when the high-priority server returns to normal, whether to automatically switch active and standby iSCSI targets of the LUN.</p> <p>Value:</p> <ul style="list-style-type: none"> ● Enabled: Automatic switch active and standby iSCSI targets. ● Disabled: No automatic switch active and standby iSCSI targets.

Example

- Sets the active-standby priority for LUN lun04a to hblock_3,hblock_1,hblock_2. If one of these servers fails later, lun04a will select the active and standby IQNs according to this priority order.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun prefer -n lun04a --priority hblock_3,hblock_1,hblock_2
Set LUN lun04a with server affinity successfully.
```

- Sets automatic active-standby failover for LUNs lun04a and lun04b to Disabled.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun prefer -n lun04a,lun04b --auto-failback Disabled
Set LUN lun04a,lun04b with server affinity successfully.
```

- Configures the active-standby priority for LUNs lun04a and lun04b based on the servers where their current active and standby targets are located.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun prefer -n lun04a,lun04b --priority CUSTOM
Set LUN lun04a,lun04b with server affinity successfully.
```

4.6.8 Trigger Active/Standby Switchover of the Target Corresponding to the LUN (Cluster Mode)

```
./stor lun { S | switch } [ { -n | --name } LUN_NAME &<1-n> ]
```

This command is used to trigger the active/standby switchover of the target corresponding to the LUN.

Note: After performing this operation, no action is required on the client side.

Parameters

Parameter	Description
<code>-n LUN_NAME &<1-n></code> or <code>--name LUN_NAME &<1-n></code>	Specifies the LUN name. You can specify one or more LUN, separated by comma (.). Note: If no LUN name is specified, it means to set all LUNs.

Examples

Trigger the active/standby switchover of the target corresponding to LUN lun01a.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun S -n lun01a
Triggered the active and standby switchover on LUN lun01a successfully.
```

Trigger the active/standby switchover of the target corresponding to all LUNs.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor lun S
Triggered the active and standby switchover on LUN lun01a,lun02a,lun03a successfully.
```

4.6.9 Recovering a LUN (Cache/Storage Mode LUN)

● Standalone Mode

```
./stor lun { RC | --recover } { -n | --name } LUN_NAME { -t | --target } TARGET_NAME --
uuid UUID [ { -w | --write-policy } WRITE_POLICY ] [ { -P | --path } PATH ] { -m | --
mode } STORAGE_MODE [ --cloud-provider CLOUD_PROVIDER ] { -B | --bucket } BUCKET_NAME
[ { -X | --prefix } PREFIX ] { -A | --ak } ACCESS_KEY { -S | --sk } SECRET_KEY { -E | --
endpoint } ENDPOINT [ --sign-version VERSION ] [ --region REGION ] [ { -C | --cloud-
storage-class } CLOUD_STORAGE_CLASS ] [ { -M | --cloud-compression } CLOUD_COMPRESSION ]
```

Cluster Mode:

```
./stor lun { RC | --recover } { -n | --name } LUN_NAME { -t | --target } TARGET_NAME [ -
-priority SERVER_ID &<1-n> [ --auto-failback AUTO_FAILBACK ] ] [ --pool POOL ] [ --
cache-pool CACHE_POOL ] --uuid UUID [ { -a | --ha } HIGH_AVAILABILITY ] [ { -c | --local-
storage-class } LOCAL_STORAGE_CLASS ] [ --min-replica MIN_REPLICA ] [ --ec-fragment-size
EC_FRAGMENT_SIZE ] [ --redundancy-overlap REDUNDANCY_OVERLAP ] [ { -w | --write-policy }
WRITE_POLICY ] { -m | --mode } STORAGE_MODE [ --cloud-provider CLOUD_PROVIDER ] { -B | --
bucket } BUCKET_NAME [ { -X | --prefix } PREFIX ] { -A | --ak } ACCESS_KEY { -S | --sk }
SECRET_KEY [ { -C | --cloud-storage-class } CLOUD_STORAGE_CLASS ] { -E | --endpoint }
ENDPOINT [ --sign-version VERSION ] [ --region REGION ] [ { -M | --cloud-compression }
CLOUD_COMPRESSION ]
```

This command is used to recover a cache or a storage mode LUN.

The following scenarios are suitable for using the recover LUN function:

- There is a storage mode LUN or a cache mode LUN in the original HBlock, and the LUN data has been uploaded to the cloud. If the HBlock fails and cannot be started at this time, the LUN can be regenerated in another HBlock through the function of recovering the LUN, and LUN data can be recovered from the cloud.
- The original cache mode LUN or storage mode LUN has been deleted, but the cloud data is retained. The LUN data can be recovered through the recover function.

Note:

- When performing recover operation, please ensure that the connection between the original HBlock and the cloud has been disconnected, and no new data will be written to the LUN.
- The LUN name to be recovered does not exist in the current HBlock.

- The LUN to be recovered must be found in the specified bucket/prefix and the data must be complete.
- The recover LUN operation is executed asynchronously. Please check the recover progress by querying the LUN function.
- Before recovering a LUN in standalone mode, ensure that no residual data of that LUN exists in its corresponding data directory.
- When a LUN is in Recovering or RecoverFailed status, read and write operations are not supported.

Parameters

Parameter	Description
-n <i>LUN_NAME</i> or --name <i>LUN_NAME</i>	The source LUN name.
-t <i>TARGET_NAME</i> or --target <i>TARGET_NAME</i>	Specifies the iSCSI target name for the LUN. The value is a string of 1 to 16 case-sensitive characters. It can contain lowercase letters, dots (.), digits, or hyphens (-). Only supports starting with a letter or a digit. Note: If the specified iSCSI target name does not exist when recovering the LUN, the iSCSI target will be created at the same time.
--priority <i>SERVER_ID</i> &<1-n>	Specifies the sequential list of server ID used to identify the affinity priority of active and standby iSCSI targets, only supported by cluster mode. The system will select the active and standby IQN of the LUN based on the sequence of the specified server IDs. Multiple server IDs can be added at one time, separated by comma (,). Prerequisite: The name of the iSCSI target already exists, and the specified server must be the same as the one where the iSCSI target is located.
--auto-failback <i>AUTO_FAILBACK</i>	Indicates whether to automatically switch active and standby iSCSI targets of the LUN based on affinity priority, only supported by cluster mode. That is, when the high-priority server returns to normal, whether to automatically switch active and standby iSCSI targets of the LUN. Value: <ul style="list-style-type: none"> ● Enabled: Automatic switch active and standby iSCSI targets.

	<ul style="list-style-type: none"> ● Disabled: No automatic switch active and standby iSCSI targets. The default value is Enabled .
--pool <i>POOL</i>	Specifies the storage pool (only supported by cluster mode). The storage pool is the final storage pool in which LUN data is stored. By default, the base storage pool in the cluster is used.
--cache-pool <i>CACHE_POOL</i>	Specifies the cache storage pool (only supported by cluster mode). If a cache storage pool is specified, LUN data is first written to the cache storage pool and then stored to the storage pool. Note : The storage pool and cache storage pool must not be the same.
--uuid <i>UUID</i>	Universally unique identifier of the source LUN.
-a <i>HIGH_AVAILABILITY</i> or --ha <i>HIGH_AVAILABILITY</i>	Sets the high availability type of the LUN, only supported by cluster mode. Value: <ul style="list-style-type: none"> ● ActiveStandby (as): Enable active and standby. The LUN is associated with all IQNs under the corresponding target. ● Disabled (off): Disable active and standby LUNs. The LUN is associated with 1 target IQN under the corresponding target. For the client's method of connecting this type of LUN, see Windows Client – Standalone , Linux Client – Standalone . Consistent with source configuration by default.
-c <i>LOCAL_STORAGE_CLASS</i> or --local-storage-class <i>LOCAL_STORAGE_CLASS</i>	Sets the redundancy mode for the LUN (only supported by cluster mode). Value: <ul style="list-style-type: none"> ● single-copy. ● 2-copy. ● 3-copy. ● EC <i>N+M</i>: Erasure code mode. Wherein, N and M are positive integers, $N \geq M$ and $N+M \leq 128$. Indicates that the data is divided into N fragments and M pieces of parity data is generated. The default value is the redundancy mode of the source LUN. Note (In the following scenarios, the cluster is available): <ul style="list-style-type: none"> ● After the EC <i>N+M</i> LUN is restored: <ul style="list-style-type: none"> ■ Data can be written to the LUN if the number of available fault domains in the storage pool

	<p>where the LUN resides is greater than or equal to the minimum number of replicas of the LUN. Data cannot be written to the LUN and an alarm is generated if the number of available fault domains in the storage pool where the LUN resides is smaller than the minimum number of replicas.</p> <ul style="list-style-type: none">■ Data in the LUN is normal and will not degrade if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to $N+M$. Data in the LUN is being degraded if the number of available fault domains in the storage pool where the LUN resides is between $[N, N+M]$. We recommend that you add or repair the fault domains as soon as possible. Data written to the storage pool is corrupted if the number of available fault domains in the storage pool where the LUN resides is less than N.● After a LUN is restored in replica mode:<ul style="list-style-type: none">■ Data can be written to the LUN if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to the minimum number of replicas of the LUN. Data cannot be written to the LUN and an alarm is generated if the number of available fault domains in the storage pool where the LUN resides is smaller than the minimum number of replicas.■ Data in the LUN is normal and will not degrade if the number of available fault domains in the storage pool where the LUN resides is greater than or equal to the number of replicas. For a two-replica or three-replica LUN, data in the LUN data is being degraded if the fault domain in the storage pool where the LUN resides is greater than or equal to 1, but less than the number of replicas. We recommend that you add or repair fault domains as soon as possible. Data written to the storage pool is corrupted if no fault domain is available in the storage pool where the LUN resides.
--	--

<p>--min-replica <i>MIN_REPLICA</i></p>	<p>Specifies the minimum replica number, only supported by cluster mode.</p> <ul style="list-style-type: none"> ● For a LUN in replica mode, assuming that the number of LUN replicas is X and the minimum replica number is Y (must satisfy $Y \leq X$), each time the data is written to LUN, at least Y replicas of data are written successfully before this write request is considered successful. ● For a LUN in EC $N + M$ mode, assuming that the minimum replica number of the LUN is set to Y (must satisfy $N \leq Y \leq N + M$), the data blocks and parity blocks that sum to at least Y blocks are written successfully before this write request is considered successful. <p>Note: If the minimum number for the LUN is specified, the redundancy mode for the LUN must be specified. For a LUN in replica mode, the value is an integer that ranges from 1 to N, where N is the number of replicas. For a LUN in EC $N + M$ mode, the value is an integer that ranges from N to $N + M$. Consistent with source configuration by default.</p>
<p>--redundancy-overlap <i>REDUNDANCY_OVERLAP</i></p>	<p>Specifies the number of copies/fragments from the same data which are allowed to be distributed in the same fault domain, only supported by cluster mode. Different copies/fragments of the same data are distributed in different fault domains generally. When the fault domain is damaged, it is allowed to place multiple copies/fragments in the same fault domain but different paths according to the redundancy overlap principle.</p> <p>Note: If fault domain level of the storage pool is path, this parameter does not take effect.</p> <p>Value: For a LUN in replica mode, the value is an integer that ranges from 1 to N, where N is the number of replicas. For a LUN in EC $N + M$ mode, the value is an integer that ranges from 1 to $N + M$. Consistent with source configuration by default.</p>
<p>--ec-fragment-size <i>EC_FRAGMENT_SIZE</i></p>	<p>The fragment size of erasure code. This parameter takes effect only when LUN data storage redundancy mode is in EC mode, otherwise it is ignored.</p> <p>Value: The value is 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048 or 4096, and the unit is KiB. Consistent with source configuration by default.</p>
<p>-w <i>WRITE_POLICY</i> or</p>	<p>Sets the write policy for the LUN:</p>

--write-policy <i>WRITE_POLICY</i>	<ul style="list-style-type: none"> ● WriteBack (wb): After the data is written to the memory, it returns to the client successfully, and then the data is written to the disk asynchronously. It is suitable for scenarios with high performance requirements and low stability requirements. ● WriteThrough (wt): The data is written to both memory and disk at the same time, and then returns to the client after successful writing. It is suitable for scenarios that the stability requirements are high, the write performance requirements are not high, and the recently written data will be read in a short time. ● WriteAround (wa): Write data directly to the disk without writing to memory. It is suitable for scenarios with high stability requirements, low performance requirements, and more writing and less reading. <p>Consistent with source configuration by default.</p>
-p <i>PATH</i> or --path <i>PATH</i>	<p>Specifies the disk path to store LUN data, only supported by standalone mode.</p> <p>Value: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; . : .</p> <p>If not specified, the default disk path of the server will be used for the LUN.</p>
-m <i>STORAGE_MODE</i> or --mode <i>STORAGE_MODE</i>	<p>Specifies the storage mode for the LUN:</p> <ul style="list-style-type: none"> ● Cache: Store part of hot data at local and store all data in cloud asynchronously. ● Storage: Store all data at local and asynchronously store it in cloud. <p>Consistent with source configuration by default.</p>
--cloud-provider <i>CLOUD_PROVIDER</i>	<p>Specifies the type of object storage:</p> <ul style="list-style-type: none"> ● OOS: eSurfing Cloud Object-Oriented Storage. ● S3: Other object storage compatible with S3. <p>The default value is OOS.</p> <p>Note: The recovered LUN and the source LUN must use the same object storage service.</p>
-B <i>BUCKET_NAME</i> or --bucket <i>BUCKET_NAME</i>	<p>The bucket where the source LUN data is stored.</p> <p>Note: Do not enable bucket lifecycle settings and compliance retention.</p>
-X <i>PREFIX</i> or --prefix <i>PREFIX</i>	<p>The prefix name of the source LUN in object storage. If the source LUN does not specify a prefix, the parameter does not need to be specified.</p> <p>The value is a string of 1 to 256 characters.</p>

-A <i>ACCESS_KEY</i> or --ak <i>ACCESS_KEY</i>	The source LUN's access key of object storage.
-S <i>SECRET_KEY</i> or --sk <i>SECRET_KEY</i>	The source LUN's secret key of object storage.
-C <i>CLOUD_STORAGE_CLASS</i> or --cloud-storage-class <i>CLOUD_STORAGE_CLASS</i>	Specifies storage class of object storage: <ul style="list-style-type: none"> ● STANDARD. ● STANDARD_IA. Consistent with source configuration by default.
-E <i>ENDPOINT</i> or --endpoint <i>ENDPOINT</i>	The source LUN's endpoint. Note: Access will be via HTTPS if only the domain name is entered.
--sign-version <i>VERSION</i>	Specifies signature authentication version: <ul style="list-style-type: none"> ● v2 ● v4 The default value is v2.
--region <i>REGION</i>	Specifies the region of endpoint. If the signature authentication version is v4, this item is required.
-M <i>CLOUD_COMPRESSION</i> or --cloud-compression <i>CLOUD_COMPRESSION</i>	Whether to compress data and upload it to object storage: <ul style="list-style-type: none"> ● Enabled (on): Compress data and upload it to object storage. ● Disabled (off): Do not compress data and upload it to object storage. Consistent with source configuration by default.

Examples

Recover LUN lun02a.

```
[root@songt-0004 CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun recover -n lun02a -t target02 -
-uuid lun-uuid-68fd608b-3654-4bfc-8e6f-d99f92e38d90 --pool default -B hblocktest3 -A
8f129a5529f202811fd0 -S ***** -E oos-cn.ctyunapi.cn
Confirm the information you set. If correct, enter Yes(Y). If not, enter No(N).
UUID: lun-uuid-68fd608b-3654-4bfc-8e6f-d99f92e38d90
LUN Name: lun02a
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap:1
Capacity: 210 GiB
iSCSI Target: target02
Local Sector Size: 4096 Bytes
High Availability: ActiveStandby
Write Policy: WriteBack
Object Storage Info:
+-----+
```

Provider	OOS	
Bucket Name	hblocktest3	
Endpoint	https://oos-cn.ctyunapi.cn	
Signature Version	v2	
Region		
Storage Class	STANDARD	
Access Key	8f129a5529f202811fd0	
Object Size	1024 KiB	
Compression	Enabled	
+-----+-----+-----+		
Y		
Start to recover LUN lun02a.		

4.6.10 Suspend a LUN (Cache/Storage Mode LUN)

```
./stor lun { SUS | suspend } { -n | --name } LUN_NAME [ -f | --force ]
```

This command is used to suspend a cache or a storage mode LUN.

Suspending a LUN is applicable to scenarios where "multiple clusters take turns writing to the same cloud LUN." For example: Cluster A suspends the LUN, then performs cloud data restore on Cluster B, after which Cluster B performs read/write operations on the LUN data. When Cluster B suspends the LUN, Cluster A can then resume the LUN. During the resume process, only incremental data needs to be synchronized, allowing direct access to the latest data.

Note:

- Suspending a LUN will immediately freeze all read and write requests to that LUN. Ensure that all data on the LUN has been persisted; that is, if the LUN has been mounted by a client, ensure that all client data has been synchronized to the LUN. Business dependent on this LUN will be unavailable until the LUN returns to Normal status.
- When a LUN is in Suspended status, only suspend, edit the LUN, edit cloud configuration, resume, delete, and query operations are permitted.
- When a LUN is in Suspended, Suspending, or SuspendFailed status, read and write operations are not supported.
- Before suspending a LUN, you must first log out to disconnect client connections; otherwise, client connections cannot be disconnected until the LUN status returns to Normal.
- This operation can only be performed when the LUN is in Normal, Suspended, Suspending, or SuspendFailed status.

Parameters

Parameter	Description
<code>-n LUN_NAME</code> or <code>--name LUN_NAME</code>	Specifies the name of the LUN that needs to be suspended.
<code>-f</code> or <code>--force</code>	Forcibly suspending the LUN. Note: Forcibly suspending a LUN carries the risk that local data may not be uploaded to the cloud. Please proceed with caution.

Examples

Suspend LUN lunb1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun SUS -n lunb1
```

Start suspending LUN lun1.
You can list lun to check whether the suspension is completed.

4.6.11 Resume a LUN from Restore Interruption or Suspension (Cache/Storage Mode LUN)

```
./stor lun { RS | resume } { -n | --name } LUN_NAME [ --sync-mode SYNC_MODE ] [ -f | --force ]
```

This command is used to resume a LUN from restore interruption or suspension.

Note:

- This operation can only be performed when the LUN is in RecoverFailed or Suspended status.
- Ensure that the LUN data files in the local data directory have not been modified; otherwise, the resumed LUN may contain incomplete data or become unusable.

Parameters

Parameter	Description
-n <i>LUN_NAME</i> or --name <i>LUN_NAME</i>	Specifies the name of the LUN to be resumed from restore interruption or suspension.
--sync-mode <i>SYNC_MODE</i>	Specifies the data scope for resuming a cache-mode LUN. Value: <ul style="list-style-type: none"> ● MetaOnly: Updates only the differences in the LUN index compared to the cloud data index. ● Full: Updates both the differences in the LUN index compared to the cloud, and the differences between locally stored hot data and the cloud. The default value is Full . Note: This parameter applies only to cache mode LUNs. If no local data exists, no local data will be updated.
-f or --force	Forcibly resuming the LUN. Note: Forcibly resuming a LUN carries the risk of partial local data loss. Please proceed with caution.

Examples

Resume LUN lunb1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun RS -n lunb1
Start resuming LUN lunb1.
You can list lun to check whether the recovery is completed.
```

4.6.12 Wipe a LUN (Local Mode LUN)

```
./stor lun { W | wipe} { -n | --name } LUN_NAME [ --scope SCOPE ]
```

This command is used to is used to wipe the specified LUN.

Note:

- If a clone LUN is wiped of all data and its associated snapshots, it will become an independent LUN.
- When a LUN is in Wiping or WipeFailed status, read and write operations are not supported.
- Before wiping a LUN, you must first log out to disconnect client connections; otherwise, client connections cannot be disconnected until the LUN status returns to Normal.

Parameters

Parameter	Description
<code>-n LUN_NAME</code> or <code>--name LUN_NAME</code>	Specifies the name of the LUN that needs to be wiped.
<code>--scope SCOPE</code>	Specifies the data scope for wiping the LUN. Value: <ul style="list-style-type: none"> ● All: Wipes all LUN data and its associated snapshots. Note: If the LUN has associated clone LUNs, this operation cannot be performed. ● NoSnapshot: Wipes LUN data while retaining its associated snapshots. The LUN snapshots can subsequently be used for LUN rollback or backup export. Note: Clone LUNs cannot perform the operation of wiping data only while retaining associated snapshots. The default value is All .

Examples

- Wipe LUN lun04a, including all LUN data and its associated snapshots.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun W -n lun04a
Start wiping LUN lun04a.
You can list lun to check whether the wipe operation is completed.
```

- Wipe LUN lun04b, wiping LUN data while retaining its associated snapshots.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun W -n lun04b --scope NoSnapshot
Start wiping LUN lun04b. Please note that all existing snapshots will be preserved.
You can list lun to check whether the wipe operation is completed.
```

4.6.13 Query LUN Information

```
./stor lun ls [ [ { -n | --name } LUN_NAME [ { -s | --snapshot } SNAPSHOT_FORMAT | --qos ] | --clone | --wwid ]
```

This command is used to query LUN information.

Parameters

Parameter	Description
<code>-n LUN_NAME</code> or <code>--name LUN_NAME</code>	Specifies the LUN name to query.
<code>-s SNAPSHOT_FORMAT</code> or <code>--snapshot SNAPSHOT_FORMAT</code>	Snapshot display format: <ul style="list-style-type: none"> ● list: Displays detailed LUN snapshot information. ● tree: Displays snapshots in a tree structure. Note: You must specify the LUN name to display the LUN snapshot information.
<code>--qos</code>	Specifies to query the QoS policy associated with the LUN. Note: You must specify the LUN name to display the QoS policy information.
<code>--clone</code>	Specifies to query all clone LUNs.
<code>--wwid</code>	Specifies the unique identifier of the query LUN. If there are multiple LUNs on the HBlock side when the client connects to a LUN, the WWID identifier can be used to confirm the LUN to be connected.

Examples

- Standalone mode: Query all LUNs information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls
```

No.	LUN Name	Storage Mode	Capacity	Status	Target	Snapshot Count	Snapshot Size	Is Clone
1.	luna1 (LUN 0)	Local	100 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)	2	471.13 MiB	
2.	luna2 (LUN 1)	Local	1 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)	1	0 B	
3.	luna2-c1 (LUN 2)	Local	1 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)	0	0 B	clone
4.	lunb1 (LUN 0)	Cache	220 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targetb.2 (192.168.0.66:3260,Active)	0	0 B	
5.	lunc1 (LUN 0)	Storage	330 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targetc.3 (192.168.0.66:3260,Active)	0	0 B	
6.	lunc2 (LUN 1)	Storage	1 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targetc.3 (192.168.0.66:3260,Active)	0	0 B	

- Standalone mode: Query all clone LUNs information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls --clone
```

No.	LUN Name	Storage Mode	Capacity	Status	Target	Snapshot Count	Snapshot Size	Is Clone
1.	luna2-c1 (LUN 2)	Local	1 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)	0	0 B	clone

- Standalone Mode: Query the identifiers of all LUNs.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls --wwid
```

No.	LUN Name	Storage Mode	Capacity	Status	Target	Snapshot Count	Snapshot Size	Is Clone	WWID
1.	luna1 (LUN 0)	Local	100 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)	2	471.13 MiB		33000000027321825

2.	luna2 (LUN 1)	Local	1 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)	1	0 B		33fffffff93cdeb
3.	luna2-c1 (LUN 2)	Local	1 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)	0	0 B	clone	33fffffff893d7109
4.	lumb1 (LUN 0)	Cache	220 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targetb.2 (192.168.0.66:3260,Active)	0	0 B		330000003ea215ef
5.	lunc1 (LUN 0)	Storage	330 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targetc.3 (192.168.0.66:3260,Active)	0	0 B		33fffffff85ef3541
6.	lunc2 (LUN 1)	Storage	1 GiB	Normal	iqn.2012-08.cn.ctyunapi.oos:targetc.3 (192.168.0.66:3260,Active)	0	0 B		33fffffff85af22e

● Cluster Mode: Query all LUNs information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls
```

No.	LUN Name	Storage Mode	Capacity	Local Storage Class	Minimum Replicas	Status	Target	Snapshot Count	Snapshot Size	Is Clone
1.	lun01a (LUN 0)	Local	110 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.1 (192.168.0.67:3260,Active)	2	471.36 MiB	
2.	lun01a-C1 (LUN 2)	Local	110 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.2 (192.168.0.64:3260,Standby)	1	4 KiB	clone
3.	lun01b (LUN 1)	Local	150 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.1 (192.168.0.67:3260,Active)	2	0 B	
4.	lun02a (LUN 0)	Cache	210 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target02.3 (192.168.0.65:3260,Active)	0	0 B	
5.	lun03a (LUN 0)	Storage	310 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target03.5 (192.168.0.64:3260,Active)	0	0 B	

● Cluster mode: Query all clone LUNs information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls --clone
```

No.	LUN Name	Storage Mode	Capacity	Local Storage Class	Minimum Replicas	Status	Target	Snapshot Count	Snapshot Size	Is Clone
1.	lun01a-C1 (LUN 2)	Local	110 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.2 (192.168.0.64:3260,Active)	1	4 KiB	clone

● Cluster Mode: Query the identifier information of all LUNs.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls --wid
```

No.	LUN Name	Storage Mode	Capacity	Local Storage Class	Minimum Replicas	Status	Target	Snapshot Count	Snapshot Size	Is Clone	WID
1.	lun01a (LUN 0)	Local	110 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.1 (192.168.0.67:3260,Active)	2	471.36 MiB		33fffffff9c6296b
2.	lun01a-C1 (LUN 2)	Local	110 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.2 (192.168.0.64:3260,Standby)	1	4 KiB	clone	33000000142740b3
3.	lun01a-C2 (LUN 3)	Local	110 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.1 (192.168.0.67:3260,Active)	0	0 B	clone	33000000904a14fb64
4.	lun01b (LUN 1)	Local	150 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.2 (192.168.0.64:3260,Standby)	2	0 B		33fffffff97405c3e
5.	lun02a (LUN 0)	Cache	210 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target02.3 (192.168.0.65:3260,Active)	0	0 B		33fffffff0be07e06

Table 1: Descriptions of information for all queried LUNs.

Item	Description
No.	Number.
LUN Name	Including LUN name and LUN number. The contents in brackets indicate the LUN number. LUN number: The number of the LUN in target, allocated by the storage system, corresponds to the LUN ID in the device address when the client mounts the storage device. If there is only one LUN in the target, the LUN number is generally 0.
Storage Mode	The storage mode for the LUN: <ul style="list-style-type: none"> ● Local: Store all data at local only. ● Cache: Store part of hot data at local and store all data in cloud asynchronously. ● Storage: Store all data at local and asynchronously store it in cloud.
Capacity	The storage capacity of LUN.
Local Storage Class	LUN redundancy mode (only supported by cluster mode): <ul style="list-style-type: none"> ● single-copy ● 2-copy

	<ul style="list-style-type: none"> ● 3-copy ● EC $N+M$ +fragementsize
Minimum Replicas	Minimum replica number (only supported by cluster mode).
Status	LUN status: <ul style="list-style-type: none"> ● Normal: The LUN is normal. ● Deleting: The LUN is being deleted. ● DeleteFailed: The LUN deletion failed. ● Recovering: The LUN is recovering. ● RecoverFailed: The LUN recovered failed. ● Rollbacking: The LUN is rollbacking. ● Flattening: The chain-breaking process between the clone LUN and the snapshot is underway. The clone LUN is copying data from the source LUN. After the copy is complete, it becomes an independent LUN that no longer depends on the snapshot or source LUN. ● Importing: The LUN is being imported with backup data. ● Wiping: The LUN is being wiped. ● WipeFailed: The LUN data wiping failed. ● Suspended: The LUN is suspended. ● Suspending: The LUN is being suspended. ● SuspendFailed: The LUN failed to be suspended.
Target	Target(s) of the LUN, including target IQN, target IP, target port, the status of the target of the LUN (Active, Standby, Offline, ColdStandby).
Snapshot Count	Number of snapshots of the LUN (only supported by local LUN).
Snapshot Size	Total size of snapshots associated with the LUN (only supported by local LUN), representing the amount of data recorded in the snapshots. Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.
Is Clone	Indicates whether the LUN is a clone LUN: <ul style="list-style-type: none"> ● clone: The LUN is a clone LUN. ● Not displayed: The LUN is not a clone LUN.
WWID	Unique identifier of the LUN. If there are multiple LUNs on the HBlock side when the client connects to a LUN, the WWID identifier can be used to confirm the LUN to be connected.

- Standalone Mode: Query the information for local mode LUN luna1 (non-clone LUN).

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n luna1
LUN Name: luna1 (LUN 0)
Storage Mode: Local
Capacity: 100 GiB
```

```
Status: Normal
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)
Create Time: 2025-12-09 14:03:58
Local Sector Size: 4096 Bytes
Data Health: 100% normal, 0% low redundancy, 0% error
Write Policy: WriteBack
WWID: 33000000027321825
UUID: lun-uuid-b5e99a16-1ec0-415b-88d1-f14c4ac8e455
Path: /mnt/stor01
Snapshot Count: 2
Snapshot Size: 471.13 MiB (Note: Snapshot size may vary due to LUN issues or parent
snapshot deletion.)
```

- Standalone Mode: Query the information for local mode LUN luna2-c1 (clone LUN).

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n l luna2-c1
LUN Name: luna2-c1 (LUN 2)
Storage Mode: Local
Capacity: 1 GiB
Status: Normal
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260,Active)
Create Time: 2025-12-22 10:50:46
Local Sector Size: 4096 Bytes
Data Health: 100% normal, 0% low redundancy, 0% error
Write Policy: WriteBack
WWID: 33ffffffff893d7109
UUID: lun-uuid-ce9ad570-c6d9-4c66-ab2a-1316dffdb9a7
Path: /mnt/stor01
Source Snapshot: luna2-snap20251218095041
Source LUN: luna2
Snapshot Count: 0
Snapshot Size: 0 B (Note: Snapshot size may vary due to LUN issues or parent snapshot
deletion.)
```

- Standalone Mode: Query the information for cache mode LUN lunb1

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lunb1
LUN Name: lunb1 (LUN 0)
Storage Mode: Cache
Capacity: 220 GiB
Status: Normal
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targetb.2 (192.168.0.66:3260,Active)
Create Time: 2025-12-09 14:05:05
Local Sector Size: 4096 Bytes
Data Health: 100% normal, 0% low redundancy, 0% error
Write Policy: WriteBack
WWID: 3300000003ea215ef
UUID: lun-uuid-6d66d545-e76b-489d-a9af-b918871b7626
Path: /mnt/stor01
Object Storage Info:
+-----+-----+
```

Provider	OOS	
Bucket Name	hblocktest3	
Endpoint	https://oos-cn.ctyunapi.cn	
Signature Version	v2	
Region		
Storage Class	STANDARD	
Access Key	8f129a5529f202811fd0	
Object Size	1024 KiB	
Compression	Enabled	
+-----+-----+-----+		

- Cluster Mode: Query the information for local mode LUN lun01a (non-clone LUN).

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun01a
LUN Name: lun01a (LUN 0)
Storage Mode: Local
Capacity: 110 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target01.1 (192.168.0.67:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target01.2 (192.168.0.64:3260,Standby)
Create Time: 2025-12-18 09:38:57
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 33ffffffffc9c6296b
UUID: lun-uuid-efa1d8b2-92ee-487e-8f5d-bf0c9d5f3d8d
Snapshot Count: 2
Snapshot Size: 471.36 MiB (Note: Snapshot size may vary due to LUN issues or parent
snapshot deletion.)
```

- Cluster Mode: Query the information for local mode LUN lun01a-C1 (clone LUN).

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun01a-C1
LUN Name: lun01a-C1 (LUN 2)
Storage Mode: Local
Capacity: 110 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target01.2 (192.168.0.64:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target01.1 (192.168.0.67:3260,Standby)
Create Time: 2025-12-18 09:53:04
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
```

```

Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 330000000142740b3
UUID: lun-uuid-28273c97-1517-4146-8aaa-8bdb5b29b51f
Source Snapshot: lun01a-s1
Source LUN: lun01a
Snapshot Count: 1
Snapshot Size: 4 KiB (Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.)

```

- Cluster Mode: Query the information for cache mode LUN lun03a.

```

[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun03a
LUN Name: lun03a (LUN 0)
Storage Mode: Storage
Capacity: 310 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target03.5 (192.168.0.64:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target03.6 (192.168.0.65:3260,Standby)
Create Time: 2025-12-18 09:44:04
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 330000000111c0f4a
UUID: lun-uuid-f7c974e4-e418-4a02-9e53-416410c81b31
Object Storage Info:
+-----+-----+
| Provider      | OOS      |
| Bucket Name   | hblocktest3 |
| Endpoint      | https://oos-cn.ctyunapi.cn |
| Signature Version | v4      |
| Region       | cn       |
| Storage Class | STANDARD |
| Access Key    | 8f129a5529f202811fd0 |
| Object Size   | 1024 KiB |
| Compression   | Enabled   |
+-----+-----+

```

- Cluster Mode: Query the information about the failed recovery of LUN lun04a.

```

[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun04a

```

```

LUN Name: lun04a (LUN 0)
Storage Mode: Storage
Capacity: 1 GiB
Status: RecoverFailed
Failed Reason: Insufficient fault domain for LUN to write (with the redundancy overlap taken into consideration), resolve the problem and resume LUN recovery.
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target04.7 (192.168.0.65:3260,Offline)
               iqn.2012-08.cn.ctyunapi.oos:target04.8 (192.168.0.64:3260,Offline)
Create Time: 2025-12-22 13:56:00
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: pool
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 33ffffffff94c72aed
UUID: lun-uuid-500f8321-5c6b-4468-ad0a-36fe7c096b01
Object Storage Info:
+-----+-----+
| Provider      | OOS      |
| Bucket Name   | hblocktest3 |
| Endpoint      | https://oos-cn.ctyunapi.cn |
| Signature Version | v2      |
| Region       |          |
| Storage Class | STANDARD |
| Access Key    | 8f129a5529f202811fd0 |
| Object Size   | 1024 KiB |
| Compression   | Enabled  |
+-----+-----+

```

Table 2: Description of information for the specified LUN being queried

Item	Description
LUN Name	Including LUN name and LUN number. The contents in brackets indicate the LUN number. LUN number: The number of the LUN in target, allocated by the storage system, corresponds to the LUN ID in the device address when the client mounts the storage device. If there is only one LUN in the target, the LUN number is generally 0.
Storage Mode	The storage mode for the LUN: <ul style="list-style-type: none"> ● Local: Store all data at local only. ● Cache: Store part of hot data at local and store all data in cloud asynchronously. ● Storage: Store all data at local and asynchronously store it in cloud.
Capacity	The storage capacity of LUN.

Status	LUN status: <ul style="list-style-type: none"> ● Normal: The LUN is normal. ● Deleting: The LUN is being deleted. ● DeleteFailed: The LUN deletion failed. ● Recovering: The LUN is recovering. ● RecoverFailed: The LUN recovered failed. ● Rollbacking: The LUN is rollbacking. ● Flattening: The chain-breaking process between the clone LUN and the snapshot is underway. The clone LUN is copying data from the source LUN. After the copy is complete, it becomes an independent LUN that no longer depends on the snapshot or source LUN. ● Importing: The LUN is being imported with backup data. ● Wiping: The LUN is being wiped. ● WipeFailed: The LUN data wiping failed. ● Suspended: The LUN is suspended. ● Suspending: The LUN is being suspended. ● SuspendFailed: The LUN failed to be suspended.
Failed Reason	The reason for the failure of LUN recovery when the LUN status is RecoverFailed, the reason for the failure of deleting cloud data of the LUN when the LUN status is DeleteFailed, the reason why LUN data wiping fails when the LUN status is WipeFailed, or the reason why the LUN failed to be suspended when the LUN status is SuspendFailed.
Auto Failback	When the high-priority server returns to normal, whether to automatically switch active and standby iSCSI targets of the LUN. <ul style="list-style-type: none"> ● Enabled: Automatic switch active and standby iSCSI targets. ● Disabled: No automatic switch active and standby iSCSI targets.
iSCSI Target	Target(s) of the LUN, including target IQN, target IP, target port, the status of the target of the LUN (Active, Standby, Offline, ColdStandby).
Create Time	The creation time of the LUN.
Local Storage Class	LUN redundancy mode (only supported by cluster mode): <ul style="list-style-type: none"> ● single-copy ● 2-copy ● 3-copy ● EC $N+M$ +fragementsize
Minimum Replicas	Minimum replica number (only supported by cluster mode).
Redundancy Overlap	Redundancy overlap (only supported by cluster mode).
Local Sector Size	LUN sector size.
Cache Storage Pool	Cache storage pool (only supported by cluster mode).

Storage Pool	Storage pool (only supported by cluster mode): Indicates the final storage pool in which LUN data is stored.
Data Health	The data health status of the LUN, including: the percentage of normal data (normal), the percentage of low redundancy data (low redundancy), and the percentage of erroneous data (error). If low redundancy data exists, low redundancy reconstruction progress will be provided.
High Availability	High Availability Type of LUN (Supported only in Cluster Edition): <ul style="list-style-type: none"> ● ActiveStandby: Enable active and standby. The LUN is associated with all IQNs under the corresponding target. ● Disabled: Disable active/standby. The LUN is associated with one IQN under the corresponding target.
Write Policy	The write policy of the LUN: <ul style="list-style-type: none"> ● WriteBack: After the data is written to the memory, it returns to the client successfully, and then the data is written to the disk asynchronously. ● WriteThrough: The data is written to both memory and disk at the same time, and then returns to the client after successful writing. ● WriteAround: Write data directly to the disk without writing to memory.
WWID	Unique identifier of the LUN. If there are multiple LUNs on the HBlock side when the client connects to a LUN, the WWID identifier can be used to confirm the LUN to be connected.
UUID	Universally unique identifier of the LUN.
Path	The path to store LUN data (only supported by standalone mode).
Source Snapshot	Snapshot associated with the clone LUN (only supported by clone LUN).
Source LUN	Source LUN name of the clone LUN (only supported by clone LUN).
Snapshot Count	Number of snapshots of the LUN (only supported by local LUN).
Snapshot Size	Total size of snapshots associated with the LUN (only supported by local LUN), representing the amount of data recorded in the snapshots. Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.
Object Storage Info	Cloud information of the LUN, including Provider, Bucket Name, Prefix, Endpoint, Signature Version, Region, Storage Class, Access Key, Object Size, Compression.

Provider	The type of object storage: <ul style="list-style-type: none"> ● OOS: eSurfing Cloud Object-Oriented Storage. ● S3: Other object storage compatible with S3.
Bucket Name	Bucket name of object storage.
Prefix	The prefix name for the LUN in object storage. If it is not set, this item will not be displayed.
Endpoint	Endpoint of object storage.
Signature Version	Signature authentication version: <ul style="list-style-type: none"> ● v2 ● v4
Region	The region of endpoint. If it is a v2 signature, this field should be left blank.
Storage Class	Storage class of object storage: <ul style="list-style-type: none"> ● STANDARD. ● STANDARD_IA.
Access Key	Access key of object storage.
Object Size	The size of data stored in object storage.
Compression	Whether to compress data and upload it to object storage: <ul style="list-style-type: none"> ● Enabled: Compress data and upload it to object storage. ● Disabled: Do not compress data and upload it to object storage.

- Local LUN: Query snapshot information of LUN lun01a, displayed in a tree structure.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun01a -s tree
LUN Name: lun01a (LUN 0)
Storage Mode: Local
Capacity: 110 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target01.1 (192.168.0.67:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos:target01.2 (192.168.0.64:3260,Standby)
Create Time: 2025-12-18 09:38:57
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 33ffffffffc9c6296b
UUID: lun-uuid-efa1d8b2-92ee-487e-8f5d-bf0c9d5f3d8d
Snapshot Count: 2
Snapshot Size: 471.36 MiB (Note: Snapshot size may vary due to LUN issues or parent
snapshot deletion.)
Snapshots:
```

```

root
├── lun01a-s1
│   └── lun01a-snap20251218095146 (*)
    
```

- Local: Query snapshot information of LUN lun01a, showing detailed snapshot information.

```

[root@hbblockserver CTYUN_HBBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun01a -s list
LUN Name: lun01a (LUN 0)
Storage Mode: Local
Capacity: 110 GiB
Status: Normal
Auto Failback: Enabled
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos.target01.1 (192.168.0.67:3260,Active)
               iqn.2012-08.cn.ctyunapi.oos.target01.2 (192.168.0.64:3260,Standby)
Create Time: 2025-12-18 09:38:57
Local Storage Class: EC 2+1+16 KiB
Minimum Replicas: 2
Redundancy Overlap: 1
Local Sector Size: 4096 Bytes
Storage Pool: default
Data Health: 100% normal, 0% low redundancy, 0% error
High Availability: ActiveStandby
Write Policy: WriteBack
WWID: 33ffffffffc9c6296b
UUID: lun-uuid-efad8b2-92ee-487e-8f5d-bf0c9d5f3d8d
Snapshot Count: 2
Snapshot Size: 471.36 MiB (Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.)
Snapshots:
-----+-----+-----+-----+-----+-----+-----+-----+
| No. | Name                               | Size      | Consistency | Snapshot Name | LUN Capacity | Reclaim Policy | Status | Create Time          |
-----+-----+-----+-----+-----+-----+-----+-----+
| 1.  | lun01a-s1                          | 471.36 MiB | -           |                | 110 GiB      | Retain         | Normal | 2025-12-18 09:51:36 |
| 2.  | lun01a-snap20251218095146 (*)     | 0 B       | cons1      |                | 110 GiB      | Retain         | Normal | 2025-12-18 09:51:46 |
-----+-----+-----+-----+-----+-----+-----+-----+
    
```

Table 3: Description of information for the specified LUN's snapshot

Item	Description
LUN Name	Including LUN name and LUN number. The contents in brackets indicate the LUN number. LUN number: The number of the LUN in target, allocated by the storage system, corresponds to the LUN ID in the device address when the client mounts the storage device. If there is only one LUN in the target, the LUN number is generally 0.
Storage Mode	The storage mode for the LUN: ● Local: Store all data at local only.
Capacity	The storage capacity of LUN.
Status	LUN status: ● Normal: The LUN is normal. ● Deleting: The LUN is being deleted. ● DeleteFailed: The LUN deletion failed. ● Recovering: The LUN is recovering. ● RecoverFailed: The LUN recovered failed. ● Rollbacking: The LUN is rollbacking. ● Flattening: The chain-breaking process between the clone LUN and the snapshot is underway. The clone LUN

	<p>is copying data from the source LUN. After the copy is complete, it becomes an independent LUN that no longer depends on the snapshot or source LUN.</p> <ul style="list-style-type: none"> ● Importing: The LUN is being imported with backup data. ● Wiping: The LUN is being wiped. ● WipeFailed: The LUN data wiping failed.
Auto Failback	<p>When the high-priority server returns to normal, whether to automatically switch active and standby iSCSI targets of the LUN.</p> <ul style="list-style-type: none"> ● Enabled: Automatic switch active and standby iSCSI targets. ● Disabled: No automatic switch active and standby iSCSI targets.
iSCSI Target	<p>Target(s) of the LUN, including target IQN, target IP, target port, the status of the target of the LUN (Active, Standby, Offline, ColdStandby).</p>
Create Time	<p>The creation time of the LUN.</p>
Local Storage Class	<p>LUN redundancy mode (only supported by cluster mode):</p> <ul style="list-style-type: none"> ● single-copy ● 2-copy ● 3-copy ● EC $N+M$ +fragementsize
Minimum Replicas	<p>Minimum replica number (only supported by cluster mode).</p>
Redundancy Overlap	<p>Redundancy overlap (only supported by cluster mode).</p>
Local Sector Size	<p>LUN sector size.</p>
Cache Storage Pool	<p>Cache storage pool (only supported by cluster mode).</p>
Storage Pool	<p>Storage pool (only supported by cluster mode): Indicates the final storage pool in which LUN data is stored.</p>
High Availability	<p>High Availability Type of LUN (Supported only in Cluster Edition):</p> <ul style="list-style-type: none"> ● ActiveStandby: Enable active and standby. The LUN is associated with all IQNs under the corresponding target. ● Disabled: Disable active/standby. The LUN is associated with one IQN under the corresponding target.
Write Policy	<p>The write policy of the LUN:</p> <ul style="list-style-type: none"> ● WriteBack: After the data is written to the memory, it returns to the client successfully, and then the data is written to the disk asynchronously. ● WriteThrough: The data is written to both memory and disk at the same time, and then returns to the client after successful writing. ● WriteAround: Write data directly to the disk without writing to memory.

WWID	<p>Unique identifier of the LUN.</p> <p>If there are multiple LUNs on the HBlock side when the client connects to a LUN, the WWID identifier can be used to confirm the LUN to be connected.</p>
UUID	Universally unique identifier of the LUN.
Path	The path to store LUN data (only supported by standalone mode).
Source Snapshot	Snapshot associated with the clone LUN (only supported by clone LUN).
Source LUN	Source LUN name of the clone LUN (only supported by clone LUN).
Snapshot Count	Number of snapshots of the LUN (only supported by local LUN).
Snapshot Size	<p>Total size of snapshots associated with the LUN (only supported by local LUN), representing the amount of data recorded in the snapshots.</p> <p>Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.</p>
Snapshots	Snapshot information.
	<p>Snapshot display in detail includes:</p> <ul style="list-style-type: none"> ● No.: Serial number. ● Name: Snapshot name. *: Indicates that the current read/write operations are based on this snapshot. If there is no * before any name, it indicates that the current read/write operations are based on the initial state of the LUN. ● Size: Snapshot size, representing the actual amount of data recorded in the snapshot. The unit is bytes. Note: Snapshot size may vary due to LUN issues or parent snapshot deletion. ● Consistency Snapshot Name: The name of the consistency snapshot it belongs to. Shown only for single-LUN snapshot in a consistency snapshot; "-" indicates it's not part of a consistency snapshot. ● LUN Capacity: Source LUN capacity at snapshot creation. ● Reclaim Policy: Snapshot reclaim policy: <ul style="list-style-type: none"> ■ Delete: Automatically deleted when the snapshot has no associated clone LUNs and has at most one child node (indicating no other snapshots depend on it or current write operations aren't based on it). ■ Retain: Retained when the snapshot has no associated clone LUNs and has at most one child node.

	<ul style="list-style-type: none"> ● Status: Snapshot status: <ul style="list-style-type: none"> ■ Normal. ■ Error. ■ Pending: The snapshot is being created. ■ Deleting. ● Create Time: Snapshot creation time.
	<p>Snapshots are displayed in a tree structure, showing only the snapshot names.</p> <p>*: Indicates that the current read/write operations of the LUN are based on this snapshot.</p>

- Standalone mode: Query the QoS policy associated with LUN luna1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n luna1 --qos
The specific QoS policy for LUN luna1
QoS Policy Name: QoS1
Reclaim Policy: Retain
IOPS (Total/Read/Write): 2500/1000/1000
Bandwidth (Total/Read/Write): 4.88 GiB/s | 3.42 GiB/s | 3.42 GiB/s
Burst IOPS (Total/Read/Write): 2000/1000/1500
Burst Bandwidth (Total/Read/Write): 4.88 GiB/s | 3.42 GiB/s | 3.42 GiB/s
Burst Duration (second): 1/1/1/1/1/1
Create Time: 2025-08-05 14:29:22
Description: It is QoS1.
```

- Cluster mode: Query the QoS policy associated with LUN lun01a.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor lun ls -n lun01a --qos
The specific QoS policy for LUN lun01a
QoS Policy Name: QoS6
Reclaim Policy: Retain
IOPS (Total/Read/Write): 1000/600/400
Bandwidth (Total/Read/Write): Unlimited
Burst IOPS (Total/Read/Write): 10000/6000/4000
Burst Bandwidth (Total/Read/Write): Unlimited
Burst Duration (second): 1/1/1/1/1/1
Create Time: 2025-07-31 16:28:51
Description: It is QoS6.

The default QoS policy in pool for LUN lun01a
Storage Pool Name: defaultpool
QoS Policy Name: Qos-Test
Reclaim Policy: Retain
IOPS (Total/Read/Write): 8000/Unlimited/Unlimited
Bandwidth (Total/Read/Write): 7.81 GiB/s | Unlimited | Unlimited
Burst IOPS (Total/Read/Write): Unlimited
Burst Bandwidth (Total/Read/Write): Unlimited
Burst Duration (second): 1/1/1/1/1/1
```

Create Time: 2025-08-06 17:02:32

Table 4: Description of the QoS policy associated with the LUN

Item	Description
The specific QoS policy for LUN <i>lunName</i>	The specific QoS policy for LUN.
The default QoS policy in pool for LUN <i>lunName</i>	Default QoS policy for LUNs within the storage pool (cluster edition only). <ul style="list-style-type: none"> ● When the LUN belongs to both a cache pool and a storage pool, the cache pool's "default QoS policy for LUNs in the pool" is displayed; if the cache pool has not set this policy, the field is omitted. ● When the LUN resides in only a storage pool, that pool's "default QoS policy for LUNs in the pool" is shown; if the pool has not set this policy, the field is omitted.
QoS Policy Name	QoS policy name.
Reclaim Policy	QoS policy reclaim policy: <ul style="list-style-type: none"> ● Delete: When all objects associated with the QoS policy are disassociated or deleted, the system will automatically remove the QoS policy. ● Retain: When all objects associated with the QoS policy are disassociated or deleted, the QoS policy itself remains and is not deleted.
IOPS (Total/Read/Write)	The limit of total/read/write operations per second.
Bandwidth (Total/Read/Write)	The limit of total/read/write throughput per second.
Burst IOPS (Total/Read/Write)	The burst limit of total/read/write operations per second.
Burst Bandwidth (Total/Read/Write)	The burst limit of total/read/write throughput per second.
Burst Duration (second)	The duration in seconds of I/O operations with burst limit. The corresponding sequence of QoS policy parameters is: burst IOPS (total/read/write), burst bandwidth (total/read/write).
Create Time	The creation time of the QoS policy.
Description	The description of the QoS policy.
Storage Pool Name	Name of the storage pool to which the LUN belongs, only supported by cluster mode.

4.7 Snapshot Operations

4.7.1 Create a Snapshot

```
./stor snapshot add { -n | --name } SNAPSHOT_NAME { -l | --lun } LUN_NAME [ --reclaim-policy RECLAIM_POLICY ]
```

This command is used to create a snapshot.

Note:

- Before performing this operation, ensure all data on the source LUN is persistent. If the source LUN is mounted by a client, make sure the client's data has been synchronized to the LUN. Before creating a snapshot:
 - For Linux clients: If the client supports **sync -f** (check with **sync --help**), run **sync -f**; otherwise, run **sync**.
 - For Windows clients: Take the disk corresponding to the source LUN offline on the client before the snapshot, and bring it back online afterward.
- The source LUN needs to be in Normal status.

Usage restrictions:

- Maximum snapshots per LUN: 512.
- Maximum snapshots per system: 100,000.
- Maximum clones per snapshot: 512.
- Maximum snapshot depth per system: 512.

Parameters

Parameter	Description
-n <i>SNAPSHOT_NAME</i> or --name <i>SNAPSHOT_NAME</i>	Specifies the snapshot name. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit. Note: Snapshot names, consistency snapshot names, and names of single-LUN snapshots in a consistency snapshot must be unique and not duplicated across each other.
-l <i>LUN_NAME</i> or --lun <i>LUN_NAME</i>	Specifies the source LUN for the snapshot. Note: Snapshots can only be created for local LUNs.
--reclaim-policy <i>RECLAIM_POLICY</i>	Specifies snapshot reclaim policy: <ul style="list-style-type: none"> ● Delete: Automatically deleted when the snapshot has no associated clone LUNs and

	<p>has at most one child node (indicating no other snapshots depend on it or current write operations aren't based on it).</p> <p>Note: If the snapshot reclaim policy is set to Delete, the following operations will trigger snapshot deletion: rolling back a consistency snapshot, rolling back a snapshot, flattening a clone LUN, deleting a clone LUN, wiping a LUN, and deleting a snapshot.</p> <ul style="list-style-type: none"> ● Retain: Retained when the snapshot has no associated clone LUNs and has at most one child node. <p>The default value is Retain.</p>
--	---

Examples

Create a snapshot named snapshot-lun01a for LUN lun01a.

Note: You can input a snapshot description of up to 256 characters. If you don't want to add a description, just press Enter.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor snapshot add -n snapshot-lun01a -
l lun01a
Enter the description for this snapshot, limited to 256 characters:
lun01a's snapshot
Start creating snapshot snapshot-lun01a of LUN lun01a. You can list snapshot to check
whether it is completed.
```

4.7.2 Modify a Snapshot

```
./stor snapshot set { -n | --name } SNAPSHOT_NAME [ --change-description ] [ --new-name NEW_NAME ] [ --reclaim-policy RECLAIM_POLICY ]
```

This command is used to modify the snapshot information, including single-LUN snapshot in a consistency snapshot.

Parameters

Parameter	Description
<code>-n SNAPSHOT_NAME</code> or <code>--name SNAPSHOT_NAME</code>	Specifies the snapshot name to be modified. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit.
<code>--change-description</code>	Modifies the description of the snapshot.
<code>--new-name NEW_NAME</code>	Specifies a new name for the snapshot name. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit. Note: Snapshot names, consistency snapshot names, and names of single-LUN snapshots in a consistency snapshot must be unique and not duplicated across each other.
<code>--reclaim-policy RECLAIM_POLICY</code>	Modifies the snapshot reclaim policy: <ul style="list-style-type: none"> ● Delete: Automatically deleted when the snapshot has no associated clone LUNs and has at most one child node (indicating no other snapshots depend on it or current write operations aren't based on it). ● Retain: Retained when the snapshot has no associated clone LUNs and has at most one child node.

Examples

- Modify the description of snapshot snapshot-lun01a.

Note: Enter the snapshot description interactively, as a string of 1-256 characters.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor snapshot set -n snapshot-lun01a -
-change-description
Enter the description for this snapshot, limited to 256 characters:
The first snapshot of lun01a.
Set snapshot snapshot-lun01a successfully
```

- Rename snapshot snapshot-lun01a to snapshot-lun01a-1 and set the retention policy to Delete.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor snapshot set -n snapshot-lun01a -  
-new-name snapshot-lun01a-1 --reclaim-policy Delete  
Set snapshot snapshot-lun01a successfully.
```

4.7.3 Roll Back a Snapshot

```
./stor snapshot { R | rollback } { -n | --name } SNAPSHOT_NAME
```

This command is used to roll back the snapshot, including single-LUN snapshot in a consistency snapshot.

Prerequisite: Both the snapshot to be rolled back and the source LUN must be in Normal status.

Note:

- This operation overwrites the source LUN's data with the snapshot's data. It is recommended to create a new snapshot of the source LUN to back up the data.
- If the source LUN of the snapshot is a clone LUN and is flattening operation, rollback using any snapshot of this clone LUN is not allowed.
- If the source LUN has a snapshot being created, you cannot perform the rollback operation.
- If the source LUN is mounted by a client, unmount it before rollback and remount it afterward:
 - For Linux clients:
 1. Before rollback, run the command on the client: **umount** *DIRECTORY_NAME_OR_PATH*
 2. After rollback, run the command on the client: **mount** */dev/sdx* *DIRECTORY_NAME_OR_PATH*
 - For Windows clients:
 1. Before rollback, take the corresponding disk offline on the client.
 2. After rollback, bring the corresponding disk back online on the client.

For a LUN in rollback status:

- Read and write operations are not allowed.
- No new snapshots can be created.
- No further rollback operations are allowed.
- You cannot delete the snapshot being rolled back.
- The snapshot cannot be modified.
- The LUN cannot be modified.
- The LUN cannot be expanded.
- If it is a clone LUN, no relationship chain break operation can be performed.

Parameters

Parameter	Description
<code>-n SNAPSHOT_NAME</code> or <code>--name SNAPSHOT_NAME</code>	Specifies the snapshot name. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit.

Examples

Roll back snapshot snapshot-lun01a-2.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor snapshot R -n snapshot-lun01a-2  
DANGER: This operation will overwrite data on the source LUN with that on the snapshot.  
Start rollbacking snapshot snapshot-lun01a-2 on LUN lun01a.
```

4.7.4 Delete a Snapshot

```
./stor snapshot rm { -n | --name } SNAPSHOT_NAME
```

This command is used to delete a snapshot, including single-LUN snapshot in a consistency snapshot.

Note:

- If a snapshot has clone LUNs that are still associated, it cannot be deleted.
- If a snapshot has two or more child nodes (indicating other snapshots depend on it or current writes are based on it), it cannot be deleted.
- When a snapshot is in the process of being deleted, only queries and re-deletion of the snapshot are allowed.

Parameters

Parameter	Description
<code>-n SNAPSHOT_NAME</code> or <code>--name SNAPSHOT_NAME</code>	Specifies the snapshot name to be deleted.

Examples

Delete snapshot snapshot-lun01a-2.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor snapshot rm -n snapshot-lun01a-2
WARNING: This operation cannot be undone and will delete all snapshot information.
Start deleting snapshot snapshot-lun01a-2.
```

4.7.5 Query Snapshot Information

```
./stor snapshot ls [ { -n | --name } SNAPSHOT_NAME ]
```

This command is used to query snapshot information.

Parameters

Parameter	Description
<code>-n SNAPSHOT_NAME</code> or <code>--name SNAPSHOT_NAME</code>	Specifies the snapshot name to be queried.

Examples

- Query all snapshots information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor snapshot ls
```

No.	Name	Size	LUN Name	LUN Capacity	Consistency Snapshot Name	Clone LUN Count	Reclaim Policy	Status
1.	luna1-s1	279.89 MiB	luna1	100 GiB	-	0	Retain	Normal
2.	luna1-snap20251218095041	191.25 MiB	luna1	100 GiB	cons1	0	Retain	Normal
3.	luna1-snapshot	44 KiB	luna1	100 GiB	-	1	Retain	Normal
4.	luna2-snap20251218095041	0 B	luna2	1 GiB	cons1	1	Retain	Normal

Table 1: Descriptions of information for all queried snapshots.

Item	Description
No.	Number
Name	The snapshot name.
Size	Snapshot size, representing the actual amount of data recorded in the snapshot. The unit is bytes. Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.
LUN Name	The source LUN name.
LUN Capacity	Source LUN capacity at snapshot creation.
Consistency Snapshot Name	The name of the consistency snapshot it belongs to. Note: Shown only for single-LUN snapshot in a consistency snapshot; "-" indicates it's not part of a consistency snapshot.
Clone LUN Count	The number of clone LUNs associated with the snapshot.
Reclaim Policy	Snapshot reclaim policy: <ul style="list-style-type: none"> ● Delete: Automatically deleted when the snapshot has no associated clone LUNs and has at most one child node (indicating no other snapshots depend on it or current write operations aren't based on it). ● Retain: Retained when the snapshot has no associated clone LUNs and has at most one child node.

Status	Snapshot status: <ul style="list-style-type: none"> ● Normal. ● Error. ● Pending: The snapshot is being created. ● Deleting.
---------------	--

- Query information for snapshot snap1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor snapshot ls -n snap1
Name: snap1
Size: 88.74 MiB (Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.)
LUN Name: lun01a
LUN Capacity: 200 GiB
Reclaim Policy: Retain
Status: Normal
Create Time: 2025-07-30 14:53:30
Clones:
+-----+-----+-----+
| No. | Name          | Capacity | Status |
+-----+-----+-----+
| 1.  | lun01a-clone1 | 200 GiB  | Normal |
+-----+-----+-----+
```

- Query information for single-LUN snapshot lun01a-snap20250801171358 in a consistency snapshot

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor snapshot ls -n lun01a-snap20250801171358
Name: lun01a-snap20250801171358
Size: 191.3 MiB (Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.)
LUN Name: lun01a
LUN Capacity: 200 GiB
Reclaim Policy: Retain
Consistency Snapshot Name: conssnap1
Status: Normal
Create Time: 2025-08-01 17:13:59
```

Table 2 Descriptions of information for the specified snapshot being queried

Item	Description
Name	The snapshot name
Size	Snapshot size, representing the actual amount of data recorded in the snapshot. The unit is bytes. Note: Snapshot size may vary due to LUN issues or parent snapshot deletion.
Description	Snapshot description.
LUN Name	The source LUN name.

LUN Capacity	Source LUN capacity at snapshot creation.
Reclaim Policy	Snapshot reclaim policy: <ul style="list-style-type: none"> ● Delete: Automatically deleted when the snapshot has no associated clone LUNs and has at most one child node (indicating no other snapshots depend on it or current write operations aren't based on it). ● Retain: Retained when the snapshot has no associated clone LUNs and has at most one child node.
Consistency Snapshot Name	The name of the consistency snapshot it belongs to. Note: Shown only for single-LUN snapshot in a consistency snapshot.
Status	Snapshot status: <ul style="list-style-type: none"> ● Normal. ● Error. ● Pending: The snapshot is being created. ● Deleting.
Create Time	The creation time of the snapshot.
Clones	Information of the clone LUN associated with the snapshot: <ul style="list-style-type: none"> ● No.: Number ● Name: The clone LUN name. ● Capacity: The clone LUN capacity. ● Status: The clone LUN status: <ul style="list-style-type: none"> ■ Normal. ■ Deleting: The clone LUN is being deleting. ■ DeleteFailed: The LUN deletion failed. ■ Rollbacking: The LUN is rollbacking. ■ Importing: The LUN is being imported with backup data. ■ Wiping: The LUN is being wiped. ■ WipeFailed: The LUN data wiping failed. ■ Flattening: The chain-breaking process between the clone LUN and the snapshot is underway. The clone LUN is copying data from the source LUN. After the copy is complete, it becomes an independent LUN that no longer depends on the snapshot or source LUN.

4.8 Consistency Snapshot Operation

4.8.1 Create a Consistency Snapshot

```
./stor conssnap add { -n | --name } CONSISTENCYSNAPSHOT_NAME { -l | --lun } LUN_NAME &<1-n> [ --reclaim-policy RECLAIM_POLICY ]
```

This command is used to create a consistency snapshot.

A consistency snapshot involves creating snapshots for all selected LUNs simultaneously at a specific moment, ensuring the snapshots reflect the data status of the same point in time.

LUN snapshots created via a consistency snapshot are by default named: *LUNName-snap-timestamp* (in seconds), e.g., *lun01-snap20240601120000*. You can't customize the name at creation but can rename the snapshot, for detail, see **Modify a Snapshot**.

Usage restrictions: Maximum number of LUNs for a consistency snapshot: 512.

Note:

- Before performing this operation, ensure all data on the source LUN is persistent. If the source LUN is mounted by a client, make sure the client's data has been synchronized to the LUN. Before creating a consistency snapshot:
 - For Linux clients: If the client supports **sync -f** (check with **sync --help**), run **sync -f**; otherwise, run **sync**.
 - For Windows clients: Take the disk corresponding to the source LUN offline on the client before the snapshot, and bring it back online afterward.
- The source LUNs need to be in Normal status.

Parameters

Parameter	Description
-n <i>CONSISTENCYSNAPSHOT_NAME</i> or --name <i>CONSISTENCYSNAPSHOT_NAME</i>	Specifies the consistency snapshot name. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit. Note: Snapshot names, consistency snapshot names, and names of single-LUN snapshots in a consistency snapshot must be unique and not duplicated across each other.
-l <i>LUN_NAME</i> &<1-n> or --lun <i>LUN_NAME</i> &<1-n>	Specifies source LUNs for the consistency snapshot. Note: <ul style="list-style-type: none"> ● Consistency snapshot can only be created for local LUNs. ● Maximum number of LUNs for a consistency snapshot: 512.

<p>--reclaim-policy <i>RECLAIM_POLICY</i></p>	<p>Specifies reclaim policy for single-LUN snapshots in the consistency snapshot:</p> <ul style="list-style-type: none"> ● Delete: Automatically deleted when the snapshot has no associated clone LUNs and has at most one child node (indicating no other snapshots depend on it or current write operations aren't based on it). ● Retain: Retained when the snapshot has no associated clone LUNs and has at most one child node. <p>The default value is Retain.</p>
--	---

Examples

Create a consistency snapshot named `consistencysnapshot3` for LUNs `lun01`, `lun01a`, `lun02`, and `lun03`.

Note: You can input a consistency snapshot description of up to 256 characters. If you don't want to add a description, just press Enter.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor consnap add -n
consistencysnapshot3 -l lun01,lun01a,lun02,lun03
Enter the description for this consistency snapshot, limited to 256 characters:
The consistencysnapshot of lun01,lun01a,lun02,lun03.
Start creating consistency snapshot consistencysnapshot3. You can list consistency
snapshot to check whether it is completed.
```

4.8.2 Modify a Consistency Snapshot

```
./stor conssnap set { -n | --name } CONSISTENCYSNAPSHOT_NAME [ --change-description ]
[ --new-name NEW_NAME ] [ --reclaim-policy RECLAIM_POLICY ]
```

This command is used to modify the consistency snapshot information.

Parameters

Parameter	Description
-n <i>CONSISTENCYSNAPSHOT_NAME</i> or --name <i>CONSISTENCYSNAPSHOT_NAME</i>	Specifies the consistency snapshot name to be modified. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit.
--change-description	Modifies the description of the consistency snapshot.
--new-name <i>NEW_NAME</i>	Specifies a new name for the consistency snapshot name. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit. Note: Snapshot names, consistency snapshot names, and names of single-LUN snapshots in a consistency snapshot must be unique and not duplicated across each other.
--reclaim-policy <i>RECLAIM_POLICY</i>	Modifies reclaim policy for single-LUN snapshots in the consistency snapshot: <ul style="list-style-type: none"> ● Delete: Automatically deleted when the snapshot has no associated clone LUNs and has at most one child node (indicating no other snapshots depend on it or current write operations aren't based on it). ● Retain: Retained when the snapshot has no associated clone LUNs and has at most one child node.

Examples

- Modify the description of consistency snapshot consistencysnapshot5.

Note: Enter the consistency snapshot description interactively, as a string of 1-256 characters.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor consnap set -n
consistencysnapshot5 --change-description
Enter the description for this consistency snapshot, limited to 256 characters:
The consistencysnapshot of lun01a, lun02.
Set consistency snapshot consistencysnapshot5 successfully.
```

- Rename consistency snapshot consistencysnapshot3 to consistencysnapshot4 and set the retention policy to Delete.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor consnap set -n
consistencysnapshot3 --new-name consistencysnapshot4 --reclaim-policy Delete
Set consistency snapshot consistencysnapshot3 successfully.
```

4.8.3 Roll Back a Consistency Snapshot

```
./stor conssnap { R | rollback } { -n | --name } CONSISTENCYSNAPSHOT_NAME
```

This command is used to roll back the consistency snapshot.

Prerequisite: Both the consistency snapshot to be rolled back and source LUNs must be in Normal status.

Note:

- This operation overwrites the source LUN's data with the consistency snapshot's data. It is recommended to create a new snapshot or consistency snapshot of source LUNs to back up the data.
- If the source LUN is mounted by a client, unmount it before rollback and remount it afterward:
 - For Linux clients:
 1. Before rollback, run the command on the client: **umount** *DIRECTORY_NAME_OR_PATH*
 2. After rollback, run the command on the client: **mount** */dev/sdx* *DIRECTORY_NAME_OR_PATH*
 - For Windows clients:
 1. Before rollback, take the corresponding disk offline on the client.
 2. After rollback, bring the corresponding disk back online on the client.
- If the source LUN of the consistency snapshot is a clone LUN and is flattening operation, rollback using any snapshot of this clone LUN is not allowed.

If a single-LUN snapshot in a consistency snapshot is in the deletion process, you can still roll back other snapshots in normal status within the same consistency snapshot. The snapshot in the deletion status won't be rolled back.

Parameters

Parameter	Description
-n <i>CONSISTENCYSNAPSHOT_NAME</i> or --name <i>CONSISTENCYSNAPSHOT_NAME</i>	Specifies the consistency snapshot name. The value is a string of 1 to 256 case-sensitive characters. It can contain letters, digits, hyphens (-), or underscores (_). Only supports starting with a letter or a digit.

Examples

Roll back consistency snapshot consistencysnapshot4.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor conssnap R -n  
consistencysnapshot4  
DANGER: This operation will overwrite data on source LUNs with that on the consistency  
snapshot.  
Start rollbacking consistency snapshot consistencysnapshot4 on LUN  
lun01,lun01a,lun02,lun03.
```

4.8.4 Delete a Consistency Snapshot

```
./stor conssnap rm { -n | --name } CONSISTENCYSNAPSHOT_NAME
```

This command is used to delete a consistency snapshot.

Note:

- If a single-LUN snapshot in a consistency snapshot has clone LUNs that are still associated, the consistency snapshot cannot be deleted.
- Deleting a consistency snapshot will delete all single-LUN snapshots under it.
- When a consistency snapshot is in the process of being deleted, only query is allowed.

Parameters

Parameter	Description
<code>-n</code> <i>CONSISTENCYSNAPSHOT_NAME</i> or <code>--name</code> <i>CONSISTENCYSNAPSHOT_NAME</i>	Specifies the consistency snapshot name to be deleted.

Examples

Delete consistency snapshot `consistencysnapshot2`.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor conssnap rm -n  
consistencysnapshot2  
WARNING: This operation cannot be undone and will delete all consistency snapshot  
information.  
Start deleteing consistency snapshot consistencysnapshot2.
```

4.8.5 Query Consistency Snapshot

```
./stor conssnap ls [ { -n | --name } CONSISTENCYSNAPSHOT_NAME ]
```

This command is used to query consistency snapshot information.

Parameters

Parameter	Description
<code>-n CONSISTENCYSNAPSHOT_NAME</code> or <code>--name CONSISTENCYSNAPSHOT_NAME</code>	Specifies the consistency snapshot name to be queried.

Examples

- Query all consistency snapshots information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor conssnap ls
+-----+-----+-----+-----+-----+
| No. | Name | Status | LUN Snapshot Count | Create Time |
+-----+-----+-----+-----+-----+
| 1. | cons1 | Normal | 2 | 2025-12-18 09:51:46 |
| 2. | cons2 | Normal | 2 | 2025-12-18 10:40:54 |
+-----+-----+-----+-----+-----+
```

Table 1: Descriptions of information for all queried consistency snapshots.

Item	Description
No.	Number
Name	The consistency snapshot name.
Status	Consistency snapshot status: <ul style="list-style-type: none"> ● Normal. ● Error. ● Pending: The consistency snapshot is being created. ● Deleting.
LUN Snapshot Count	Number of single-LUN snapshots in the consistency snapshot.
Create Time	The creation time of the consistency snapshot.

- Query information for consistency snapshot cons1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor conssnap ls -n cons1
Name: cons1
Status: Normal
Create Time: 2025-12-18 09:51:46
LUN Snapshot Count: 2
```

Lun Snapshots:						
No.	Name	LUN Name	LUN Capacity	Clone LUN Count	Reclaim Policy	Status
1.	lun01a-snap20251218095146	lun01a	110 GiB	1	Retain	Normal
2.	lun01b-snap20251218095146	lun01b	150 GiB	0	Retain	Normal

Table 2 Descriptions of information for the specified consistency snapshot being queried

Item	Description
Name	The consistency snapshot name
Status	Consistency snapshot status: <ul style="list-style-type: none"> ● Normal. ● Error. ● Pending: The consistency snapshot is being created. ● Deleting.
Description	Consistency snapshot description.
Create Time	The creation time of the consistency snapshot.
LUN Snapshot Count	Number of LUN snapshots in the consistency snapshot.
Lun Snapshots	The information of single-LUN snapshots in the consistency snapshot: <ul style="list-style-type: none"> ● No.: Number ● Name: The single-LUN snapshot name. ● LUN Name: The source LUN name. ● LUN Capacity: Source LUN capacity at snapshot creation. ● Clone LUN Count: The number of clone LUNs. ● Reclaim Policy: Snapshot reclaim policy: <ul style="list-style-type: none"> ■ Delete: Automatically deleted when the snapshot has no associated clone LUNs and has at most one child node (indicating no other snapshots depend on it or current write operations aren't based on it). ■ Retain: Retained when the snapshot has no associated clone LUNs and has at most one child node. ● Status: Snapshot status: <ul style="list-style-type: none"> ■ Normal. ■ Error. ■ Pending. ■ Deleting.

4.9 Backup Operation (Local LUN)

4.9.1 Export Backup File

```
./stor backup { E | export } [ { -f | --from-snap } FROM_SNAPSHOT_NAME ] { -s | --snap } SNAPSHOT_NAME { -o | --out } OUTPUT_FILE [ -M | --compression ] [ --allow-modify ]
```

This command is used to export backup file.

Backups are divided into full backups and incremental backups:

- Full backup: Exports all data from the LUN's creation up to the selected snapshot into a backup file.
- Incremental backup: Exports only the changed data between two snapshots into a backup file. For example, creating an incremental backup from snap1 to snap2 exports the LUN's data between the time of snap1 and the time of snap2.

Note:

- Backups can only be exported when all snapshots are in Normal status.
- Backups can only be exported when the LUN is in Normal, Flattening, or Rollbacking status.
- When exporting a backup, the associated snapshots and LUN must not be deleted.
- Each server can export a maximum of 4 backups simultaneously. The server here refers to the server where the snapshot source LUN's Active IQN resides.

Parameters

Parameter	Description
<code>-f FROM_SNAPSHOT_NAME</code> or <code>--from-snap FROM_SNAPSHOT_NAME</code>	Snapshot name, indicating that data created after this snapshot point will be exported. Note: This parameter must be provided for an incremental backup, where it serves as the starting snapshot. If omitted, the backup starts from the beginning of the LUN.
<code>-s SNAPSHOT_NAME</code> or <code>--snap SNAPSHOT_NAME</code>	Snapshot name, indicating that the data captured at the snapshot point will be exported.
<code>-o OUTPUT_FILE</code> or <code>--out OUTPUT_FILE</code>	Specifies the backup file path and filename. Note: Read and write permissions for this path are required.
<code>-M</code> or <code>--compression</code>	Compress the backup file.
<code>--allow-modify</code>	Allow the current backup task to overwrite an existing backup file with the same name. Note: If a file with the same name exists but was not generated by an HBlock export task, it cannot be overwritten. If a backup file with the same name already exists:

- If compression is enabled for this export, the backup file will be re-exported from the beginning and overwrite the existing file.
- If compression is not enabled for this export:
 - If the existing backup file is compressed, it will be re-exported from the beginning and overwrite the existing file.
 - If the existing backup file is not compressed, the system will determine the point of interruption based on the size of the existing file, resume the export from that point, and continue writing to the same file.

Examples

Export a full backup file.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor backup E --snap lun02a-snap1 --out /mnt/stor01/lun02a-snap1 --compression
Exporting backup file: 100%
Exported backup file successfully.
```

Export an incremental backup file.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor backup E --from-snap lun02a-snap1 --snap lun02a-snap2 --out /mnt/stor01/luna2-snap1_luna2-snap2 --compression --allow-modify
Exporting backup file: 100%
Exported backup file successfully.
```

4.9.2 Import Backup File

```
./stor backup { I | import } { -f | --file } FILE { -l | --lun } LUN_NAME [ -w | --wipe ] [ --wipe-scope WIPE_SCOPE ] ]
```

This command is used to import a backup file.

When importing a full backup, it is recommended to use a new LUN, or ensure that the existing data on the LUN has been wiped. The LUN data can be wiped using the wipe LUN function; decide whether to retain the LUN's snapshot data based on actual requirements.

Note:

- Only one backup file can be imported at a time. If multiple backup files are involved, import them in the exact order of their snapshot creation.
- The target LUN must be free of any read or write operations while the backup is being imported.
- While importing backup files—especially when dealing with multiple full and incremental backups—do not mount the target LUN to any client until all required backups have been fully imported. Premature mounting may cause file-system corruption.

Only mount the LUN after every planned backup has been imported, and do not import any additional backups once the LUN is mounted.

If the LUN is already mounted, unmount it first, wipe the LUN, and then proceed with the backup import.

- If the user splits the full backup file into multiple segments on their own before performing the import, they must ensure that there is data overlap between the previous segment and the next one, and that the starting position of each segment is a multiple of 256 MiB. For example, the first segment is 0–515 MiB, and the second segment is 512 MiB–1024 MiB.

Parameters

Parameter	Description
<code>-f FILE</code> or <code>--file FILE</code>	The backup file to be imported.
<code>-l LUN_NAME</code> or <code>--lun LUN_NAME</code>	Name of the LUN into which the backup will be imported. Note: The LUN must be larger than the source LUN's size at the moment the snapshot was taken.
<code>-w</code> or <code>--wipe</code>	Wipe the LUN and retain only the imported backup data. Note: <ul style="list-style-type: none"> ● For LUNs that have already been written with data, it is recommended to wipe them, retaining only the data imported from the backup files.

	<ul style="list-style-type: none"> ● For scenarios involving multiple imports using full backup files and incremental backup files: select to wipe existing LUN data during the first import. For subsequent imports, existing LUN data must be retained. ● Newly created LUNs are not affected by this parameter.
--wipe-scope <i>WIPE_SCOPE</i>	Specifies the data scope for wiping the LUN. Value: <ul style="list-style-type: none"> ● All: Wipes all LUN data and its associated snapshots. Note: If the LUN has associated clone LUNs, this operation cannot be performed. ● NoSnapshot: Wipes LUN data while retaining its associated snapshots. The LUN snapshots can subsequently be used for LUN rollback or backup export. Note: Clone LUNs cannot perform the operation of wiping data only while retaining associated snapshots. The default value is All .

Examples

- Import the full backup file.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor backup I -l lun02d -f
/mnt/stor01/lun02a-snap1 -w
Wiping LUN lun02d.
Wiped LUN successfully.
Sending backup file: 100%
Importing.
Imported backup file successfully.
```

- Import the incremental backup file.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor backup I -l lun02d -f
/mnt/stor01/luna2-snap1_luna2-snap2
Sending backup file: 100%
Importing.
Imported backup file successfully.
```

4.10 iSCSI Target Operations

4.10.1 Create an iSCSI Target

Standalone Mode

```
./stor target add { -n | --name } TARGET_NAME [ --max-sessions MAX_SESSIONS ] [ --reclaim-policy RECLAIM_POLICY ] [ { -c | --chap-name } CHAP_NAME { -p | --password } CHAP_PASSWORD { -s | --status } STATUS ] [ [ --initiator [ IP<1-n>][:NAME<1-n> ] [ IP<1-n>][:NAME<1-n> ] ] [ --target [ IP<1-n>][:NIC<1-n> ] [ IP<1-n>][:NIC<1-n> ] ] | --allow-file ALLOW_FILE ]
```

Cluster Mode

```
./stor target add { -n | --name } TARGET_NAME [ --max-sessions MAX_SESSIONS ] [ --reclaim-policy RECLAIM_POLICY ] [ { -c | --chap-name } CHAP_NAME { -p | --password } CHAP_PASSWORD { -s | --status } STATUS ] [--num SERVER_NUMBER] [ --server SERVER_ID &<1-n> ] [ [ --initiator [ IP<1-n>][:NAME<1-n> ] [ IP<1-n>][:NAME<1-n> ] ] [ --target [ IP<1-n>][:NIC<1-n> ] [ IP<1-n>][:NIC<1-n> ] ] | --allow-file ALLOW_FILE ]
```

This command is used to create an iSCSI target.

Note: HBlock supports a maximum of 32766 target IQNs. A target can be associated with up to 256 LUNs, but each LUN can only be associated with one target.

Parameters

Parameter	Description
-n TARGET_NAME or --name TARGET_NAME	Specifies the iSCSI target name. The value is a string of 1 to 16 case-sensitive characters. It can only contain lowercase letters, dots (.), digits, or hyphens (-). Only supports starting with a letter or a digit. Note: A target can be associated with up to 256 LUNs, but a LUN can only be associated with one target.
--max-sessions MAX_SESSIONS	The maximum number of sessions allowed to be established per IQN under the iSCSI target. The value is an integer that ranges from 0 to 1024, the default value is 1. 0 means the client cannot discover the target. Note: If multiple clients are connected to the same target IQN, they can read at the same time, but cannot write at the same time.
-c CHAP_NAME or	Specifies the CHAP authentication name.

--chap-name <i>CHAP_NAME</i>	The value is a string of 3 to 64 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), underscores (_), or colons (:). Only supports starting with a letter or a digit.
--reclaim-policy <i>RECLAIM_POLICY</i>	Specifies the reclaim policy of the iSCSI target. Value: <ul style="list-style-type: none"> ● Delete: The iSCSI target is automatically deleted when all associated LUNs are deleted. ● Retain: The iSCSI target remains retained when all associated LUNs are deleted. The default value is Retain. Note: If you specify a non-existent iSCSI target when creating a LUN, the reclaim policy for creating a new iSCSI target will default to Delete.
-p <i>CHAP_PASSWORD</i> or --password <i>CHAP_PASSWORD</i>	Specifies the client CHAP authentication password. The value is a string of 12 to 16 case-sensitive characters. It must contain at least two of the following: lowercase letters, uppercase letters, digits, or underscores (_).
-s <i>STATUS</i> or --status <i>STATUS</i>	Specifies the status of CHAP authentication. Value: <ul style="list-style-type: none"> ● Enabled (on): Enable CHAP authentication. ● Disabled (off): Disable CHAP authentication.
--num <i>SERVER_NUMBER</i>	Specifies the number of servers where target is located, only supported by cluster mode. The value is an integer that ranges from 2 to <i>n</i> , <i>n</i> is the number of servers in the cluster, the default value is 2.
--server <i>SERVER_ID</i> &<1- <i>n</i> >	Specifies the server ID where target is located, only supported by cluster mode. You can Specifies up to <i>n</i> server IDs, separated by commas. <i>n</i> is the number of servers in the cluster.
--initiator [<i>IP</i> &<1- <i>n</i> >] [<i>:NAME</i> &<1- <i>n</i> >]	Set the iSCSI initiator allowlist for every IQN under the target. You can define several allowlists, the relationship among them is logical OR . Within one allowlist you may simultaneously specify both IP address and initiator name; these two conditions are combined with logical AND . If no allowlist is configured, all accesses are permitted. Note: Mounted LUNs retain read-write access even after their initiators and targets are removed from the

	<p>allowlist; once the session is disconnected, initiators outside the allowlist are blocked from re-mounting.</p> <ul style="list-style-type: none"> ● <i>IP&<1-n></i>: Set the allowlist for iSCSI initiators based on their IP addresses. Supports IPv4, IPv6, and CIDR subnets; multiple entries can be set, separated by commas. Note: The IP address cannot be 'localhost' addresses. ● <i>NAME&<1-n></i>: Set the allowlist for iSCSI initiators based on their names. The value is a string of 1 to 223 case-sensitive characters. It can only contain letters, digits, dots (.), colons (:), or hyphens (-). Wildcards * and ? are supported. Multiple entries can be configured and must be separated by commas.
<p>--target [<i>IP&<1-n></i>] [<i>:NIC&<1-n></i>]</p>	<p>Set the target allowlist for every IQN under the target. You can define several allowlists, the relationship among them is logical OR. Within one allowlist you may simultaneously specify both IP address and NIC name; these two conditions are combined with logical AND.</p> <ul style="list-style-type: none"> ● <i>IP&<1-n></i>: Set the allowlist for targets based on their IP addresses. Supports IPv4, IPv6, and CIDR subnets; multiple entries can be set, separated by commas. Note: The IP address cannot be 'localhost' addresses. ● <i>NIC&<1-n></i>: Set the allowlist for targets based on their NICs. The value is a string of 1 to 100 characters. It can contain letters, digits, dots (.), hyphens (-), or underscores (_). Wildcards * and ? are supported. Multiple entries can be configured and must be separated by commas.
<p>--allow-file <i>ALLOW_FILE</i></p>	<p>Allowlist file containing the permitted initiator and target access lists. The file must be UTF-8-encoded JSON; see iSCSI Target Allowlist File for details. Note: If an allowlist file is provided, the parameters -initiator and --target will be ignored.</p>

Examples

- Standalone Mode: Create targetc, allow up to three sessions to be established, and enable CHAP authentication.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor target add -n targetc --max-  
sessions 3 -c chap-test -p ***** -s on  
Created target targetc successfully.  
iqn = iqn.2012-08.cn.ctyunapi.oos:targetc.2(192.168.0.32:3260)
```

- Cluster Mode: Create target02, allow up to six sessions to be established, enable CHAP authentication, and specify the server IDs as hblock_1, hblock_2, and hblock_3.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor target add -n target02 --max-  
sessions 6 -c chap-test -p ***** -s on --num 3 --server hblock_1,hblock_2,hblock_3  
Created target target02 successfully.  
iqn = iqn.2012-08.cn.ctyunapi.oos:target02.5(192.168.0.110:3260)  
iqn = iqn.2012-08.cn.ctyunapi.oos:target02.6(192.168.0.192:3260)  
iqn = iqn.2012-08.cn.ctyunapi.oos:target02.7(192.168.0.102:3260)
```

- Cluster Mode: Create target05 and set up its allowlist.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor target add -n target05 --  
initiator 192.168.0.70:iqn.1991-05.com.microsoft:songt-0001 192.168.0.64  
Created target target05 successfully.  
iqn = iqn.2012-08.cn.ctyunapi.oos:target05.7 (192.168.0.67:3260)  
iqn = iqn.2012-08.cn.ctyunapi.oos:target05.8 (192.168.0.65:3260)
```

4.10.2 Delete an iSCSI Target

```
./stor target rm { -n | --name } TARGET_NAME [ { -c | --connection }
ISCSI_INITIATOR_NAME ] [ { -i | --target-ip } TARGET_IP ]
```

This command is used to delete the specified iSCSI target or disconnect the specified iSCSI target.

Note: Only iSCSI targets that are not associated with any LUNs can be deleted.

Parameters

Parameter	Description
<code>-n TARGET_NAME</code> or <code>--name TARGET_NAME</code>	Specifies the iSCSI target name.
<code>-c ISCSI_INITIATOR_NAME</code> or <code>--connection ISCSI_INITIATOR_NAME</code>	Specifies the iSCSI initiator name to be deleted. The value is a string. <ul style="list-style-type: none"> ● If the client is Windows operating system, <code>ISCSI_INITIATOR_NAME</code> is the "Initiator Name" in Configuration of iSCSI Initiator Properties. ● If the client is Linux operating system, <code>ISCSI_INITIATOR_NAME</code> is the "InitiatorName" acquired by running command <code>cat /etc/iscsi/initiatorname.iscsi</code>.
<code>-i TARGET_IP</code> or <code>--target-ip TARGET_IP</code>	Specifies target IP.

Examples

- Delete the target named targettest.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target rm -n targettest
Removed target targettest successfully.
```

- Delete target's initiator connection. And you need to disconnect from the iSCSI target on the client. For details, see **Client Operations**.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target rm -n target01 -c
iqn.1994-05.com.redhat:265a95d81ae
Removed the connection 'iqn.1994-05.com.redhat:265a95d81ae' with target target01
successfully.
```

4.10.3 Set CHAP Authentication of the iSCSI Target

```
./stor target set { -i | --item } chap { -n | --name } TARGET_NAME [ { -c | --chap-name }
CHAP_NAME ] [ { -p | --password } CHAP_PASSWORD ] { -s | --status } STATUS
```

This command is used to set CHAP authentication of iSCSI target.

Parameters

Parameter	Description
-i chap or --item chap	Sets CHAP authentication.
-n TARGET_NAME or --name TARGET_NAME	Specifies the iSCSI target name.
-c CHAP_NAME or --chap-name CHAP_NAME	Specifies the CHAP authentication name. The value is a string of 3 to 64 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), underscores (_), or colons (:). Only supports starting with a letter or a digit.
-p CHAP_PASSWORD or --password CHAP_PASSWORD	Specifies the client CHAP authentication password. The value is a string of 12 to 16 case-sensitive characters. It must contain at least two of the following: lowercase letters, uppercase letters, digits, or underscores (_).
-s STATUS or --status STATUS	Specifies the status of CHAP authentication: <ul style="list-style-type: none"> ● Enabled (on): Enable CHAP authentication. ● Disabled (off): Disable CHAP authentication.

Examples

Set the CHAP authentication name and authentication password of target02.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target set -i chap -n target02 -c
chap3 -p ***** -s on
Set target target02 successfully.
```

4.10.4 Migrate the iSCSI Target (Cluster Mode)

```
./stor target set { -i | --item } server { -n | --name } TARGET_NAME { -m | --migrate }
{ SOURCE_SERVER_ID:DEST_SERVER_ID }&<1-n> { -f | --force }
```

The command is used to migrate iSCSI target from one sever to the other server.

You can migrate only one server which the target located, or you can migrate multiple servers which the target located at the same time.

Note:

- Before migrating iSCSI target, ensure that the cluster is in working status and the destination server is properly connected.
- Currently, you can only forcibly migrate iSCSI target, which may cause data loss.
- If iSCSI target to migrate has been connected to a LUN and the LUN has been mounted to a client, disconnect the client from the original iSCSI target IQN before migrating iSCSI target. After the migration, ensure that the original iSCSI target IQN cannot be discovered, and reconnect the client to the migrated iSCSI target IQN.
- After the migration is complete, please check and adjust the target's allowlist to ensure it meets the access-control requirements.

Parameters

Parameter	Description
-i server or --item server	Sets iSCSI target migration.
-n TARGET_NAME or --name TARGET_NAME	Specifies the iSCSI target name.
-m { SOURCE_SERVER_ID:DEST_SERVER_ID }&<1-n> or --migrate { SOURCE_SERVER_ID:DEST_SERVER_ID }&<1-n>	Specifies the source server ID and destination server ID. Note: Specify multiple sets of SOURCE_SERVER_ID:DEST_SERVER_ID, separated by a comma (,).
-f or --force	Forcibly migrate iSCSI target. Note: You can only forcibly migrate iSCSI target, which may cause data loss.

Examples

Migrate a server for target03.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target set -i server -n target03 -m
hblock_2:hblock_3 --force
```


4.10.5 Edit Maximum Session Number per IQN Under the iSCSI Target

```
./stor target set { -i | --item } session { -n | --name } TARGET_NAME --max-sessions
MAX_SESSIONS
```

This command is used to edit the maximum number of sessions allowed to be established for each IQN under the iSCSI target.

Parameters

Parameter	Description
<code>-i session</code> or <code>--item session</code>	Indicates setting the maximum number of sessions allowed to be established for each IQN under the iSCSI target.
<code>-n TARGET_NAME</code> or <code>--name TARGET_NAME</code>	Specifies iSCSI target name.
<code>--max-sessions MAX_SESSIONS</code>	Specifies the maximum number of sessions allowed to be established per IQN under the iSCSI target. The value is an integer that ranges from 0 to 1024, the default value is 1. 0 means the client cannot discover the target. It is recommended to change it to a number larger than the number of existing sessions, because if it is changed to a number smaller than the number of existing sessions, it may cause the existing client connection to be disconnected and the connection cannot be established again. Note: If multiple clients are connected to the same target IQN, they can read at the same time, but cannot write at the same time.

Examples

Edit the maximum number of sessions allowed for each IQN under target01 to 10.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target set -i session -n target01
--max-sessions 10
Set target target01 successfully.
```

4.10.6 Modify the Remain Policy of the iSCSI Target

```
./stor target set { -i | --item } generic { -n | --name } TARGET_NAME --reclaim-policy
RECLAIM_POLICY
```

This command is used to modify the remain policy of the iSCSI target.

Parameters

Parameter	Description
-i generic or --item generic	Indicates setting the reclaim policy of the iSCSI target.
-n TARGET_NAME or --name TARGET_NAME	iSCSI target name.
--reclaim-policy RECLAIM_POLICY	Specifies the reclaim policy of the iSCSI target. Value: <ul style="list-style-type: none"> ● Delete: The iSCSI target is automatically deleted when all associated LUNs are deleted. ● Retain: The iSCSI target remains retained when all associated LUNs are deleted.

Examples

Modify the remain policy of **target05** to **Retain**.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target set -i generic -n target05
--reclaim-policy Retain
Set target target05 successfully.
```

4.10.7 Set iSCSI Target AllowList

```
./stor target { S | setallow } { -n | --name } TARGET_NAME {-a | --action } ACTION [ { -q | --iqn } IQN_NAME<1-n> ] [ [ --initiator [ IP<1-n>[:NAME<1-n> ] [ IP<1-n>[:NAME<1-n> ] ] [ --target [ IP<1-n>[:NIC<1-n> ] [ IP<1-n>[:NIC<1-n> ] ] | --allow-file ALLOW_FILE ]
```

This command is used to set iSCSI target allowlist.

Parameters

Parameter	Description
<code>-n TARGET_NAME</code> or <code>--name TARGET_NAME</code>	iSCSI target name.
<code>-a ACTION</code> or <code>--action ACTION</code>	Set access permissions for the iSCSI target. Value: <ul style="list-style-type: none"> ● add: Add to the allowlist. ● delete: Delete the allowlist. ● replace: Replace the existing allowlist.
<code>-q IQN_NAME<1-n></code> or <code>--iqn IQN_NAME<1-n></code>	Specifies the target IQN name. Multiple target IQN can be set, separated by commas. If an IQN is specified, the access rule applies only to that IQN; if omitted, the rule is applied to every IQN under the target.
<code>--initiator [IP<1-n>[:NAME<1-n>]</code>	Set the iSCSI initiator allowlist for every IQN under the target. You can define several allowlists, the relationship among them is logical OR . Within one allowlist you may simultaneously specify both IP address and initiator name; these two conditions are combined with logical AND . If no allowlist is configured, all accesses are permitted. Note: Mounted LUNs retain read-write access even after their initiators and targets are removed from the allowlist; once the session is disconnected, initiators outside the allowlist are blocked from re-mounting. <ul style="list-style-type: none"> ● <code>IP<1-n></code>: Set the allowlist for iSCSI initiators based on their IP addresses. Supports IPv4, IPv6, and CIDR subnets; multiple entries can be set, separated by commas. Note: The IP address cannot be 'localhost' addresses.

	<ul style="list-style-type: none"> ● <i>NAME</i>&<1-<i>n</i>>: Set the allowlist for iSCSI initiators based on their names. The value is a string of 1 to 223 case-sensitive characters. It can only contain letters, digits, dots (.), colons (:), or hyphens (-). Wildcards * and ? are supported. Multiple entries can be configured and must be separated by commas. <p>Note: At least one entry must be provided for either <i>initiator</i> or <i>target</i>.</p>
<p>--target [<i>IP</i>&<1-<i>n</i>>] [:<i>NIC</i>&<1-<i>n</i>>]</p>	<p>Set the target allowlist for every IQN under the target. You can define several allowlists, the relationship among them is logical OR. Within one allowlist you may simultaneously specify both IP address and NIC name; these two conditions are combined with logical AND.</p> <ul style="list-style-type: none"> ● <i>IP</i>&<1-<i>n</i>>: Set the allowlist for targets based on their IP addresses. Supports IPv4, IPv6, and CIDR subnets; multiple entries can be set, separated by commas. <p>Note: The IP address cannot be 'localhost' addresses.</p> <ul style="list-style-type: none"> ● <i>NIC</i>&<1-<i>n</i>>: Set the allowlist for targets based on their NICs. The value is a string of 1 to 100 characters. It can contain letters, digits, dots (.), hyphens (-), or underscores (_). Wildcards * and ? are supported. Multiple entries can be configured and must be separated by commas. <p>Note: At least one entry must be provided for either <i>initiator</i> or <i>target</i>.</p>
<p>--allow-file <i>ALLOW_FILE</i></p>	<p>Allowlist file containing the permitted initiator and target access lists. The file must be UTF-8-encoded JSON; see iSCSI Target Allowlist File for details.</p> <p>Note: If an allowlist file is provided, the parameters -initiator and --target will be ignored.</p>

Examples

- Add an allowlist for target05.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target S -n target5 -a add --
initiator 192.168.0.66 192.168.0.65
Set target target5 allowlist successfully.
```

- Add an allowlist for iqn.2012-08.cn.ctyunapi.oos:target5.13 under target5.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target S -n target5 -q iqn.2012-08.cn.ctyunapi.oos:target5.13 -a add --initiator 192.168.0.67
Set target iqn iqn.2012-08.cn.ctyunapi.oos:target5.13 allowlist successfully.
```

4.10.8 Delete an iSCSI Target Allowlist

```
./stor target { R | rallow } { -n | --name } TARGET_NAME [ { -q | --iqn } IQN_NAME&<1-n> ]
```

This command is used to delete the iSCSI target allowlist.

Parameters

Parameter	Description
<code>-n TARGET_NAME</code> or <code>--name TARGET_NAME</code>	iSCSI target name.
<code>-q IQN_NAME&<1-n></code> or <code>--iqn IQN_NAME&<1-n></code>	Specifies the target IQN name. Multiple target IQN can be set, separated by commas. If a target IQN is specified, the allowlist for that IQN is deleted; if no IQN is specified, the allowlists for all IQNs under the target are deleted.

Examples

- Delete the allowlist for `iqn.2012-08.cn.ctyunapi.oos:target5.13` under `target05`.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target R -n target5 -q iqn.2012-08.cn.ctyunapi.oos:target5.13
Removed target iqn iqn.2012-08.cn.ctyunapi.oos:target5.13 allowlist successfully.
```

- Delete all allowlists under `target05`.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target R -n target5
Removed target target5 allowlist successfully.
```

4.10.9 Query Target Information

```
./stor target ls [ -c | --connection ] [ { -n | --name } TARGET_NAME ]
```

This command is used to query iSCSI target information.

Parameters

Parameter	Description
<code>-c</code> or <code>--connection</code>	Query iSCSI target connections.
<code>-n TARGET_NAME</code> or <code>--name TARGET_NAME</code>	Specifies the iSCSI target to be queried.

Examples

- Query all iSCSI target information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls
```

No.	Target Name	Max Sessions	iSCSI Target	CHAP
1.	target01	1	iqn.2012-08.cn.ctyunapi.oos:target01.3(192.168.0.117:3260) iqn.2012-08.cn.ctyunapi.oos:target01.4(192.168.0.102:3260)	test,Disabled
2.	target02	1	iqn.2012-08.cn.ctyunapi.oos:target02.5(192.168.0.192:3260) iqn.2012-08.cn.ctyunapi.oos:target02.6(192.168.0.117:3260)	Disabled
3.	target03	1	iqn.2012-08.cn.ctyunapi.oos:target03.7(192.168.0.192:3260) iqn.2012-08.cn.ctyunapi.oos:target03.8(192.168.0.102:3260)	Disabled
4.	target05(Deleting)	1	iqn.2012-08.cn.ctyunapi.oos:target05.9(192.168.0.117:3260) iqn.2012-08.cn.ctyunapi.oos:target05.10(192.168.0.192:3260)	Disabled
5.	tg001	2	iqn.2012-08.cn.ctyunapi.oos:tg001.1(192.168.0.117:3260) iqn.2012-08.cn.ctyunapi.oos:tg001.2(192.168.0.192:3260)	Disabled

Table 1: Description of the queried iSCSI target information

Item	Description
No.	Number.
Target Name	iSCSI target name. Note: If the iSCSI target is in the process of being deleted, "Deleting" will be displayed after the iSCSI target name.
Max Sessions	The maximum number of sessions allowed to be established per IQN under the iSCSI target.
iSCSI Target	Target IQN, IP, and port number.
CHAP	CHAP authentication information, including CHAP name and CHAP status. The CHAP name will be displayed only if CHAP authentication information is configured. CHAP status: <ul style="list-style-type: none"> ● Enabled: Enable CHAP authentication. ● Disabled: CHAP authentication is not enabled.

- Standalone Mode: Query information about iSCSI target named targeta.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -n targeta
Target Name: targeta
Max Sessions: 3
Create Time: 2025-07-22 16:18:57
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260)
LUN: luna1 (LUN 0),luna2 (LUN 1),luna3 (LUN 2)
Reclaim Policy: Delete
CHAP: chanptest,Hblockhblock,Disabled
Allowlist for iqn.2012-08.cn.ctyunapi.oos:targeta.1 (192.168.0.66:3260):
  Initiator allowlist:
  +-----+-----+-----+
  | No. | IPs           | Names |
  +-----+-----+-----+
  | 1.  | 192.168.0.68 |       |
  | 2.  | 192.168.0.70 |       |
  +-----+-----+-----+
  Target allowlist:
  +-----+-----+-----+
  | No. | IPs           | NICs |
  +-----+-----+-----+
  | 1.  | 192.168.0.66 |       |
  +-----+-----+-----+
```

- Cluster Mode: Query information about the iSCSI target named target03.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -n target03
Target Name: target03
Max Sessions: 1
Create Time: 2025-07-10 09:42:00
Number of Servers: 2
iSCSI Target: iqn.2012-08.cn.ctyunapi.oos:target03.5 (192.168.0.64:3260)
              iqn.2012-08.cn.ctyunapi.oos:target03.6 (192.168.0.65:3260)
LUN: lun03a (LUN 0),lun03b (LUN 1)
Reclaim Policy: Delete
CHAP: chap-test,Hblockhblock,Disabled
Server ID: hblock_1,hblock_2
Allowlist for iqn.2012-08.cn.ctyunapi.oos:target03.5 (192.168.0.64:3260):
  Initiator allowlist:
  +-----+-----+-----+
  | No. | IPs           | Names |
  +-----+-----+-----+
  | 1.  | 192.168.0.70 | iqn.1991-05.com.microsoft:songt-0001 |
  | 2.  | 192.168.0.66 |       |
  +-----+-----+-----+
  Target allowlist:
  +-----+-----+-----+
  | No. | IPs           | NICs |
  +-----+-----+-----+
  | 1.  | 192.168.0.64 |       |
  +-----+-----+-----+
```

```

| 2. | 192.168.0.65 | |
| 3. | 192.168.0.67 | |
+-----+-----+-----+
Allowlist for iqn.2012-08.cn.ctyunapi.oos:target03.6 (192.168.0.65:3260):
Initiator allowlist:
+-----+-----+-----+
| No. | IPs           | Names                               |
+-----+-----+-----+
| 1.  | 192.168.0.70 | iqn.1991-05.com.microsoft:songt-0001 |
| 2.  | 192.168.0.66 |                                         |
+-----+-----+-----+
Target allowlist:
+-----+-----+-----+
| No. | IPs           | NICs |
+-----+-----+-----+
| 1.  | 192.168.0.64 | |
| 2.  | 192.168.0.65 | |
| 3.  | 192.168.0.67 | |
+-----+-----+-----+

```

Table 2: Description of the specified iSCSI target information

Item	Description
Target Name	iSCSI target name.
Status	The status of the iSCSI target: <ul style="list-style-type: none"> ● Deleting: The iSCSI target is being deleted. This item is returned only when the iSCSI target is in the process of being deleted.
Max Sessions	The maximum number of sessions allowed to be established per IQN under the iSCSI target.
Create Time	Target creation time.
Number of Servers	The number of servers where target is located (only supported by cluster mode).
iSCSI Target	Target IQN, Client IP, and port number.
LUN	The LUN corresponding to target. Content within parentheses indicates the LUN number.
Reclaim Policy	The reclaim policy of the iSCSI target: <ul style="list-style-type: none"> ● Delete: The iSCSI target is automatically deleted when all associated LUNs are deleted. ● Retain: The iSCSI target remains retained when all associated LUNs are deleted.
CHAP	CHAP authentication information, including CHAP name, CHAP password, and CHAP status. <p>Note: The CHAP name will be displayed only if CHAP authentication information is configured.</p> CHAP status: <ul style="list-style-type: none"> ● Enabled: Enable CHAP authentication.

	<ul style="list-style-type: none"> ● Disabled: CHAP authentication is not enabled.
ServerID	The server ID corresponding to the iSCSI target (only supported by cluster mode).
Initiator allowlist	The iSCSI initiator allowlist. <ul style="list-style-type: none"> ● No.: Number. ● IPs: The allowlist for iSCSI initiators based on their IP addresses. ● Names: The allowlist for iSCSI initiators based on their names.
Target allowlist	The target allowlist. <ul style="list-style-type: none"> ● No.: Number. ● IPs: The allowlist for targets based on their IP addresses. ● Names: The allowlist for targets based on their NICs.

- Query iSCSI target connections information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -c
```

SessionId	Target IQN	Target IP	Initiator name	Client IP	Client Port
0x8	iqn.2012-08.cn.ctyunapi.oos:target01.3	192.168.0.102	iqn.1991-05.com.microsoft:ecs-e16f-0915299	192.168.0.116	61314
0x3	iqn.2012-08.cn.ctyunapi.oos:target01.4	192.168.0.192	iqn.1991-05.com.microsoft:ecs-e16f-0915299	192.168.0.116	61416
0x2	iqn.2012-08.cn.ctyunapi.oos:target03.5	192.168.0.192	iqn.1994-05.com.redhat:ca375039f35f	127.0.0.1	58978
0x1	iqn.2012-08.cn.ctyunapi.oos:target1.1	192.168.0.192	iqn.1994-05.com.redhat:ca375039f35f	127.0.0.1	58976

Table 3: Description of the iSCSI target connection

Item	Description
SessionId	iSCSI session ID.
Target IQN	iSCSI target IQN.
Target IP	iSCSI target IP.
Initiator name	The name of the initiator to which the connection belongs.
Client IP	Client IP.
Client Port	Client port number.

- Query the iSCSI target connection information of target02.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor target ls -c -n target02
```

SessionId	Target IQN	Target IP	Initiator name	Client IP	Client Port
0x5	iqn.2012-08.cn.ctyunapi.oos:target02.7	192.168.0.121	iqn.1991-05.com.microsoft:ecs-28f3	192.168.0.46	52594
0x7	iqn.2012-08.cn.ctyunapi.oos:target02.8	192.168.0.72	iqn.1991-05.com.microsoft:ecs-28f3	192.168.0.46	52795

Table 4: Description of the iSCSI target connection

Item	Description
SessionId	iSCSI session ID.
Target IQN	iSCSI target IQN.
Target IP	iSCSI target IP.
Initiator name	The name of the initiator to which the connection belongs.
Client IP	Client IP.
Client Port	Client port number.

4.11 Storage Pool Operations (Cluster Mode)

4.11.1 Create a Storage Pool

```
./stor storagepool add { -n | --name } POOL_NAME [ --fault-domain FAULT_DOMAIN ] [ [ --qos-name QOS_NAME ] [ --iops IOPS ] [ --read-iops READ_IOPS ] [ --write-iops WRITE_IOPS ] [ --bps BPS ] [ --read-bps READ_BPS ] [ --write-bps WRITE_BPS ] [ --iops-burst IOPS_BURST ] [ --read-iops-burst READ_IOPS_BURST ] [ --write-iops-burst WRITE_IOPS_BURST ] [ --bps-burst BPS_BURST ] [ --read-bps-burst READ_BPS_BURST ] [ --write-bps-burst WRITE_BPS_BURST ] [ --iops-burst-secs IOPS_BURST_SECS ] [ --read-iops-burst-secs READ_IOPS_BURST_SECS ] [ --write-iops-burst-secs WRITE_IOPS_BURST_SECS ] [ --bps-burst-secs BPS_BURST_SECS ] [ --read-bps-burst-secs READ_BPS_BURST_SECS ] [ --write-bps-burst-secs WRITE_BPS_BURST_SECS ] ] ]
```

This command is used to create a storage pool.

Note: A maximum of 32768 storage pools can be created in an HBlock cluster.

Parameter

Parameter	Description
<code>-n POOL_NAME</code> or <code>--name POOL_NAME</code>	Specifies the name of a storage pool. The value is a string of 1 to 16 case-sensitive characters. It can contain letters, digits, hyphens (-), and underscores (_). Only supports starting with a letter or a digit.
<code>--fault-domain FAULT_DOMAIN</code>	Specifies the fault domain level of the storage pool. Value: <ul style="list-style-type: none"> ● path: disk path level. ● server: server level. ● rack: rack level. ● room: room level. The default value is server.
<code>--qos-name QOS_NAME</code>	Specifies the name of a QoS policy. Note: If the specified QoS policy name does not exist when creating a storage pool, the QoS policy will be created at the same time, the reclaim policy of the QoS policy is Delete. Value: The value is a string of 1 to 64 case-sensitive characters. It can contain letters, digits, and hyphens (-). Only supports starting with a letter or a digit.

	<p>Note:</p> <ul style="list-style-type: none"> ● If QoS policy name is not specified but any IOPS/Bps parameter is set, a new QoS policy is automatically created based on the specified IOPS/Bps. The system assigns the policy name in the format: <code>pool-poolname-qos-timestamp</code>. ● If QoS policy name is specified and any IOPS/Bps parameter is set, QoS policy name must be unique and cannot match any existing QoS policy. ● If QoS policy name is specified and no IOPS/Bps parameter is set, QoS policy name must be the name of an existing QoS policy.
<code>--iops IOPS</code>	Specifies the limit of I/O operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
<code>--read-iops READ_IOPS</code>	Specifies the limit of read operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
<code>--write-iops WRITE_IOPS</code>	Specifies the limit of write operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
<code>--bps BPS</code>	<p>Specifies the limit of I/O throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
<code>--read-bps READ_BPS</code>	Specifies the limit of read throughput per second.

	<p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
<p>--write-bps <i>WRITE_BPS</i></p>	<p>Specifies the limit of write throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
<p>--iops-burst <i>IOPS_BURST</i></p>	<p>Specifies the burst limit of I/O operations per second.</p> <p>This setting only takes effect when --iops <i>IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>IOPS</i>, 999,999,999]. The default value is -1, indicating no limit.</p>
<p>--read-iops-burst <i>READ_IOPS_BURST</i></p>	<p>Specifies the burst limit of read operations per second.</p> <p>This setting only takes effect when --read-iops <i>READ_IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range</p>

	<p>(<i>READ_IOPS</i> 999,999,999]. The default value is -1, indicating no limit.</p>
<p>--write-iops-burst <i>WRITE_IOPS_BURST</i></p>	<p>Specifies the burst limit of write operations per second.</p> <p>This setting only takes effect when --write-iops <i>WRITE_IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>WRITE_IOPS</i>, 999,999,999]. The default value is -1, indicating no limit.</p>
<p>--bps-burst <i>BPS_BURST</i></p>	<p>Specifies the burst limit of I/O throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --bps <i>BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3].
<p>--read-bps-burst <i>READ_BPS_BURST</i></p>	<p>Specifies the burst limit of read throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --read-bps <i>READ_BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p>

	<ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3].
<p>--write-bps-burst <i>WRITE_BPS_BURST</i></p>	<p>Specifies the burst limit of write throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --write-bps <i>WRITE_BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i> 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3].
<p>--iops-burst-secs <i>IOPS_BURST_SECS</i></p>	<p>Specifies the duration of I/O operations with burst limit.</p>

	<p>Note: This setting is only effective when the IOPS Burst feature is enabled.</p> <p>Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.</p>
<p>--read-iops-burst-secs <i>READ_IOPS_BURST_SECS</i></p>	<p>Specifies the duration of read operations with burst limit.</p> <p>Note: This setting is only effective when the read IOPS Burst feature is enabled.</p> <p>Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.</p>
<p>--write-iops-burst-secs <i>WRITE_IOPS_BURST_SECS</i></p>	<p>Specifies the duration of write operations with burst limit.</p> <p>Note: This setting is only effective when the write IOPS Burst feature is enabled.</p> <p>Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.</p>
<p>--bps-burst-secs <i>BPS_BURST_SECS</i></p>	<p>Specifies the duration of I/O throughput with burst limit.</p> <p>Note: This setting is only effective when the Bps Burst feature is enabled.</p> <p>Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.</p>
<p>--read-bps-burst-secs <i>READ_BPS_BURST_SECS</i></p>	<p>Specifies the duration of read throughput with burst limit.</p> <p>Note: This setting is only effective when the read Bps Burst feature is enabled.</p> <p>Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.</p>
<p>--write-bps-burst-secs <i>WRITE_BPS_BURST_SECS</i></p>	<p>Specifies the duration of write throughput with burst limit.</p> <p>Note: This setting is only effective when the write Bps Burst feature is enabled.</p> <p>Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.</p>

Example

Create storage pool pool4 and add its description.

Note: If you do not want to add the description, press Enter.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor storagepool add -n pool4 --fault-domain path
Enter the description for this pool, limited to 50 characters:
pool4's fault domain is path level.
Created storage pool pool4 successfully.
```

4.11.2 Add Nodes to a Storage Pool

```
./stor storagepool { A | addnode } { -n | --name } POOL_NAME --node NODE <1-n>
```

This command is used to add nodes to a storage pool.

All disk path nodes added during initialization belong to nodes of the base storage pool. If it is necessary to add nodes from the base storage pool to other storage pool, they need to be removed from the base storage pool first.

Note:

- The nodes to add must belong to the cluster topology, and each cluster topology node can be added to only one storage pool.
- When a node at the path level is added, the disk path cannot belong to another storage pool or in removed status. The server to which the disk path belongs cannot be in removed status.
- If the type of the node to add is lower than the level of the fault domain of the storage pool, the ancestor nodes of the node must have a mapping node of the same level as the fault domain of the storage pool. Otherwise, an error message is reported. For example, if the fault domain of the storage pool is rack, for node root:room1:server1, an error occurs when server1 is added to the storage pool.
- If the type of the node to add is higher than the level of the fault domain of the storage pool, the child nodes of the node must have a mapping node of the same level as the fault domain of the storage pool. Otherwise, an error message is reported. For example, if the fault domain of a storage pool is rack and topology node root:room1:server1:path1 exists, an error occurs when room1 is added to the storage pool.

Parameter

Parameter	Description
<code>-n POOL_NAME</code> or <code>--name POOL_NAME</code>	Specifies the name of a storage pool.
<code>--node NODE <1-n></code>	Specifies the node to be added to the storage pool. The node must belong to the cluster topology. Multiple nodes can be added to the storage pool at a time. Separate them by commas (.). Note: <ul style="list-style-type: none"> ● If there are child nodes under the node to be added, the child nodes that have not been allocated to any storage pool will be uniformly added to the storage pool.

	<ul style="list-style-type: none">● The node name can be used in the full path of the cluster topology, in the format of: <i>name:name:name</i>, specified step by step from the root node; Partial paths can also be used, but they must be unique within the cluster topology. For example, if <code>default:room4:hblock-4</code>, <code>room4:hblock-4</code>, <code>hblock-4</code> all point to the same node in the cluster topology, and the node name is unique in the cluster topology, then one of the node names can be selected.
--	---

Examples

Add node `default:hblock_4:/mnt/storage01` (default is the name of the root node) to `pool4`.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor storagepool addnode -n pool4 --node
default:hblock_4:/mnt/storage01
Added node default:hblock_4:/mnt/storage01 to storage pool pool4 successfully.
```

4.11.3 Modify a Storage Pool

```
./stor storagepool set { -n | --name } POOL_NAME [ --new-name NEW_NAME ] [ --change-  
description ]
```

This command is used to modify the information about a storage pool.

Parameter

Parameter	Description
<code>-n POOL_NAME</code> or <code>--name POOL_NAME</code>	Specifies the name of a storage pool.
<code>--new-name NEW_NAME</code>	Specifies a new name for the storage pool. The value is a string of 1 to 16 case-sensitive characters. It can contain letters, digits, hyphens (-), and underscores (_). Only supports starting with a letter or a digit.
<code>--change-description</code>	Changes the description of the storage pool.

● Examples

Modify the name of storage pool from pool3 to rackpool3.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor storagepool set -n pool3 --new-name  
rackpool3  
Set storage pool pool3 successfully.
```

Modify the description of pool4.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor storagepool set -n pool4 --change-  
description  
Enter the description for this pool, limited to 50 characters:  
pool4's fault domain is path level.  
Set storage pool pool4 successfully.
```

4.11.4 Remove a Node from a Storage Pool

```
./stor storagepool { R | rmnode } { -n | --name } POOL_NAME --node NODE &<1-n> [ -d | --del-data ] [ -f | --force ]
```

This command is used to remove a node from a storage pool.

Note:

- If only one fault domain is available in the base storage pool, you cannot remove any node from the fault domain.
- A node involving multiple fault domains of the base storage pool cannot be removed.

Parameters

Parameter	Description
<code>-n POOL_NAME</code> or <code>--name POOL_NAME</code>	Specifies the name of the storage pool from which you want to remove a node.
<code>--node NODE &<1-n></code>	Specifies a node name to be removed from the storage pool. Note: You can remove multiple nodes in the same fault domain at a time. Nodes are separated by commas (.). Nodes in multiple fault domains cannot be removed at a time.
<code>-d</code> or <code>--del-data</code>	Removes a node from the storage pool and delete the HBlock data stored on the node.
<code>-f</code> or <code>--force</code>	Forcibly removes a node from the storage pool. Note: Forcibly removing a node from the storage pool may cause data loss.

Example

Remove root:room1:rack1:hblock_2:/mnt/storage01 from rackpool3.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor storagepool rmnode -n rackpool3 --node root:room1:rack1:hblock_2:/mnt/storage01
Start removing node root:room1:rack1:hblock_2:/mnt/storage01 from storage pool rackpool3.
```

4.11.5 Delete a Non-base Storage Pool

```
./stor storagepool rm { -n | --name } POOL_NAME
```

This command is used to delete a non-base storage pool.

Note:

- You cannot delete a base storage pool.
- A storage pool containing LUNs cannot be deleted.
- After a storage pool is deleted, the topology nodes in the storage pool do not belong to any storage pool, but remain in the cluster topology.

Parameter

Parameter	Description
<code>-n POOL_NAME</code> or <code>--name POOL_NAME</code>	Specifies the name of the storage pool to delete.

Example

Delete storage pool pool4.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor storagepool rm -n pool4  
Start deleting storage pool pool4.
```

4.11.6 Query Storage Pool Information

```
./stor storagepool ls [ { -n | --name } POOL_NAME [ --qos ] ]
```

This command is used to query storage pool information.

Parameters

Parameter	Description
-n <i>POOL_NAME</i> or --name <i>POOL_NAME</i>	Specifies the storage pool to be queried.
--qos	Specifies to query the QoS policy associated with the storage pool.

Examples

- Query information about all storage pools.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor storagepool ls
```

No.	Pool Name	Pool Status	Fault Domain	Total Capacity	Used Capacity
1.	default (*)	Normal	path	278.57 GiB	23.42 GiB
2.	pool2	Normal	server	92.86 GiB	2.15 GiB

Table1: Descriptions of information for all queried storage pools

Item	Description
No.	Number.
Pool Name	The name of the storage pool. *: Indicates a base storage pool.
Pool Status	Status of the storage pool: <ul style="list-style-type: none"> ● Normal. ● Deleting.
Fault Domain	Fault domain level: <ul style="list-style-type: none"> ● path. ● server ● rack. ● room.
Total Capacity	Total capacity of the storage pool.
Used Capacity	Used capacity of the storage pool.

- Query the information of a specified storage pool.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor storagepool ls -n pool2
Storage Pool Name: default
Pool Status: Normal
Fault Domain: path
Base Pool: true
Total Capacity: 278.57 GiB
Used Capacity: 23.98 GiB
Created Time: 2026-01-12 17:34:40
LUN:
  LUN for CachePool: lun05a
  LUN for Pool: lun02,lun02a,lun03a,lun01a,lun01b,lun01a-C1
  status: 7 total, 1.33 TiB (cache mode:1 total, 210 GiB; storage mode:1 total, 310 GiB; local mode:5 total, 841 GiB)
Data Health: 100% normal, 0% low redundancy, 0% error
Topology:
Node Name: default:room1:rack1:hblock_2:/mnt/stor01
Node Type: path
Status: Healthy
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+
| No. | Path | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail |
+-----+-----+-----+-----+-----+-----+-----+
| 1. | default:room1:rack1:hblock_2:/mnt/stor01 | 8.4 GiB | 92.86 GiB | 1000.26 MiB | 90 GiB | Healthy | |
+-----+-----+-----+-----+-----+-----+-----+
Node Name: default:room1:hblock_3:/mnt/stor01
Node Type: path
Status: Healthy
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+
| No. | Path | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail |
+-----+-----+-----+-----+-----+-----+-----+
| 1. | default:room1:hblock_3:/mnt/stor01 | 7.82 GiB | 92.86 GiB | 984.27 MiB | 90 GiB | Healthy | |
+-----+-----+-----+-----+-----+-----+-----+
Node Name: default:hblock_1:/mnt/stor02
Node Type: path
Status: Healthy
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+
| No. | Path | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail |
+-----+-----+-----+-----+-----+-----+-----+
| 1. | default:hblock_1:/mnt/stor02 | 7.76 GiB | 92.86 GiB | 1004.29 MiB | 90 GiB | Healthy | |
+-----+-----+-----+-----+-----+-----+-----+
```

Table 2: Description of information for the specified storage pool being queried

Item	Description
Storage Pool Name	The name of the storage pool.
Pool Status	Status of the storage pool: <ul style="list-style-type: none"> ● Normal. ● Deleting.
Fault Domain	Fault domain level of the storage pool: <ul style="list-style-type: none"> ● path. ● server. ● rack. ● room.
Base Pool	Whether it is a base storage pool: <ul style="list-style-type: none"> ● true: It is a base storage pool. ● false: It is not a base storage pool.
Total Capacity	Total capacity of the storage pool.
Used Capacity	Used capacity of the storage pool.
Created Time	Time when the storage pool is created.
Description	Description of the storage pool.
LUN	Information about LUNs, including LUN for CachePool, LUN for Pool, and status.

	LUN for CachePool	List of LUNs using this storage pool as a cache storage pool.
	LUN for Pool	List of LUNs using this storage pool as the final storage pool.
	status	LUN status information, including the total number of LUNs, total capacity, and number and capacity of each node's LUN.
Data Health	The data health status of the storage pool, including: the percentage of normal data (normal) for all LUNs in this storage pool, the percentage of low redundancy data (low redundancy) for all LUNs in this storage pool, and the percentage of erroneous data (error) for all LUNs in this storage pool. If low redundancy data exists, low redundancy reconstruction progress will be provided.	
Topology	Topology of the storage pool, including Node Name, Node Type, Status, Disk Path(s), and Removing Details.	
	Node Name	Node name.
	Node Type	Node Type: <ul style="list-style-type: none"> ● path. ● server. ● rack. ● room.
	Status	Node health status: <ul style="list-style-type: none"> ● Healthy: A node is healthy and can be read and written normally. ● Warning: A node is in warning status and can be read. ● Error: A node is in error status and cannot be accessed.
	Disk Path(s)	Disk path information: <ul style="list-style-type: none"> ● No.: Number. ● Path: Disk path. ● Used Capacity: The used capacity of the disk where the disk path resides. ● Total Capacity: The total capacity of the disk where the disk path resides. ● Used Capacity Quota: The total amount of data that has been written to the disk path by HBlock. ● Capacity Quota: The capacity quota of the disk path, that is, for each disk path, the total amount of data that can be written by HBlock. Once the space used by HBlock reaches the quota, data writing will

		<p>be blocked immediately, and space exceeding the quota is not allowed to be used.</p> <ul style="list-style-type: none"> ● Health Status: Health status of the disk path. <ul style="list-style-type: none"> ■ Healthy: The disk path is in a healthy status and can be read and written normally, and the disk usage of the disk path does not exceed the threshold (the system default value is 95%). ■ Warning: The disk path is warning and is readable, but one or more of the following situations exist: slow disk, the utilization rate of the disk where the disk path resides exceeds the threshold (the system default value is 95%), the remaining disk space is less than 1GiB, HBlock stops writing to this path, the capacity quota usage of the disk path exceeds the threshold (the system default value is 95%), or the capacity quota of the disk path is set to 0. ■ Error: The disk path is in an error status and cannot be accessed. The reasons may be that an I/O error occurs on the disk, resulting in the inability to read or write, or the disk path is not mounted correctly, etc. ● Health Detail: <ul style="list-style-type: none"> ■ If the health status is Healthy, this column is empty. ■ If the health status is Warning or Error, detailed information about the warning or error is displayed.
	<p>Removing Details</p>	<p>Details of the removed node:</p> <ul style="list-style-type: none"> ● No.: Number. ● Path: Disk path of a node. ● Stage: The stage in which a node is removed: <ul style="list-style-type: none"> ■ Reconfiguration. ■ CheckingData. ■ Executing. ● Details: <ul style="list-style-type: none"> ■ FaultDomains: Details about the fault domains, including healthy (number of healthy fault domains), warning (number of warning fault domains), and error (number of error fault domains).

		<ul style="list-style-type: none"> ■ Data: safe (percentage of safe data), await reconstruction (percentage of data that needs to be reconstructed), await more faultdomains (percentage of data that requires additional fault domains to be reconstructed), single-copy (percentage of single-copy data), corrupted (corrupted data percentage).
--	--	---

● Query the QoS policy of storage pool defaultpool.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor storagepool ls -n defaultpool --qos
The specific QoS policy for storage pool defaultpool
QoS Policy Name: QoS6
Reclaim Policy: Retain
IOPS (Total/Read/Write): 1000/600/400
Bandwidth (Total/Read/Write): Unlimited
Burst IOPS (Total/Read/Write): 10000/6000/4000
Burst Bandwidth (Total/Read/Write): Unlimited
Burst Duration (second): 1/1/1/1/1/1
Create Time: 2025-07-31 16:28:51
Description: It is QoS6.

The default QoS policy for LUNs in the pool:
QoS Policy Name: QoS-Test
Reclaim Policy: Retain
IOPS (Total/Read/Write): 8000/Unlimited/Unlimited
Bandwidth (Total/Read/Write): 7.81 GiB/s | Unlimited | Unlimited
Burst IOPS (Total/Read/Write): Unlimited
Burst Bandwidth (Total/Read/Write): Unlimited
Burst Duration (second): 1/1/1/1/1/1
Create Time: 2025-08-06 17:02:32
```

Table3: Description of querying QoS policy information associated with the storage pool

Item	Description
The specific QoS policy for storage pool <i>poolName</i>	QoS policy of the storage pool.
The default QoS policy for LUNs in the pool	The default QoS policy for LUN in the pool.
QoS Policy Name	QoS policy name.
Reclaim Policy	QoS policy reclaim policy: <ul style="list-style-type: none"> ● Delete: When all objects associated with the QoS policy are disassociated or deleted, the system will automatically remove the QoS policy. ● Retain: When all objects associated with the QoS policy are disassociated or deleted, the QoS policy itself remains and is not deleted.
IOPS (Total/Read/Write)	The limit of total/read/ write operations per second.
Bandwidth (Total/Read/Write)	The limit of total/read/write throughput per second.

Burst IOPS (Total/Read/Write)	The burst limit of total/read/write operations per second.
Burst Bandwidth (Total/Read/Write)	The burst limit of total/read/write throughput per second.
Burst Duration (second)	The duration in seconds of I/O operations with burst limit. The corresponding sequence of QoS policy parameters is: burst IOPS (total/read/write), burst bandwidth (total/read/write).
Create Time	The creation time of the QoS policy.
Description	The description of the QoS policy.

4.12 QoS Policy

4.12.1 Create a QoS Policy

```
./stor qos add { -n | --name } QOS_NAME [ --iops IOPS ] [ --read-iops READ_IOPS ] [ --write-iops WRITE_IOPS ] [ --bps BPS ] [ --read-bps READ_BPS ] [ --write-bps WRITE_BPS ] [ --iops-burst IOPS_BURST ] [ --read-iops-burst READ_IOPS_BURST ] [ --write-iops-burst WRITE_IOPS_BURST ] [ --bps-burst BPS_BURST ] [ --read-bps-burst READ_BPS_BURST ] [ --write-bps-burst WRITE_BPS_BURST ] [ --iops-burst-secs IOPS_BURST_SECS ] [ --read-iops-burst-secs READ_IOPS_BURST_SECS ] [ --write-iops-burst-secs WRITE_IOPS_BURST_SECS ] [ --bps-burst-secs BPS_BURST_SECS ] [ --read-bps-burst-secs READ_BPS_BURST_SECS ] [ --write-bps-burst-secs WRITE_BPS_BURST_SECS ] [ --reclaim-policy RECLAIM_POLICY ] [ --description ]
```

This command is used to create a QoS policy.

Parameter

Parameter	Description
--qos-name <i>QOS_NAME</i>	Specifies the name of a QoS policy. The value is a string of 1 to 64 case-sensitive characters. It can contain letters, digits, and hyphens (-). Only supports starting with a letter or a digit.
--iops <i>IOPS</i>	Specifies the limit of I/O operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
--read-iops <i>READ_IOPS</i>	Specifies the limit of read operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
--write-iops <i>WRITE_IOPS</i>	Specifies the limit of write operations per second. The value is an integer that ranges from -1 to 999,999,999, the default value is -1. -1 indicates no limit.
--bps <i>BPS</i>	Specifies the limit of I/O throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value,

	<p>representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
<p>--read-bps <i>READ_BPS</i></p>	<p>Specifies the limit of read throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
<p>--write-bps <i>WRITE_BPS</i></p>	<p>Specifies the limit of write throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814.

	<ul style="list-style-type: none"> ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
--iops-burst <i>IOPS_BURST</i>	<p>Specifies the burst limit of I/O operations per second.</p> <p>This setting only takes effect when --iops <i>IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>IOPS</i>, 999,999,999]. The default value is -1, indicating no limit.</p>
--read-iops-burst <i>READ_IOPS_BURST</i>	<p>Specifies the burst limit of read operations per second.</p> <p>This setting only takes effect when --read-iops <i>READ_IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>READ_IOPS</i> 999,999,999]. The default value is -1, indicating no limit.</p>
--write-iops-burst <i>WRITE_IOPS_BURST</i>	<p>Specifies the burst limit of write operations per second.</p> <p>This setting only takes effect when --write-iops <i>WRITE_IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>WRITE_IOPS</i>, 999,999,999]. The default value is -1, indicating no limit.</p>
--bps-burst <i>BPS_BURST</i>	<p>Specifies the burst limit of I/O throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --bps <i>BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3,906,250].

	<ul style="list-style-type: none"> ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3].
<p>--read-bps-burst <i>READ_BPS_BURST</i></p>	<p>Specifies the burst limit of read throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --read-bps <i>READ_BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3].
<p>--write-bps-burst <i>WRITE_BPS_BURST</i></p>	<p>Specifies the burst limit of write throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. The default value is -1. -1 indicates no limit.</p> <p>Only when --write-bps <i>WRITE_BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p>

	<ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i> 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3].
--iops-burst-secs <i>IOPS_BURST_SECS</i>	Specifies the duration of I/O operations with burst limit. Note: This setting is only effective when the IOPS Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--read-iops-burst-secs <i>READ_IOPS_BURST_SECS</i>	Specifies the duration of read operations with burst limit. Note: This setting is only effective when the read IOPS Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--write-iops-burst-secs <i>WRITE_IOPS_BURST_SECS</i>	Specifies the duration of write operations with burst limit. Note: This setting is only effective when the write IOPS Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--bps-burst-secs <i>BPS_BURST_SECS</i>	Specifies the duration of I/O throughput with burst limit. Note: This setting is only effective when the Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.

--read-bps-burst-secs <i>READ_BPS_BURST_SECS</i>	Specifies the duration of read throughput with burst limit. Note: This setting is only effective when the read Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--write-bps-burst-secs <i>WRITE_BPS_BURST_SECS</i>	Specifies the duration of write throughput with burst limit. Note: This setting is only effective when the write Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999, the default value is 1. The unit is seconds.
--reclaim-policy <i>RECLAIM_POLICY</i>	Specifies the QoS policy reclaim policy. Value: <ul style="list-style-type: none"> ● Delete: When all objects associated with the QoS policy are disassociated or deleted, the system will automatically remove the QoS policy. ● Retain: When all objects associated with the QoS policy are disassociated or deleted, the QoS policy itself remains and is not deleted. The default value is Retain.
--description	Specifies the description of the QoS policy.

Example

Create QoS policy qostest1 and add its description.

Note: If a description is provided, it must be entered interactively; the value must be a string 1–256 characters long.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos add -n qostest1 --iops 9999999 --bps 3800g --description
Enter the description for this QoS policy, limited to 256 characters:
It is a test QoS policy.
Created QoS policy qostest1 successfully.
```

4.12.2 Modify a QoS Policy

```
./stor qos set { -n | --name } QOS_NAME [ --new-name NEW_NAME ] [ --iops IOPS ] [ --read-iops READ_IOPS ] [ --write-iops WRITE_IOPS ] [ --bps BPS ] [ --read-bps READ_BPS ] [ --write-bps WRITE_BPS ] [ --iops-burst IOPS_BURST ] [ --read-iops-burst READ_IOPS_BURST ] [ --write-iops-burst WRITE_IOPS_BURST ] [ --bps-burst BPS_BURST ] [ --read-bps-burst READ_BPS_BURST ] [ --write-bps-burst WRITE_BPS_BURST ] [ --iops-burst-secs IOPS_BURST_SECS ] [ --read-iops-burst-secs READ_IOPS_BURST_SECS ] [ --write-iops-burst-secs WRITE_IOPS_BURST_SECS ] [ --bps-burst-secs BPS_BURST_SECS ] [ --read-bps-burst-secs READ_BPS_BURST_SECS ] [ --write-bps-burst-secs WRITE_BPS_BURST_SECS ] [ --reclaim-policy RECLAIM_POLICY ] [ --change-description ]
```

This command is used to modify the QoS policy.

Parameter

Parameter	Description
<code>--qos-name QOS_NAME</code>	Specifies the name of a QoS policy to be modified.
<code>--new-name NEW_NAME</code>	Specifies the new name of a QoS policy. The value is a string of 1 to 64 case-sensitive characters. It can contain letters, digits, and hyphens (-). Only supports starting with a letter or a digit.
<code>--iops IOPS</code>	Specifies the limit of I/O operations per second. The value is an integer that ranges from -1 to 999,999,999. -1 indicates no limit.
<code>--read-iops READ_IOPS</code>	Specifies the limit of read operations per second. The value is an integer that ranges from -1 to 999,999,999. -1 indicates no limit.
<code>--write-iops WRITE_IOPS</code>	Specifies the limit of write operations per second. The value is an integer that ranges from -1 to 999,999,999. -1 indicates no limit.
<code>--bps BPS</code>	Specifies the limit of I/O throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. -1 indicates no limit. <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000.

	<ul style="list-style-type: none"> ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
--read-bps <i>READ_BPS</i>	<p>Specifies the limit of read throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
--write-bps <i>WRITE_BPS</i>	<p>Specifies the limit of write throughput per second. The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. -1 indicates no limit.</p> <ul style="list-style-type: none"> ● If the unit is B/s, the value is an integer that ranges from -1 to 4,096,000,000,000. ● If the unit is KiB/s, the value is an integer that ranges from -1 to 4,000,000,000. ● If the unit is MiB/s, the value is an integer that ranges from -1 to 3,906,250. ● If the unit is GiB/s, the value is an integer that ranges from -1 to 3,814. ● If the unit is TiB/s, the value is an integer that ranges from -1 to 3.
--iops-burst <i>IOPS_BURST</i>	<p>Specifies the burst limit of I/O operations per second.</p> <p>This setting only takes effect when --iops <i>IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>IOPS</i>, 999,999,999]. -1 indicates no limit.</p>

<p>--read-iops-burst <i>READ_IOPS_BURST</i></p>	<p>Specifies the burst limit of read operations per second.</p> <p>This setting only takes effect when --read-iops <i>READ_IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>READ_IOPS</i> 999,999,999]. -1 indicates no limit.</p>
<p>--write-iops-burst <i>WRITE_IOPS_BURST</i></p>	<p>Specifies the burst limit of write operations per second.</p> <p>This setting only takes effect when --write-iops <i>WRITE_IOPS</i> is greater than or equal to 1. It must be set to -1 or to a positive integer within the range (<i>WRITE_IOPS</i>, 999,999,999]. -1 indicates no limit.</p>
<p>--bps-burst <i>BPS_BURST</i></p>	<p>Specifies the burst limit of I/O throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. -1 indicates no limit.</p> <p>Only when --bps <i>BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 3].
<p>--read-bps-burst <i>READ_BPS_BURST</i></p>	<p>Specifies the burst limit of read throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. -1 indicates no limit.</p>

	<p>Only when --read-bps <i>READ_BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3,814]. ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>READ_BPS</i>, 3].
<p>--write-bps-burst <i>WRITE_BPS_BURST</i></p>	<p>Specifies the burst limit of write throughput per second.</p> <p>The value is an integer. The unit abbreviation B/b, K/k, M/m, G/g, T/t can be entered after the value, representing B/s, KiB/s, MiB/s, GiB/s, TiB/s respectively. The default unit is B/s. -1 indicates no limit.</p> <p>Only when --write-bps <i>WRITE_BPS</i> is greater than or equal to 1, and the units are the same, can this setting take effect within the following range of values:</p> <ul style="list-style-type: none"> ● If the unit is B/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 4,096,000,000,000]. ● If the unit is KiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 4,000,000,000]. ● If the unit is MiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3,906,250]. ● If the unit is GiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3,814].

	<ul style="list-style-type: none"> ● If the unit is TiB/s, it must be set to -1 or to a positive integer within the range (<i>WRITE_BPS</i>, 3].
--iops-burst-secs <i>IOPS_BURST_SECS</i>	Specifies the duration of I/O operations with burst limit. Note: This setting is only effective when the IOPS Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999. The unit is seconds.
--read-iops-burst-secs <i>READ_IOPS_BURST_SECS</i>	Specifies the duration of read operations with burst limit. Note: This setting is only effective when the read IOPS Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999. The unit is seconds.
--write-iops-burst-secs <i>WRITE_IOPS_BURST_SECS</i>	Specifies the duration of write operations with burst limit. Note: This setting is only effective when the write IOPS Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999. The unit is seconds.
--bps-burst-secs <i>BPS_BURST_SECS</i>	Specifies the duration of I/O throughput with burst limit. Note: This setting is only effective when the Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999. The unit is seconds.
--read-bps-burst-secs <i>READ_BPS_BURST_SECS</i>	Specifies the duration of read throughput with burst limit. Note: This setting is only effective when the read Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999. The unit is seconds.
--write-bps-burst-secs <i>WRITE_BPS_BURST_SECS</i>	Specifies the duration of write throughput with burst limit. Note: This setting is only effective when the write Bps Burst feature is enabled. Value: The value is an integer that ranges from 1 to 999,999,999. The unit is seconds.
--reclaim-policy <i>RECLAIM_POLICY</i>	Specifies the QoS policy reclaim policy. Value: <ul style="list-style-type: none"> ● Delete: When all objects associated with the QoS policy are disassociated or deleted, the

	system will automatically remove the QoS policy. <ul style="list-style-type: none"> ● Retain: When all objects associated with the QoS policy are disassociated or deleted, the QoS policy itself remains and is not deleted.
--change-description	Change the description of the QoS policy.

Example

Modify QoS policy qostest1.

Note: If a description is provided, it must be entered interactively; the value must be a string 1-256 characters long.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos set -n qostest1 --new-name qos1 --
iops 800000 --bps 5500g --change-description
Enter the description for this QoS policy, limited to 256 characters:
It is qos1.
Set QoS policy qostest1 successfully.
```

4.12.3 Associate the QoS Policy with Objects

```
./stor qos { A | assoc } { -n | --name } QOS_NAME { -o | --object } OBJECT --list list<1-n>
```

This command is used to associate the QoS policy with objects.

Note:

- If the object type to be associated is LUN, the LUN status cannot be Deleting or Deletefailed; if the object type to be associated is storagepool or storagepoolforlun, the storage pool status cannot be Deleting.
- When multiple QoS policies are associated with the same object, the last operation takes precedence.
- For cluster mode, QoS policy enforcement follows these rules:
 - If a LUN has a QoS policy attached, that policy takes precedence.
 - If a LUN has no QoS policy attached:
 - ◆ If the LUN has both a `cachePool` and a `pool`, its QoS policy is the default QoS policy for LUNs in the `cachePool`. If the `cachePool` has not defined such a default policy, the LUN has no QoS policy.
 - ◆ If the LUN has only a `pool`, its QoS policy is the default QoS policy for LUNs in that `pool`. If the `pool` has not defined such a default policy, the LUN has no QoS policy.

Parameter

Parameter	Description
<code>-n QOS_NAME</code> or <code>--name QOS_NAME</code>	Specifies the name of a QoS policy.
<code>-o OBJECT</code> or <code>--object OBJECT</code>	Specifies the associated object type. Value: <ul style="list-style-type: none"> ● LUN. ● storagepool (only supported by cluster mode): Storage pool. ● storagepoolforlun (only supported by cluster mode): Default QoS policy for LUNs within the storage pool.
<code>--list list<1-n></code>	Specifies the object list. Multiple specific objects can be added at a time, separated by commas (,). Note: <ul style="list-style-type: none"> ● If the associated object type is LUN, enter the specific LUN name(s).

	<ul style="list-style-type: none">● If the associated object type is storagepool, enter the storage pool name(s).● If the associated object type is storagepoolforlun, enter the storage pool name(s).
--	---

Example

Associate the QoS policy with LUNs.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos assoc -n qos1 -o LUN --list luna1,lunb1,lunb2
Associate QoS policy qos1 with LUN luna1,lunb1,lunb2 successfully.
```

Associate the QoS policy with storage pools.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos assoc -n qos1 -o storagepool --list default,pool1
Associate QoS policy qos1 with storage pool default,pool1 successfully.
```

Disassociate the default QoS policy for LUNs in the pool.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos assoc -n qos1 -o storagepoolforlun -list default
Associate QoS policy qos1 with storage pool default for setting the default QoS policy for LUNs in the pool successfully.
```

4.12.4 Disassociate the QoS Policy from Objects

```
./stor qos { D | disass } { -n | --name } QOS_NAME { -o | --object } OBJECT --list
list<1-n>
```

This command is used to disassociate the QoS policy from objects.

Parameter

Parameter	Description
<code>-n QOS_NAME</code> or <code>--name QOS_NAME</code>	Specifies the name of a QoS policy.
<code>-o OBJECT</code> or <code>--object OBJECT</code>	Specifies the disassociated object type. Value: <ul style="list-style-type: none"> ● LUN. ● storagepool (only supported by cluster mode): Storage pool. ● storagepoolforlun (only supported by cluster mode): Default QoS policy for LUNs within the storage pool.
<code>--list list<1-n></code>	Specifies the object list. Multiple specific objects can be added at a time, separated by commas (,). Note: <ul style="list-style-type: none"> ● If the disassociated object type is LUN, enter the specific LUN name(s). ● If the disassociated object type is <code>storagepool</code>, enter the storage pool name(s). ● If the disassociated object type is <code>storagepoolforlun</code>, enter the storage pool name(s).

Example

Disassociate the QoS policy from LUNs.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos disass -n qos1 -o LUN --list
luna1,lunb1
Disassociate QoS policy qos1 with LUN luna1,lunb1 successfully.
```

Disassociate the QoS policy from storage pools.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos D -n qos1 -o storagepool --list
default,pool1
Disassociate QoS policy qos1 with storage pool default,pool1 successfully.
```

Disassociate the default QoS policy for LUNs in the pool.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos D -n qos1 -o storagepoolforlun --  
list default
```

```
Disassociate QoS policy qos1 with storage pool default for setting the default QoS policy for  
LUNs in the pool successfully. Associate QoS policy qos1 with storage pool default for setting the  
default QoS policy for LUNs in the pool successfully.
```

4.12.5 Delete a QoS Policy

```
./stor qos rm { -n | --name } QOS_NAME
```

This command is used to delete the specified QoS policy.

Note: A QoS policy can only be deleted when it is not associated with any objects.

Parameter

Parameter	Description
<code>-n QOS_NAME</code> or <code>--name QOS_NAME</code>	Specifies the name of the QoS policy to be deleted.

Example

Delete QoS policy Qos1.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos rm -n Qos1  
Delete QoS policy Qos1 successfully.
```

4.12.6 Query QoS Policy Information

```
./stor qos ls [ { -n | --name } QOS_NAME [ { -o | --object } OBJECT [ { -s | --status } STATUS ] ] ]
```

This command is used to query QoS policy information.

Parameter

Parameter	Description
<code>-n QOS_NAME</code> or <code>--name QOS_NAME</code>	Specifies the name of the QoS policy to be queried.
<code>-o OBJECT</code> or <code>--object OBJECT</code>	Specifies the associated object type. Value: <ul style="list-style-type: none"> ● LUN. ● storagepool (only supported by cluster mode): Storage pool. ● storagepoolforlun (only supported by cluster mode): Default QoS policy for LUNs within the storage pool.
<code>-s STATUS</code> or <code>--status STATUS</code>	Specifies the association status between the object and the QoS policy. Value: <ul style="list-style-type: none"> ● associated: Already associated. ● available: Available for association. The default value is associated.

Example

Query all QoS policies.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos ls
```

No.	QoS Policy Name	Reclaim Policy	IOPS (T/R/W)	Bandwidth (T/R/W)	Burst IOPS (T/R/W)	Burst Bandwidth (T/R/W)	Burst Duration (second)
1.	QoS-Test	Retain	8000/Unlimited/Unlimited	7.81 GiB/s Unlimited Unlimited	Unlimited	Unlimited	1/1/1/1/1/1
2.	QoS6	Retain	1000/600/400	1.95 GiB/s 1000 MiB/s 1.46 GiB/s	10000/6000/4000	2.05 GiB/s 1.95 GiB/s 1.95 GiB/s	1/1/1/1/1/1
3.	QoSPolicy20250806162435	Retain	1500/1000/1000	2.44 GiB/s Unlimited Unlimited	Unlimited	Unlimited	1/1/1/1/1/1
4.	QoSPolicy20250813142827	Retain	Unlimited	2 TiB/s Unlimited Unlimited	Unlimited	Unlimited	1/1/1/1/1/1
5.	Qostest2	Retain	Unlimited	Unlimited	Unlimited	Unlimited	1/1/1/1/1/1

Table1: Descriptions of information for all queried QoS policies

Item	Description
No.	Number.
QoS Policy Name	QoS policy name.
Reclaim Policy	QoS policy reclaim policy: <ul style="list-style-type: none"> ● Delete: When all objects associated with the QoS policy are disassociated or deleted, the system will automatically remove the QoS policy.

	<ul style="list-style-type: none"> ● Retain: When all objects associated with the QoS policy are disassociated or deleted, the QoS policy itself remains and is not deleted.
IOPS (T/R/W)	The limit of total/read/ write operations per second.
Bandwidth (T/R/W)	The limit of total/read/write throughput per second.
Burst IOPS (T/R/W)	The burst limit of total/read/write operations per second.
Burst Bandwidth (T/R/W)	The burst limit of total/read/write throughput per second.
Burst Duration (second)	The duration in seconds of I/O operations with burst limit. The corresponding sequence of QoS policy parameters is: burst IOPS (total/read/write), burst bandwidth (total/read/write).

Query QoS policy QoS6.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos ls -n QoS6
QoS Policy Name: QoS6
Reclaim Policy: Retain
IOPS (Total/Read/Write): 1000/600/400
Bandwidth (Total/Read/Write): Unlimited
Burst IOPS (Total/Read/Write): 10000/6000/4000
Burst Bandwidth (Total/Read/Write): Unlimited
Burst Duration (second): 1/1/1/1/1/1
Create Time: 2025-07-31 16:28:51
Description: It is QoS6.
```

Table2: Descriptions of information for the specified QoS policy being queried

Item	Description
QoS Policy Name	QoS policy name.
Reclaim Policy	QoS policy reclaim policy: <ul style="list-style-type: none"> ● Delete: When all objects associated with the QoS policy are disassociated or deleted, the system will automatically remove the QoS policy. ● Retain: When all objects associated with the QoS policy are disassociated or deleted, the QoS policy itself remains and is not deleted.
IOPS (Total/Read/Write)	The limit of total/read/ write operations per second.
Bandwidth (Total/Read/Write)	The limit of total/read/write throughput per second.
Burst IOPS (Total/Read/Write)	The burst limit of total/read/write operations per second.
Burst Bandwidth (Total/Read/Write)	The burst limit of total/read/write throughput per second.
Burst Duration (second)	The duration in seconds of I/O operations with burst limit. The corresponding sequence of QoS policy parameters is: burst IOPS (total/read/write), burst bandwidth (total/read/write).

Create Time	The creation time of the QoS policy.
Description	The description of the QoS policy.

Standalone mode: Query all LUNs associated with the QoS policy QoS1.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos ls -n QoS1 -o lun
```

No.	LUN Name	Storage Mode	Capacity	Status
1.	luna1	Local	100 GiB	Normal
2.	lunc3	Cache	390 GiB	Normal

Standalone mode: Query all LUNs that can still be associated with QoS policy QoS1.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos ls -n QoS1 -o LUN -s available
```

No.	LUN Name	Storage Mode	Capacity	Status
1.	lunb1	Local	202 GiB	Normal

Cluster mode: Query all LUNs associated with the QoS policy QoS6.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos ls -n QoS6 -o lun
```

No.	LUN Name	Storage Mode	Capacity	Local Storage Class	Status
1.	lun01a	Local	200 GiB	EC 2+1+16KiB	Normal
2.	lunn01a-clone1	Local	200 GiB	EC 2+1+16KiB	Normal
3.	lun03a	Local	303 GiB	EC 2+1+16KiB	Normal
4.	lunc03b	Cache	333 GiB	EC 2+1+16KiB	Normal

Cluster mode: Query all LUNs that can still be associated with QoS policy QoS6.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos ls -n QoS6 -o LUN -s available
```

No.	LUN Name	Storage Mode	Capacity	Local Storage Class	Status
1.	lun11	Local	1 GiB	EC 2+1+16KiB	Normal
2.	lun1	Local	1 GiB	EC 2+1+16KiB	Normal
3.	lun-sttest01	Local	123 GiB	EC 2+1+16KiB	Normal
4.	2a21fca428422904	Local	100 GiB	EC 2+1+16KiB	Normal
5.	clone-01	Local	77 GiB	EC 2+1+16KiB	Normal
6.	123	Local	11 GiB	EC 2+1+16KiB	Normal
7.	74a607cb8c8273fd	Local	333 GiB	EC 2+1+16KiB	Normal

Table3: Descriptions of LUN information associated with the QoS policy

Item	Description
No.	Number.
LUN Name	LUN name.
Storage Mode	The storage mode for the LUN: <ul style="list-style-type: none"> ● Local: Store all data at local only. ● Cache: Store part of hot data at local and store all data in cloud asynchronously. ● Storage: Store all data at local and asynchronously store it in cloud.
Capacity	The storage capacity of LUN.
Local Storage Class	LUN redundancy mode (only supported by cluster mode): <ul style="list-style-type: none"> ● single-copy ● 2-copy ● 3-copy ● EC N+M.
Status	LUN status: <ul style="list-style-type: none"> ● Normal: The LUN is normal. ● Deleting: The LUN is being deleted. ● DeleteFailed: The LUN deletion failed. ● Recovering: The LUN is recovering. ● RecoverFailed: The LUN recovered failed. ● Rollbacking: The LUN is rollbacking. ● Flattening: The chain-breaking process between the clone LUN and the snapshot is underway. The clone LUN is copying data from the source LUN. After the copy is complete, it becomes an independent LUN that no longer depends on the snapshot or source LUN. ● Importing: The LUN is being imported with backup data. ● Wiping: The LUN is being wiped. ● WipeFailed: The LUN data wiping failed.

Cluster mode: Query all storage pools associated with the QoS policy QoS6.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos ls -n QoS6 -o storagepool
+-----+-----+-----+-----+-----+
| No. | Pool Name | Fault Domain | Total Capacity | Used Capacity |
+-----+-----+-----+-----+-----+
| 1. | default | server | 278.57 GiB | 65.37 GiB |
| 2. | pool1 | server | 0 B | 0 B |
+-----+-----+-----+-----+-----+
```

Cluster mode: Query the storage pools whose default QoS policy for LUNs in the pool is set to QoS6.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor qos ls -n QoS6 -o storagepoolforlun
+-----+-----+-----+-----+-----+
```

No.	Pool Name	Fault Domain	Total Capacity	Used Capacity
1.	pool1	server	0 B	0 B
2.	default	server	278.57 GiB	65.37 GiB

Table4: Descriptions of storage pool information associated with the QoS policy

Item	Description
No.	Number.
Pool Name	The name of the storage pool.
Fault Domain	Fault domain level of the storage pool: <ul style="list-style-type: none"> ● path. ● server. ● rack. ● room.
Total Capacity	Total capacity of the storage pool.
Used Capacity	Used capacity of the storage pool.

4.13 Cluster Topology (Cluster Mode)

4.13.1 Create a Topology Node

```
./stor topology add { -n | --name } NODE_NAME { -t | --type } NODE_TYPE [ --parent-node PARENT_NODE ]
```

This command is used to create a topology node.

Parameters

Parameter	Description
<code>-n NODE_NAME</code> or <code>--name NODE_NAME</code>	Specifies the name of a topology node. The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit.
<code>-t NODE_TYPE</code> or <code>--type NODE_TYPE</code>	Specifies the type of a topology node. Value: <ul style="list-style-type: none">● rack: rack type.● room: room type.
<code>--parent-node PARENT_NODE</code>	Specifies the name of the parent node. The root node is used by default.

Example

Create the topology node room3 of room type.

Note: If you do not want to add node description, press Enter.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor topology add --name room3 --type room
Enter the description for this node, limited to 50 characters:
room3 is new node.
Created node room3 successfully.
```

4.13.2 Modify Topology Node Information

```
./stor topology set { -n | --name } NODE_NAME [ --change-description ] [ --new-name
NEW_NAME ] [ --parent-node PARENT_NODE ]
```

This command is used to modify the topology node information.

Note: You cannot modify the topology node information at the path level.

Parameters

Parameter	Description
<code>-n NODE_NAME</code> or <code>--name NODE_NAME</code>	Specifies the name of the topology node to be modified. The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit.
<code>--change-description</code>	Change the description of the topology node.
<code>--new-name NEW_NAME</code>	Specifies the new name of the topology node. The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit.
<code>--parent-node PARENT_NODE</code>	Change the parent node of a topology node.

Examples

- Change the parent node of hblock_4 to room3.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor topology set --name hblock_4 --
parent-node room3
Set node hblock_4 successfully.
```

- Change the name of node room3 to room-3.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor topology set --name room3 --new-
name room-3
Set node room3 successfully.
```

4.13.3 Delete a Topology Node

```
./stor topology rm { -n | --name } NODE_NAME [ -d | --del-data ] [ -f | --force ]
```

This command is used to delete a topology node.

Note:

- For a server node, all paths of this node do not belong to any storage pool. You can delete this server from the cluster. Otherwise, you can only forcibly remove this node, which may cause data loss.
- A room or rack node can be deleted from the cluster only if it does not have any child nodes. Otherwise, the node cannot be removed, nor can it be forcibly removed.
- For nodes with the same name, such as room1:rack1 and room2:rack1, *NODE_NAME* can carry the name of the parent node. Otherwise the node may not be found.
- You cannot use this command to delete the path node. If you want to delete the path node, please see **Remove a Disk Path**.

Parameters

Parameter	Description
<code>-n <i>NODE_NAME</i></code> or <code>--name <i>NODE_NAME</i></code>	Specifies the name of the topology node to delete. The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit.
<code>-d</code> or <code>--del-data</code>	Deletes a server node and deletes HBlock data on the disk paths of the server. This parameter is supported when only server nodes are deleted.
<code>-f</code> or <code>--force</code>	Forcibly deletes a topology node. This parameter is supported when only server nodes are deleted. Note: Forcibly deleting topology nodes may cause data loss.

Examples

- Delete topology node room3.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor topology rm -n room3
Deleted node room3 successfully.
```

- Delete topology node hblock_4.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor topology rm -n  
room2:rack2:hblock4  
Deleted node room2:rack2:hblock4 successfully.
```

4.13.4 Query Topology Information

```
./stor topology ls [ { -n | --name } NODE_NAME ]
```

This command is used to query topology information.

Parameter

Parameter	Description
<code>-n NODE_NAME</code> or <code>--name NODE_NAME</code>	Specifies the topology of the node to be queried. Note: Not supporting querying path nodes.

Examples

- Query the information of a whole topology.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor topology ls
Type Name
root default
  room room1
    rack rack1
      server hblock_2
        path /mnt/stor
  room room2
    rack rack2
      server hblock_1
        path /mnt/stor
      server hblock_3
        path /mnt/stor
        path /mnt/storage01
```

- Query the topology information of node rack2.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor topology ls -n rack2
Type Name
rack rack2
  server hblock_1
    path /mnt/stor
  server hblock_3
    path /mnt/stor
    path /mnt/storage01
```

Description of topology information

Item	Description
Type	Node type:

	<ul style="list-style-type: none">● root: root node.● room.● rack.● server.● path.
Name	The name of the node.

4.14 Server Operations

4.14.1 Add a Server (Cluster Mode)

```
./stor server add {-s | --server } SERVER_IP[:PORT ] [ --parent-node PARENT_NODE ] [ { -n
| --name } NODE_NAME ] [ { -p | --path } PATH&<1-n> ] [ --capacity-quota CAPACITY_QUOTA ]
[--port-range PORT1-PORT2 ] [ --data-port1 DATA_PORT1 ] [ --iscsi-port ISCSI_PORT ] [ --
management-port1 MANAGEMENT_PORT1 ] [ --management-port2 MANAGEMENT_PORT2 ] [ --
management-port3 MANAGEMENT_PORT3 ] [ --management-port4 MANAGEMENT_PORT4 ] [ --
management-port5 MANAGEMENT_PORTS5 ] [ --management-port6 MANAGEMENT_PORT6 ]
```

This command is used to add a server.

Note:

- The server added to the cluster can only be added after installing HBlock.
- Please ensure that the Linux user has permission for the required ports. By default, Linux systems do not open ports less than 1024 to ordinary users without root privileges.
- When setting the port range (**--port-range** *PORT1-PORT2*), please avoid overlapping with the local temporary port (*ip_local_port_range*) range of the Linux system, otherwise, the port used by the HBlock service may be occupied. Run the command **cat /proc/sys/net/ipv4/ip_local_port_range** to view the local temporary port range.

Parameters

Parameter	Description
-s <i>SERVER_IP[:PORT]</i> or --server <i>SERVER_IP[:PORT]</i>	Specifies the IP address (<i>SERVER_IP</i>) and API port (<i>PORT</i>) of the added server. Server IP: IPv4 or [IPv6]. API port: The value is an integer that ranges from 1 to 65535, the default value is 1443. It needs to be consistent with the API port number set when installing HBlock on the server.
--parent-node <i>PARENT_NODE</i>	Specifies a parent node. The root node is used by default.
-n <i>NODE_NAME</i> or --name <i>NODE_NAME</i>	Specifies the name of a server node. The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit. The server ID is used by default as node name.

<p>-p <i>PATH</i> &<1-n> or --path <i>PATH</i> &<1-n></p>	<p>Specifies a disk path. Multiple disk paths are supported. Value: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; : . Note: The disk path is used to store data and it is not recommended to share disks or file systems with the operating system.</p>
<p>--capacity-quota <i>CAPACITY_QUOTA</i></p>	<p>Specifies the capacity quota of the disk path, that is, for each disk path added to server, the total amount of data that can be written by HBlock. Once the space used by HBlock reaches the quota, data writing is immediately blocked, and no more space is allowed to be used beyond the quota. Value: Integer. The value is smaller than the total capacity of the disk where the disk path resides. The units are K/k, M/m, G/g, T/t, and P/p. The default unit is G/g. A negative integer means unlimited writing and 0 means write inhibit. Writes are not limited by default. Note: If multiple disk paths are added at one time, only one capacity quota can be configured, and the configured capacity quota applies to all added disk paths.</p>
<p>--port-range <i>PORT1-PORT2</i></p>	<p>Specifies the port range. Storage services and services without specified ports will automatically be assigned ports from this range. The value is an integer that ranges from 1 to 65535. <i>PORT1</i> is the minimum value of the port range, <i>PORT2</i> is the maximum value of the port range, and <i>PORT1</i><<i>PORT2</i>. The default value of <i>PORT1</i> is 20000, and the default value of <i>PORT2</i> is 20500. Note: It is recommended that the specified port range contains at least 500 ports.</p>
<p>--data-port1 <i>DATA_PORT1</i></p>	<p>Specifies the data port 1 The value is an integer that ranges from 1 to 65535.</p>
<p>--iscsi-port <i>ISCSI_PORT</i></p>	<p>Specifies the iSCSI port. The value is an integer that ranges from 1 to 65535. The default value is 3260.</p>
<p>--management-port1 <i>MANAGEMENT_PORT1</i></p>	<p>Specifies the management service port 1. The value is an integer that ranges from 1 to 65535.</p>
<p>--management-port2 <i>MANAGEMENT_PORT2</i></p>	<p>Specifies the management service port 2. The value is an integer that ranges from 1 to 65535.</p>

--management-port3 <i>MANAGEMENT_PORT3</i>	Specifies the management service port 3. The value is an integer that ranges from 1 to 65535.
--management-port4 <i>MANAGEMENT_PORT4</i>	Specifies the management service port 4. The value is an integer that ranges from 1 to 65535.
--management-port5 <i>MANAGEMENT_PORT5</i>	Specifies the management service port 5. The value is an integer that ranges from 1 to 65535.
--management-port6 <i>MANAGEMENT_PORT6</i>	Specifies the management service port 6. The value is an integer that ranges from 1 to 65535.

Examples

- Add server 192.168.0.202 and specify /mnt/storage01 as the disk path.

Note: If you do not want to add the description of the new server, press Enter.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor server add -s 192.168.0.202 -p /mnt/storage01
Enter the description for this node, limited to 50 characters:
new server
Added server 192.168.0.202 with path /mnt/storage01 successfully.
```

- Add server 192.168.0.117 with parent node rack1, name server node as server4, and specify /mnt/storage01 as the disk path.

Note: If you do not want to add the description of the new server, press Enter.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor server add -s 192.168.0.117 --parent-node rack1 -n server4 -p /mnt/storage01
Enter the description for this node, limited to 50 characters:
rack1
Added server 192.168.0.117 with path /mnt/storage01 successfully.
```

4.14.2 Modify Server Port Range

```
./stor server set { -i | --item } parameter [ { -n | --server } { SERVER_ID &<1-n> | all | default }* ] --port-range PORT1-PORT2
```

This command is used to modify the port range of the HBlock server for HBlock-related services.

Note:

- The newly added data service ports will be selected from the modified port range, and the used port values will remain unchanged.
- When setting the port range (**--port-range** *PORT1-PORT2*), please avoid overlapping with the local temporary port (*ip_local_port_range*) range of the Linux system, otherwise, the port used by the HBlock service may be occupied. Run the command `cat /proc/sys/net/ipv4/ip_local_port_range` to view the local temporary port range.

Parameters

Parameter	Description
-i parameter or --item parameter	Set the parameters of the HBlock server.
-n { <i>SERVER_ID</i> &<1-n> all default }* or --server { <i>SERVER_ID</i> &<1-n> all default }*	<p>Specifies the HBlock server for which to modify the port range.</p> <p>Value:</p> <ul style="list-style-type: none"> ● <i>SERVER_ID</i> &<1-n>: Specifies the HBlock server ID(s) for which to modify the port range. Multiple server IDs can be specified at once, separated by commas (,). ● all: Specifies to modify the port range for all HBlock servers. ● default: Specifies the port range for nodes that will join the cluster subsequently. <p>The default value is the server ID of the server where the command is executed.</p> <p>Note: Multiple parameters can be specified simultaneously, separated by commas, e.g., -n hblock_1,hblock_2,default.</p>
--port-range <i>PORT1-PORT2</i>	<p>Specifies the port range for related services.</p> <p>The value is an integer that ranges from 1 to 65535. <i>PORT1</i> is the minimum value of the port range, <i>PORT2</i> is the maximum value of the port range, and <i>PORT1</i><<i>PORT2</i>.</p>

Note: It is recommended that the specified port range contains at least 500 ports.

Examples

Modify the port range for HBlock-related services on server hblock_2 to 19000-20500.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server set -i parameter -n  
hblock_2 --port-range 19000-22500  
Set successfully.
```

4.14.3 Set Server Target Portal IP

```
./stor server set { -i | --item } parameter [ { -n | --server } SERVER_ID ] { -s | --status } STATUS [ { -e | --target-portal-ip } TARGET_PORTAL_IP { -P | --port } PORT ]
```

This command is used to set server target portal IP for the specified HBlock server.

If the server and client are not in the same network segment (for example, the server is on the intranet and the client is on the extranet) and are connected through a NAT device (such as a router), you need to add the extranet address and port of the NAT device to the server, so that clients on the extranet can normally establish iSCSI connections with the target of the server.

Parameters

Parameter	Description
-i parameter or --item parameter	Set the parameters of the HBlock server.
-n SERVER_ID or --server SERVER_ID	Specifies the HBlock server ID. The default value is the server ID of the server where the command is executed.
-s STATUS or --status STATUS	Specifies the status of target portal IP: <ul style="list-style-type: none"> ● Enabled (on): Enable target portal IP. ● Disabled (off): Disable target portal IP.
-e TARGET_PORTAL_IP or --target-portal-ip TARGET_PORTAL_IP	Target portal IP, in IPv4 or [IPv6] format.
-P PORT or --port PORT	Target portal port number. The value is an integer that ranges from 1 to 65535.

Examples

Set target portal IP for server hblock_1: the target portal IP is 101.89.213.5, and the target portal port number is 3261.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server set -i parameter -n hblock_1 -s on -e 101.89.213.5 -P 3261
Set successfully.
```

4.14.4 Set the Memory Available for HBlock on the Server

```
./stor server set { -i | --item } parameter [ { -n | --server } { SERVER_ID &<1-n> | all | default }* ] { --max-memory-ratio MAX_MEMORY_RATIO | --max-memory-size MAX_MEMORY_SIZE }*
```

This command is used to set the memory available for HBlock on the server.

Note: If specify the server IDs or all servers, after executing this command, you must run **./stor restart --scope cs,ps** on the corresponding server for the changes to take effect. If the available memory for HBlock is adjusted from less than 8 GiB to greater than or equal to 8 GiB, or vice versa, you must also execute **./stor restart -scope ms** on the corresponding server:

- Specific server: If you modify the HBlock memory for a specific server, execute the relevant commands on that server only.
- **all**: If the target servers for the HBlock memory parameter are set to **all**, execute the relevant commands on all servers.

Parameters

Parameter	Description
-i parameter or --item parameter	Set the parameters of the HBlock server.
-n { <i>SERVER_ID</i> &<1-n> all default }* or --server { <i>SERVER_ID</i> &<1-n> all default }*	<p>Specifies the HBlock server for which to set memory parameters.</p> <p>Value:</p> <ul style="list-style-type: none"> ● <i>SERVER_ID</i> &<1-n>: Specifies the HBlock server ID(s) for which to modify memory parameters. Multiple server IDs can be specified at once, separated by commas (,). ● all: Specifies to modify memory parameters for all HBlock servers. ● default: Specifies the default HBlock memory usage for nodes that will join the cluster subsequently. <p>The default value is the server ID of the server where the command is executed.</p> <p>Note: Multiple parameters can be specified simultaneously, separated by commas, e.g., -n hblock_1,hblock_2,default.</p>
--max-memory-ratio <i>MAX_MEMORY_RATIO</i>	Specifies the upper limit of the proportion of system total memory that HBlock can use.

	<p>The value range is [0.01, 1], supporting up to 2 decimal places.</p> <p>Note:</p> <ul style="list-style-type: none"> ● The actual memory occupied by HBlock is related to the number of disks added to the storage pool and may exceed the set proportion limit. Once exceeded, the system event "InsufficientMemory" will be triggered. ● When both --max-memory-ratio <i>MAX_MEMORY_RATIO</i> and --max-memory-size <i>MAX_MEMORY_SIZE</i> are set, the minimum value of the two is used.
<p>--max-memory-size <i>MAX_MEMORY_SIZE</i></p>	<p>Specifies the maximum memory available for HBlock. Values: The value is an integer. Unit suffixes M/m, G/g, or T/t can be appended to the number, representing MiB, GiB, and TiB respectively. The default unit is MiB.</p> <ul style="list-style-type: none"> ● If the unit is MiB, the value is an integer that ranges from 2048 to 8,796,093,022,207. ● If the unit is GiB, the value is an integer that ranges from 2 to 8,589,934,591. ● If the unit is TiB, the value is an integer that ranges from 1 to 8,388,607. <p>Note:</p> <ul style="list-style-type: none"> ● The actual memory occupied by HBlock is related to the number of disks added to the storage pool and may exceed the set proportion limit. Once exceeded, the system event "InsufficientMemory" will be triggered. ● When both --max-memory-ratio <i>MAX_MEMORY_RATIO</i> and --max-memory-size <i>MAX_MEMORY_SIZE</i> are set, the minimum value of the two is used.

Examples

Set the maximum memory available for HBlock on server hblock_1 to 8388000 TiB.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server set -i parameter -n
hblock_1 --max-memory-size 8388000T
Set successfully.
```

4.14.5 Set Server Default Disk Path (Standalone Mode)

```
./stor server set { -i | --item } parameter { -p | --path } DEFAULT_PATH
```

This command is used to set the default disk path for the HBlock server.

Parameters

Parameter	Description
<code>-i parameter</code> or <code>--item parameter</code>	Set the parameters of the HBlock server.
<code>-p DEFAULT_PATH</code> or <code>--path DEFAULT_PATH</code>	Specifies the default disk path. The disk path must have been added to the HBlock system and the status is Normal.

Examples

Set the default disk path of the HBlock server to /mnt/stor01.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server set -i parameter -p /mnt/stor01  
Set successfully.
```

4.14.6 Migrate Base Services on the Server (Cluster Mode)

```
./stor server set { -i | --item } service { -n | --server } SERVER_ID --dest-server
DEST_SERVER_ID --migrate SERVICE [ --meta-dir META_DIR ]
```

This command is used to migrate base services on the HBlock server, including mdm (metadata management service), ls (log service), and cs (coordination service).

Application scenario: If the server where base services are located is damaged or broken down, you can migrate base services on the server to prevent HBlock service interruption.

You can run the `./stor server ls { -n | --server } SERVER_ID` command to query the status of base services. Only one base service can be migrated at a time. If there is an ongoing migration process for a base service, you must wait for it to complete before performing other service migration operations.

Note:

- When migrating the ls service, ensure that the two mdm services and the other two ls services are up and that ms services (management services) on all nodes except the source server are normal. The service status can be obtained by querying the server command.
- When migrating the mdm service, ensure that another mdm service is up and that the ps service (protocol resolution service) and ms service (management service) on all nodes except the source server are normal. The service status can be obtained by querying the server command.
- When migrating the cs service, ensure that the other two cs services are up and that the ps service (protocol resolution service) and ms service (management service) on all nodes except the source server are normal. The service status can be obtained by querying the server command.

Parameter

Parameter	Description
<code>-i service</code> or <code>--item service</code>	Migrate base services on the server.
<code>-n SERVER_ID</code> or <code>--server SERVER_ID</code>	Specifies the source HBlock server ID.
<code>--dest-server DEST_SERVER_ID</code>	Specifies the destination server ID.
<code>--migrate SERVICE</code>	Specifies the base services to be migrated. Value: <ul style="list-style-type: none"> ● mdm: metadata management service. ● ls: log service. ● cs: coordination service.

<code>--meta-dir META_DIR</code>	<p>The disk path of the migration service is used to store relevant data information of the basic service.</p> <p>Note: In order to improve read-write performance, it is recommended that the disk path of the migration service be independent of the installation disk path and the disk path where data is stored.</p> <p>Value: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; . : . The default value is the installation disk path.</p>
----------------------------------	---

Examples

Migrate the mdm service from server hblock_1 to server hblock_3.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server set -i service -n hblock_1
--dest-server hblock_3 --migrate mdm --meta-dir /mnt/stor
Start migrating service mdm on hblock_1 to hblock_3. You can list server to check whether
it is completed.
```

4.14.7 Remove a Server (Cluster Mode)

```
./stor server rm {-n | --server } SERVER_ID [ -d | --del-data ] [ -f | --force ]
```

This command is used to remove a server from the HBlock cluster.

Note:

- If there is an iSCSI target on the server to be removed, and the high availability type of the LUN in the iSCSI target is ActiveStandby. When the server is removed, the service will not be interrupted, and the iSCSI target of the LUN on the server will be switched to another server, the client needs to reconnect to the iSCSI target.
- If there is an iSCSI target on the server to be removed, and the high availability type of the LUN in the iSCSI target is Disabled. When the server is removed, the service will be interrupted, and the iSCSI target of the LUN on the server will be switched to another server, the client needs to reconnect to the iSCSI target. When removing the server, there is a risk of data loss.
- If the collected log is saved in the server installation directory after log collection, it will be deleted after the server is removed. If the collected log is saved in the disk path of HBlock, and you delete the HBlock data on the disk paths of the server when the server is removed, the log will also be deleted.
- While a server is being removed, other servers cannot be removed. If it must be removed, please use force removal, but there is a risk of data loss.
- All disk paths on this node do not belong to any storage pool. Therefore, you can remove this server. Otherwise, the server cannot be removed. If you must remove the server, forcibly remove it, which may cause data loss.
- If a disk path of a server to remove belongs to the base storage pool and is the only node in the available fault domain in the base storage pool, the disk path cannot be removed.
- If base services are running on the server, the server must not be removed.

Parameters

Parameter	Description
<code>-n SERVER_ID</code> or <code>--server SERVER_ID</code>	Specifies ID of the server to be removed.
<code>-d</code> or <code>--del-data</code>	When removing a server, delete the data in the HBlock disk path on the server.
<code>-f</code> or <code>--force</code>	Forcibly remove the server. Note: Forcibly removing the server may cause a risk of data loss.

Examples

Remove the server with server ID hblock_4.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server rm -n hblock_4
The software will be deleted on the specified server. Are you sure to remove this server?
[Yes/No]
y
Processing...
Start removing server hblock_4. You can list server to check whether it is completed.
You can reconnect the following iSCSI LUN(s):
1.iSCSI LUN name:lun04a
iSCSI target: iqn.2012-08.cn.ctyunapi.oos:target04.7(192.168.0.209:3260,ColdStandby)
iSCSI target: iqn.2012-08.cn.ctyunapi.oos:target04.8(192.168.0.72:3260,Active)
```

Note: The target associated with the LUN is in the ColdStandby status, which means that the target is being migrated from the deleted server to other servers in the cluster. During the migration process, the target is in the ColdStandby status. After the migration is completed, target will change to Active or Standby status.

4.14.8 Query Server Information

```
./stor server ls [ { -n | --server } { SERVER_ID &<1-n> | all | default }* ] [ --port ]
```

This command is used to query the relevant information of the HBlock server.

Parameters

Parameter	Description
-n { SERVER_ID &<1-n> all default }* or --server { SERVER_ID &<1-n> all default }*	Specifies the server to query. Value: <ul style="list-style-type: none"> ● <i>SERVER_ID</i> &<1-n>: Specifies the HBlock server ID(s) for which to query. Multiple server IDs can be specified at once, separated by commas (,). ● <i>all</i>: Specifies to to query all servers. ● <i>default</i>: Specifies to query the default parameters of the cluster. Note: Multiple parameters can be specified simultaneously, separated by commas, e.g., <i>-n hblock_1,hblock_2,default</i> .
--port	Specifies to query the port numbers used by HBlock on the server.

Examples

- Standalone Mode: Query HBlock server information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls
+-----+-----+-----+-----+-----+-----+-----+
| No. | Server ID | Server Name | Status | Public Address | Cluster Address | Recent Start Time |
+-----+-----+-----+-----+-----+-----+-----+
| 1. | hblock_1 | hblockserver | Connected | 192.168.0.66:3260 | 192.168.0.66 | 2025-02-17 17:51:27 |
+-----+-----+-----+-----+-----+-----+-----+
```

- Cluster Mode: Query HBlock server information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls
+-----+-----+-----+-----+-----+-----+-----+
| No. | Server ID | Server Name | Status | Public Address | Cluster Address | Recent Start Time |
+-----+-----+-----+-----+-----+-----+-----+
| 1. | hblock_1 (M**) | hblockserver | Connected | 192.168.0.192:3260 | 192.168.0.192 | 2025-01-03 10:47:27 |
| 2. | hblock_2 (**) | server2 | Connected | 192.168.0.202:3260 | 192.168.0.202 | 2025-01-03 10:46:02 |
| 3. | hblock_3 (**) | server3 | Connected | 192.168.0.102:3260 | 192.168.0.102 | 2025-01-03 10:47:16 |
+-----+-----+-----+-----+-----+-----+-----+
```

Table 1: HBlock server information description.

Item	Description
No.	Number.
Server ID	Server ID.

	<ul style="list-style-type: none"> ● A server ID followed by (M) indicates that it is the master server. ● A server ID followed by (**) indicates that it is a base node server. ● If only the server ID is present, it belongs to a non-base node server.
Server Name	Server name.
Status	Server status: <ul style="list-style-type: none"> ● Connected: The server is connected. ● Disconnected: The server is disconnected. ● Removing: The server is being removed from the cluster.
Public Address	Public network IP and port.
Cluster Address	Cluster network IP and port.
Recent Start Time	The last successful startup time of the HBlock service on this node. -: Indicates that the HBlock service is in a stopped status.

- Standalone Mode: Query the server information of the server ID hblock_1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls -n hblock_1
Server Name: songt-005
Server ID: hblock_1
Status: Connected
Public Address: 192.168.0.66:3260
Cluster Address: 192.168.0.66
Recent Start Time: 2026-01-04 10:11:14
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.9,"maxMemorySize":9223372036854775807}
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+-----+
| No. | Path          | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.  | /mnt/stor01 (*) | 2.35 GiB     | 93.29 GiB     | 1.02 GiB           | Unlimited      | Healthy      | -             |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

- Cluster Mode: Query the server information with server ID hblock_1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls -n hblock_1
Server Name: hblockserver
Server ID: hblock_1
Node Name: default:hblock_1
Parent Node: default
Status: Connected
Master Server: true
Base Server: true
Base Service: mdm (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
               ls (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
               cs (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
Public Address: 192.168.0.65:3260
Cluster Address: 192.168.0.65
Recent Start Time: 2026-01-04 09:19:45
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.9,"maxMemorySize":9223372036854775807}
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+-----+
| No. | Path          | Storage Pool | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail | Data Service |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.  | /mnt/stor01 | pool1       | 1.69 GiB     | 92.86 GiB     | 1.05 GiB           | Unlimited      | Healthy      | -             | ds-1         |
| 2.  | /mnt/stor02 | pool2       | 124 KiB      | 92.86 GiB     | 28 KiB             | Unlimited      | Healthy      | -             | ds-2         |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

- Cluster Mode: Query detailed information of all servers and default parameters.

```
[root@k8s-master CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls -n all,default
Default Properties:
=====
{"default":{"portRange":"19000-25005","maxMemoryRatio":0.9,"maxMemorySize":9223372036854775807}}

Server Properties for hblock_1:
```

```

=====
Server Name: hblockserver
Server ID: hblock_1
Node Name: default:hblock_1
Parent Node: default
Status: Connected
Master Server: true
Base Server: true
Base Service: ls (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
               cs (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
Public Address: 192.168.0.65:3260
Cluster Address: 192.168.0.65
Recent Start Time: 2026-01-04 09:19:45
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.9,"maxMemorySize":9222703533785088000}
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| No. | Path      | Storage Pool | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail | Data Service |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.  | /mnt/stor01 | pool1       | 7.36 GiB     | 92.86 GiB     | 6.69 GiB          | Unlimited      | Healthy      | -             | ds-1         |
| 2.  | /mnt/stor02 | pool2       | 124 KiB      | 92.86 GiB     | 28 KiB            | Unlimited      | Healthy      | -             | ds-2         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Ports:
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Port Classification | Port |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Port Range         | 19000-25000 |
| iSCSI Port         | 3260         |
| API Port           | 1443         |
| Web Port           | 2443         |
| Data Port          | 20007        |
| Storage Port       | 20017,20018,20019,20020,20021,20022 |
| Management Port    | 20000,20003,20010,20013,20014,20015,20016 |
| Metadata Port      | 20001,20002,20006,20008,20009,20012 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

Server Properties for hblock_2:
=====
Server Name: k8s-master
Server ID: hblock_2
Node Name: default:hblock_2
Parent Node: default
Status: Connected
Master Server: false
Base Server: true
Base Service: mdm (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
               ls (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
               cs (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
Public Address: 192.168.0.64:3260
Cluster Address: 192.168.0.64
Recent Start Time: 2026-01-07 17:59:47
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.9,"maxMemorySize":9223372036854775807}
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| No. | Path      | Storage Pool | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail | Data Service |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.  | /mnt/stor01 | pool2       | 16.98 GiB    | 39.98 GiB     | 8 KiB             | Unlimited      | Healthy      | -             | ds-2         |
| 2.  | /mnt/stor02 | pool1       | 6.73 GiB     | 92.86 GiB     | 6.69 GiB          | Unlimited      | Healthy      | -             | ds-1         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Ports:
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Port Classification | Port |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Port Range         | 19000-25000 |
| iSCSI Port         | 3260         |
| API Port           | 1443         |
| Web Port           | 2443         |
| Data Port          | 20007        |
| Storage Port       | 20017,20018,20019,20020,20021,20022 |
| Management Port    | 20000,20003,20010,20013,20014,20015,20016 |
| Metadata Port      | 20001,20002,20004,20005,20006,20008,20009,20011,20012 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

Server Properties for hblock_3:
=====
Server Name: songt-0006
Server ID: hblock_3
Node Name: default:hblock_3
Parent Node: default
Status: Connected
Master Server: false
Base Server: true

```

```

Base Service: mdm (Up, /mnt/stor01)
    ls (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
    cs (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64)
Public Address: 192.168.0.67:3260
Cluster Address: 192.168.0.67
Recent Start Time: 2026-01-04 15:15:03
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.9,"maxMemorySize":9223372036854775807}
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| No. | Path       | Storage Pool | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail | Data Service |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.  | /mnt/stor01 | pool1       | 6.79 GiB     | 92.86 GiB     | 6.69 GiB          | Unlimited      | Healthy      | -            | ds-1        |
| 2.  | /mnt/stor02 | pool2       | 587.02 MiB   | 92.86 GiB     | 28 KiB            | Unlimited      | Healthy      | -            | ds-2        |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Ports:
+-----+-----+
| Port Classification | Port |
+-----+-----+
| Port Range         | 19000-25000 |
| iSCSI Port         | 3260 |
| API Port            | 1443 |
| Web Port            | 2443 |
| Data Port           | 20007 |
| Storage Port        | 20004,20005,20011,20017,20018,20019 |
| Management Port     | 20000,20003,20010,20013,20014,20015,20016 |
| Metadata Port       | 20001,20002,20006,20008,20009,20012,20020,20021,20022 |
+-----+-----+
    
```

Table2: Description of the specified server information

Item	Description
Default Properties	Default HBlock parameters for servers that will join the cluster subsequently.
Server Properties for hblock_x	Information of the specific server.
Server Name	Server name.
Server ID	Server ID.
Node Name	Node name.
Parent Node	The name of a parent node.
Description	Description of the server. If no description is provided, no information is displayed in this field.
Status	Server status: <ul style="list-style-type: none"> ● Connected: The server is connected. ● Disconnected: The server is disconnected. ● Removing: The server is being removed from the cluster.
Master Server	Whether the server is a Master node (only supported by cluster mode): <ul style="list-style-type: none"> ● true: The server is the Master node. ● false: The server is not a Master node.
Base Server	Whether the server is a base node (only supported by cluster mode): <ul style="list-style-type: none"> ● true: The server is a base node. ● false: The server is not a base node.

Base Service	The base services on the server (only supported by cluster mode), including the basic service name, status, and the corresponding disk path.
Public Address	Public network IP and port.
Cluster Address	Cluster network IP and port.
Target portal IP and Port	iSCSI target portal and port.
Recent Start Time	The last successful startup time of the HBlock service on this node. -: Indicates that the HBlock service is in a stopped status.
Version	HBlock version number.
Parameters	Memory information available for HBlock on the server: <ul style="list-style-type: none"> ● maxMemoryRatio: The upper limit of the proportion of system total memory that HBlock can use. ● maxMemorySize: The maximum memory available for HBlock, in bytes.
Disk Paths	Disk path information: <ul style="list-style-type: none"> ● No.: Number. ● Path: Disk path. For the standalone mode, if there is (*) after the disk path, it indicates default disk path of the HBlock server. ● Storage Pool: The storage pool to which the disk path belongs. ● Used Capacity: The used capacity of the disk where the disk path resides. ● Total Capacity: The total capacity of the disk where the disk path resides. ● Used Capacity Quota: The total amount of data that has be written to the disk path by HBlock. ● Capacity Quota: The capacity quota of the disk path, that is, for each disk path, the total amount of data that can be written by HBlock. Once the space used by HBlock reaches the quota, data writing will be blocked immediately, and space exceeding the quota is not allowed to be used. ● Health Status: The health status of the disk path: <ul style="list-style-type: none"> ■ Healthy: The disk path is in a healthy status and can be read and written normally, and the disk usage of the disk path does not exceed the threshold (the system default value is 95%). ■ Warning: The disk path is warning and is readable, but one or more of the following situations exist: slow disk, the utilization rate of the disk where the disk path resides exceeds the threshold (the system default value is 95%), the remaining disk

	<p>space is less than 1GiB, HBlock stops writing to this path, the capacity quota usage of the disk path exceeds the threshold (the system default value is 95%), or the capacity quota of the disk path is set to 0.</p> <ul style="list-style-type: none"> ■ Error: The disk path is in an error status and cannot be accessed. The reasons may be that an I/O error occurs on the disk, resulting in the inability to read or write, or the disk path is not mounted correctly, etc. ● Health Detail: <ul style="list-style-type: none"> ■ If the health status is Healthy, this column is empty. ■ If the health status is Warning or Error, detailed information about the warning or error is displayed. ● Data Service: The DS process corresponding to the disk path, only supported by cluster mode. <p>Note: The DS process is displayed only after the disk path has been added to the storage pool.</p>
Ports	<p>Interfaces used by HBlock in the server:</p> <ul style="list-style-type: none"> ● Port Classification. ● Port: Specific port number.

- Cluster Mode: There are three servers in the cluster. Remove a disk path of hblock_1 and query the server information of hblock_1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls -n hblock_1
Server Name: pm-006
Server ID: hblock_1
Node Name: default:room2:rack2:hblock_1
Parent Node: default:room2:rack2
Status: Connected
Master Server: true
Base Server: true
Base Service: ls (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0)
               cs (Up, /mnt/storage01/CTYUN_HBlock_Plus_4.0.0)
Public Address: 192.168.0.192:3260
Cluster Address: 192.168.0.192
Recent Start Time: 2025-01-13 13:46:11
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.9,"maxMemorySize":9223372036854775807}
Disk Path(s):
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| No. | Path                | Storage Pool | Used Capacity | Total Capacity | Used Capacity Quota | Capacity Quota | Health Status | Health Detail | Data Service |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.  | /mnt/stor           | default     | 10.31 GiB    | 93.29 GiB     | 869.26 MiB        | Unlimited      | Healthy      | -             | ds-1         |
| 2.  | (Removing)/tmp/pathCV9V | default     | 62.1 MiB     | 5 GiB         | 62.09 MiB         | Unlimited      | Warning      | StopWritingWhenRemoving | ds-2         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Removing Details (from storage pool):
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| No. | Path                | Stage      | Details                                             |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.  | /tmp/pathCV9V      | CheckingData | FaultDomains: 2 healthy, 1 warning, 0 error
|      |                    |             | Data: 33.33% safe, 66.67% await reconstruction, 0% await more faultdomains, 0% single-copy, 0% corrupted |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

- Cluster Mode: There are four servers in the cluster, remove server hblock_4.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls -n hblock_4
Server Name: ecs-9689-0915142
Server ID: hblock_4
Node Name: default:room2:rack2:stor
```

```

Parent Node: default:room2:rack2
Status: Removing
Master Server: false
Base Server: false
Public Address: 192.168.0.117:20050
Cluster Address: 192.168.0.117
Recent Start Time: 2024-06-26 11:30:04
Version: 4.0.0
Parameters: {"maxMemoryRatio":0.9,"maxMemorySize":9223372036854775807}
Disk Path(s):

```

No.	Path	Storage Pool	Used Capacity	Total Capacity	Used Capacity Quota	Capacity Quota	Health Status	Health Detail	Data Service
1.	(Removing)/mnt/storage02	-	52.13 MiB	52.18 GiB	-	Unlimited	-	-	ds-1

Table 3: The description of the specified server information

Item	Description
Server Name	Server name.
Server ID	Server ID.
Node Name	Node name.
Parent Node	The name of a parent node.
Description	Description of the server. If no description is provided, no information is displayed in this field.
Status	Server status: <ul style="list-style-type: none"> ● Connected: The server is connected. ● Disconnected: The server is disconnected. ● Removing: The server is removing from the cluster.
Master Server	Whether the server is a Master node (supported in cluster mode only): <ul style="list-style-type: none"> ● true: The server is the Master node. ● false: The server is not a Master node.
Base Server	Whether the server is a base node (supported in cluster mode only): <ul style="list-style-type: none"> ● true: The server is a base node. ● false: The server is not a base node.
Base Service	The base services on the server (only supported by cluster mode), including the basic service name, status, and the corresponding disk path.
Public Address	Public network IP and port.
Cluster Address	Cluster network IP and port.
Target portal IP and Port	iSCSI target portal and port.
Recent Start Time	The last successful startup time of the HBlock service on this node. -: Indicates that the HBlock service is in a stopped status.
Version	HBlock version number.
Parameters	Memory information available for HBlock on the server: <ul style="list-style-type: none"> ● maxMemoryRatio: The upper limit of the proportion of system total memory that HBlock can use. ● maxMemorySize: The maximum memory available for HBlock, in bytes.
Disk Paths	Disk path information: <ul style="list-style-type: none"> ● No: number.

	<ul style="list-style-type: none"> ● Path: Disk path. For the standalone mode, if there is (*) after the disk path, it indicates default disk path of the HBlock server. ● Storage pool: The storage pool that the disk path is affiliated with, only supported by cluster mode. ● Used Capacity: The used capacity of the disk where the disk path resides. ● Total Capacity: The total capacity of the disk where the disk path resides. ● Used Capacity Quota: The total amount of data that has be written to the disk path by HBlock. ● Capacity Quota: The capacity quota of the disk path, that is, for each disk path, the total amount of data that can be written by HBlock. Once the space used by HBlock reaches the quota, data writing will be blocked immediately, and space exceeding the quota is not allowed to be used. ● Health Status: The health status of the disk path: <ul style="list-style-type: none"> ■ Healthy: The disk path is in a healthy status and can be read and written normally, and the disk usage of the disk path does not exceed the threshold (the system default value is 95%). ■ Warning: The disk path is warning and is readable, but one or more of the following situations exist: slow disk, the utilization rate of the disk where the disk path resides exceeds the threshold (the system default value is 95%), the remaining disk space is less than 1GiB, HBlock stops writing to this path, the capacity quota usage of the disk path exceeds the threshold (the system default value is 95%), or the capacity quota of the disk path is set to 0. ■ Error: The disk path is in an error status and cannot be accessed. The reasons may be that an I/O error occurs on the disk, resulting in the inability to read or write, or the disk path is not mounted correctly, etc. ● Health Detail: <ul style="list-style-type: none"> ■ If the health status is Heathy, this column is empty. ■ If the health status is Warning or Error, detailed information about the warning or error is displayed. ● Data Service: The DS process corresponding to the disk path, only supported by cluster mode. Note: The DS process is displayed only after the disk path has been added to the storage pool.
Removing Details	<p>When removing a server or removing a disk path, the details of the removed disk path:</p> <ul style="list-style-type: none"> ● No: number. ● Path: Specific disk path. ● Stage: The stage of the disk removed:

	<ul style="list-style-type: none"> ■ Reconfiguration ■ CheckingData ■ Executing ● Details: Detailed data: <ul style="list-style-type: none"> ■ FaultDomains: Fault domain details, including healthy (number of healthy ones), warning (number of warnings), and error (number of errors). ■ Data: safe (percentage of safe data), await reconstruction (percentage of data that needs to be reconstructed), await more faultdomains (percentage of data that requires additional fault domains to be reconstructed), single-copy (percentage of single-copy data), corrupted (corrupted data percentage).
--	---

- Standalone Mode: Query the port number used by HBlock in the server.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls --port
```

No.	Server ID	Server Name	Public Address	Cluster Address	Port Range	iSCSI Port	API&Web Port	Management Port
1.	hblock_1	hblockserver	192.168.0.32	192.168.0.32	20000-20500	3260	1443,2443	20000,20001,20002 20003,20004

- Cluster Mode: Query the port number used by HBlock in the server.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls --port
```

No.	Server ID	Server Name	Public Address	Cluster Address	Port Range	iSCSI Port	API&Web Port	Data Port	Storage Port	Management Port	Metadata Port
1.	hblock_1 (M**)	songt-0004	192.168.0.65	192.168.0.65	20000-20500	3260	1443,2443	20007	20017,20018,20019	20000,20003,20010 20013,20014,20015 20016	20001,20002,20004 20005,20006,20008 20009,20011,20012
2.	hblock_2 (**)	hblockserver	192.168.0.64	192.168.0.64	20000-20500	3260	1443,2443	20007	20017,20018,20019	20000,20003,20010 20013,20014,20015 20016	20001,20002,20004 20005,20006,20008 20009,20011,20012
3.	hblock_3 (**)	songt-0006	192.168.0.67	192.168.0.67	20000-20500	3260	1443,2443	20007	20004,20005,20011	20000,20003,20010 20013,20014,20015 20016	20001,20002,20006 20008,20009,20012

- Cluster Mode: Query the port number used by HBlock in the server hblock_1.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor server ls -n hblock_1 --port
```

```
Server Name: songt-0004
Server ID: hblock_1
Node Name: default:hblock_1
Parent Node: default
Public Address: 192.168.0.65
Cluster Address: 192.168.0.65
Ports:
```

Port Classification	Port
Port Range	20000-20500
iSCSI Port	3260
API Port	1443

Web Port	2443	
Data Port	20007	
Storage Port	20017,20018,20019	
Management Port	20000,20003,20010,20013,20014,20015,20016	
Metadata Port	20001,20002,20004,20005,20006,20008,20009,20011,20012	
-----+	-----+	-----+

Table 4: The description of the specified server port number

Item	Description
No.	Number.
Server Name	Server name.
Server ID	Server ID.
Node Name	Node name.
Parent Node	The name of a parent node.
Public Address	Public network IP and port.
Cluster Address	Cluster network IP and port.
Port Range	Port range.
iSCSI Port	iSCSI port.
API Port	API port.
Web Port	Web port.
Data Port	Data port (only supported by cluster mode).
Storage Port	Data storage port (only supported by cluster mode).
Management Port	Management service port.
Metadata Port	Metadata port (only supported by cluster mode).

4.14.9 Add Disk Paths

```
./stor server { A | addpath } { -p | --path } PATH &<1-n> [ { -n | --server } SERVER_ID ]
[ --capacity-quota CAPACITY_QUOTA ]
```

This command is used to add disk paths to the specified server of HBlock. Multiple disk paths can be added at one time, separated by commas (,).

Note:

- For the newly added disk path, it is recommended to set up automatic mounting at startup, or use a path or sub-path that has been set up to be automatically mounted.
- Each HBlock server can only add up to 100 disk paths.

Parameters

Parameter	Description
<code>-p PATH &<1-n></code> or <code>--path PATH &<1-n></code>	Specifies the disk path. Multiple disk paths can be added at one time, separated by comma (,). Note: This disk path is used to store data and it is not recommended to share a disk or file system with the operating system. Value: It can only contain letters, numbers, Chinese characters, or the special characters ~ ! @ \$ () _ + - ; : .
<code>-n SERVER_ID</code> or <code>--server SERVER_ID</code>	HBlock server ID of the disk path to be added. If the server ID is not specified, the disk path is added for the server executing the command.
<code>--capacity-quota CAPACITY_QUOTA</code>	Specifies the capacity quota of the disk path, that is, for each disk path added to server, the total amount of data that can be written by HBlock. Once the space used by HBlock reaches the quota, data writing is immediately blocked, and no more space is allowed to be used beyond the quota. Value: Integer. The value is smaller than the total capacity of the disk where the disk path resides. The units are K/k, M/m, G/g, T/t, and P/p. The default unit is G/g. A negative integer means unlimited writing and 0 means write inhibit. Writes are not limited by default. Note: If multiple disk paths are added at one time, only one capacity quota can be configured, and the configured capacity quota applies to all added disk paths.

Examples

Add the disk path /mnt/storage01 to the HBlock server hblock_1.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor server A -p /mnt/storage01 -n hblock_1 --capacity-quota 20G
```

Added path /mnt/storage01 on server hblock_1 successfully.

4.14.10 Modify Capacity Quota of Disk Path

```
./stor server { S | setpath } { -p | --path } PATH &<1-n> [ { -n | --server } SERVER_ID ]
--capacity-quota CAPACITY_QUOTA
```

This command is used to modify the capacity quota of the disk path.

Note: For the newly added disk path, it is recommended to set up automatic mounting at startup, or use a path or sub-path that has been set up to be automatically mounted.

Parameters

Parameter	Description
<code>-p PATH &<1-n></code> or <code>--path PATH &<1-n></code>	Specifies the disk path. The disk path cannot contain commas (.). Multiple disk paths can be added at one time, separated by comma (.). This disk path is used to store data and it is not recommended to share a disk or file system with the operating system.
<code>-n SERVER_ID</code> or <code>--server SERVER_ID</code>	HBlock Server ID of the disk path to be modified. If the server ID is not specified, the disk path is modified for the server executing the command.
<code>--capacity-quota CAPACITY_QUOTA</code>	Specifies the capacity quota of the disk path, that is, for each disk path added to server, the total amount of data that can be written by HBlock. Once the space used by HBlock reaches the quota, data writing is immediately blocked, and no more space is allowed to be used beyond the quota. Value: Integer. The value is smaller than the total capacity of the disk where the disk path resides. The units are K/k, M/m, G/g, T/t, and P/p. The default unit is G/g. A negative integer means unlimited writing and 0 means write inhibit. Writes are not limited by default. Note: If multiple disk paths are modified at one time, only one capacity quota can be configured, and the configured capacity quota applies to all specified disk paths.

Examples

Modify the capacity quota of the disk path /mnt/storage01 for the HBlock server hblock_1.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor server S -p /mnt/storage01 -n
hblock_1 --capacity-quota 80G
Set path /mnt/storage01 on server hblock_1 successfully.
```

4.14.11 Remove a Disk Path

```
./stor server { R | rmpath } { -p | --path } PATH &<1-n> [ { -n | --server } SERVER_ID ]
[ { -d | --del-data } ] [ -f | --force ]
```

This command is used to remove the disk path for the HBlock server. Multiple disk paths can be removed at one time.

Note:

- For the standalone mode, specify another disk path as default before removing a default disk path.
- If the collected logs are saved in the disk path of HBlock after log collection is executed, and you choose to delete the HBlock data in the disk path when the disk path is removed, the log will be deleted.
- While a disk path is being removed, other disk paths cannot be removed. If it must be removed, please use force removal, but there is a risk of data loss.
- For the cluster mode, the disk path can be removed only if it does not belong to any storage pool. Otherwise, the disk path cannot be removed. If you must remove the disk path, forcibly remove it, which may cause data loss.

Parameters

Parameter	Description
<code>-p PATH&<1-n></code> or <code>--path PATH &<1-n></code>	Specifies HBlock server disk path, multiple disk paths can be removed at one time, separated by commas (,).
<code>-n SERVER_ID</code> or <code>--server SERVER_ID</code>	Specifies the HBlock server ID of the removed disk path. If no HBlock server ID is specified, the disk path is removed for the server on which the command was executed.
<code>-d</code> or <code>--del-data</code>	Delete the HBlock data on the disk path when the disk path is removed.
<code>-f</code> or <code>--force</code>	Forcibly remove the disk path. Note: Forcibly removing the disk path may cause risk of data loss.

Examples

Remove the disk path /home/stor01 for HBlock server hblock_2.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor server R -p /home/stor01 -n
hblock_2
Start removing path /home/stor01.
```

4.15 Query HBlock Information

```
./stor info [ { -i | --stor-id } | { -n | --stor-name } | { -u | --user-name } | { -S | --serial-id } | { -T | --trial } ]
```

This command is used to query HBlock information.

Parameters

Parameter	Description
<code>-i</code> or <code>--stor-id</code>	Query information about the HBlock ID.
<code>-n</code> or <code>--stor-name</code>	Query information about the HBlock name.
<code>-u</code> or <code>--user-name</code>	Query information about the HBlock administrator username.
<code>-S</code> or <code>--serial-id</code>	Query information about the serial number of HBlock. This serial number is required to apply for a license.
<code>-T</code> or <code>--trial</code>	Query information about HBlock trial period.

Examples

- Standalone Mode - Commercial Edition: Query HBlock information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info
HBlock ID: 0A769770-1528-49B8-A82D-D1CB71333FFB
HBlock name: storuser
HBlock serial ID: 0A769770-1528-49B8-A82D-D1CB71333FFB-0201-040000
User name: storuser
Setup date: 2025-12-09 13:43:43
License ID: qws2b6a9-f3fb-4098-a6b3-3652a5a76530
License type: Perpetual
License status: Effective
License maintenance expire time: 2026-12-26 11:25:05
```

- Standalone Mode - Free Edition (Pro trial enabled): Query HBlock information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info
HBlock ID: C3238159-40D5-4E6E-A719-1A5205BB688D
HBlock name: stor1
HBlock serial ID: C3238159-40D5-4E6E-A719-1A5205BB688D-0201-040000
User name: storuser
Setup date: 2026-01-12 10:39:27
You are using a free edition with basic features to get started. The deadline for the upgrade support is 2028-01-12 10:39:27. Additionally, the Pro trial (started on 2026-01-12 11:29:43) ends on 2026-02-11 11:29:43. To ensure business continuity and access the full product experience, please contact your software vendor to get a license.
```

- Cluster Mode - Commercial Edition: Query HBlock information.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info
HBlock ID: 1E46239F-BD8B-4CB3-A151-C3A58892AEF9
```

```

HBlock name: stor1
HBlock serial ID: 1E46239F-BD8B-4CB3-A151-C3A58892AEF9-0202-040000
User name: storuser
Base Pool: default
Fault domain of base pool: server
Setup date: 2025-12-18 09:33:52
License ID: ehc2b6a9-f3fb-4098-a6b3-3652a5d71232
License type: Subscription
License status: Effective
License expire time: 2026-12-26 11:24:30
    
```

● Cluster Mode - Free Edition (Pro trial not enabled): Query HBlock information.

```

[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info
HBlock ID: 10266612-3B09-4111-968A-CBB4D271F358
HBlock name: stor2
HBlock serial ID: 10266612-3B09-4111-968A-CBB4D271F358-0202-040000
User name: storuser
Base Pool: default
Fault domain of base pool: server
Setup date: 2026-01-12 11:14:10
You are using a free edition with basic features to get started. The deadline for the
upgrade support is 2028-01-12 11:14:10. To ensure business continuity and access the full
product experience, please contact your software vendor to get a license.
    
```

Table 1: Description of HBlock information

Item	Description
HBlock ID	HBlock ID.
HBlock name	The name of HBlock.
HBlock serial ID	The serial number of HBlock.
User name	HBlock Administrator Username.
Base Pool	Base storage pool name, only supported by cluster mode.
Fault domain of base pool	Fault domain type of the base storage pool, only supported by cluster mode: <ul style="list-style-type: none"> ● path. ● server. ● rack. ● room.
Setup date	Initialized time of HBlock.
License ID	The ID of the last successfully imported software license (commercial edition).
License type	Types of software license subscriptions: <ul style="list-style-type: none"> ● Subscription ● Perpetual
License status	Software license status: <ul style="list-style-type: none"> ● Effective: The software license is effective.

	<ul style="list-style-type: none"> ● Expired: The software license has expired. ● Invalid: The software license is invalid.
License maintenance expire time	The maintenance expiration time of the perpetual license.
License expire time	The expiration time of the subscription license.

- Query HBlock ID.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info --stor-id
HBlock ID: CECD1458-7405-4825-BF6A-75C04F1114ED
```

- Query HBlock Name.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info --stor-name
HBlock name: test
```

- Query the Administrator Username of HBlock.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info --user-name
User name: storuser
```

- Query the Serial Number of HBlock.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info --serial-id
HBlock serial ID: CECD1458-7405-4825-BF6A-75C04F1114ED-00000003-00000001-00000015
```

- Query HBlock trial period information (software license has been imported)

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info --trial
It is not trial version. You can run "stor license ls" to view the license information
```

- Query HBlock trial period information (software license is not imported)

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor info --trial
You are using a 30-day trial version, the expiration date is 2024-08-06. After it expires, the management function will not be able to use. Please contact your software vendor to get a license.
```

4.16 Query HBlock Service Status

```
./stor status [ { -t | --type } TYPE ]
```

This command is used to query the HBlock service status.

Parameters

Parameter	Description
-t <i>TYPE</i> or --type <i>TYPE</i>	View the status of the specified service. The values of <i>TYPE</i> are as follows: <ul style="list-style-type: none"> ● system: Query HBlock system status. ● storagepool: Query storage pool status, only supported by cluster mode. ● server: Query HBlock server status. ● disk: Query HBlock disk status, only supported by cluster mode. ● LUN: Query HBlock LUN status. If <i>TYPE</i> is not specified, all information is displayed.

Examples

- Standalone Mode: Query the HBlock service status.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor status
System:
  status: Working, licenseStatus: Effective
Server:
  status: 1 total, 1 connected, 0 disconnected, 0 removing
Disk:
  usage: 60.15 MiB used, 93.29 GiB total
  status: 1 total, 1 healthy, 0 warning, 0 error
LUN:
  status: 3 total, 230 GiB (cache mode:0 total, 0 B; storage mode:1 total, 110 GiB;
local mode:2 total, 120 GiB)
  data: 100% normal, 0% low redundancy, 0% error
```

- Cluster Mode: Query the HBlock service status.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor status
System:
  status: Working, licenseStatus: Effective
Server:
  status: 3 total, 3 connected, 0 disconnected, 0 removing
LUN:
  status: 10 total, 1.18 TiB (cache mode:1 total, 210 GiB; storage mode:3 total, 312
GiB; local mode:6 total, 691 GiB)
  data: 100% normal, 0% low redundancy, 0% error
Storage Pool:
```

```

number: 2
base pool: default
fault domain of base pool:
  level: server
  status: 3 total, 3 healthy, 0 warning, 0 error
data health of base pool: 100% normal, 0% low redundancy, 0% error
disk of base pool:
  usage: 1.83 GiB used, 278.57 GiB total
  status: 3 total, 3 healthy, 0 warning, 0 error
  
```

- Query the HBlock server status.

```

[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor status -t server
Server:
  status: 3 total, 2 connected, 1 disconnected, 0 removing
  
```

Table: Description of HBlock Service Status Information

Item	Description
System	System status, including: <ul style="list-style-type: none"> ● status: System running status. ● licenseStatus: Software license status. <ul style="list-style-type: none"> ■ Effective: The software license is effective. ■ Expired: The software license has expired. ■ Invalid: The software license is invalid. ■ None: Software license not yet loaded.
Server	Server status, including number of servers, the number of servers whose status is connected, the number of servers whose status is disconnected, and the number of servers whose status is being removed.
Disk	Disk status (only supported by standalone mode), including: <ul style="list-style-type: none"> ● usage: The usage of the disk where the disk path resides, including the used capacity of the disk and the total capacity of the disk. ● status: The status of the disk where the disk path resides, including the total number of disks, the number of disks with healthy status, the number of disks with warning status, and the number of disks with error status.
LUN	The status of the LUN, including: <ul style="list-style-type: none"> ● status: Number of LUNs, total capacity, and number and capacity of each mode's LUN. ● data: The data health status of LUNs, including: the percentage of normal data (normal), the percentage of low redundancy data (low redundancy), and the percentage of erroneous data (error). If low redundancy data exists, low redundancy reconstruction progress will be provided.
Storage Pool	Storage pool status, only supported by cluster mode: <ul style="list-style-type: none"> ● number: The number of storage pools. ● base pool: The name of the base storage pool.

	<ul style="list-style-type: none">● fault domain of basis pool: The fault domain of a base storage pool:<ul style="list-style-type: none">■ level: The level of a fault domain:<ul style="list-style-type: none">◆ path.◆ server.◆ rack.◆ room.■ status: The status of a fault domain.<ul style="list-style-type: none">◆ Healthy: Indicates the health status, can be read and written normally.◆ Warning: Indicates the warning status, can be read.◆ Error: Indicates the error status, cannot be accessed.● Data health of base pool: The data health status of the base storage pool, including: the percentage of normal data (normal), the percentage of low redundancy data (low redundancy), and the percentage of erroneous data (error). If low redundancy data exists, low redundancy reconstruction progress will be provided.● disk of base pool: Indicates the disk path status of the base storage pool.<ul style="list-style-type: none">■ usage: The usage of the disk where the disk path resides, including the used capacity of the disk and the total capacity of the disk.■ status: The status of the disk where the disk path resides, including the total number of disks, the number of disks with healthy status, the number of disks with warning status, and the number of disks with error status.
--	---

4.17 Monitoring

4.17.1 View Real-time Performance Data

```
./stor monitor { V | view } [ { -d | --dimension } DIMENSION ] [ { -i | --instance }
INSTANCE_ID <1-n> ] [ { -m | --metric } METRIC <1-n> ]
```

This command is used to view HBlock's real-time performance data.

Note: Monitoring data is recorded based on the server system time. The time is adjusted, or the server time in the cluster is not synchronized, which may lead to inaccurate monitoring data.

However, users' business data will not be affected.

Parameters

Parameter	Description
-d <i>DIMENSION</i> or --dimension <i>DIMENSION</i>	Specifies the monitoring object. Value: <ul style="list-style-type: none"> ● system. ● pool (only supported by cluster mode). ● server. ● disk (diskPath). ● LUN. The default value is <code>system</code> for standalone mode, the default value is <code>pool</code> for cluster mode.
-i <i>INSTANCE_ID</i> or --instance <i>INSTANCE_ID</i>	Specifies the unique identifier of the monitoring object instance. Multiple instances can be filled in each time, separated by commas (,): <ul style="list-style-type: none"> ● Monitoring object is <code>system</code> and there is no instance. ● Monitoring object is <code>pool</code>, and the instance value is the storage pool name. ● Monitoring object is <code>server</code>, and the instance value is the server ID. ● Monitoring object is <code>disk (diskPath)</code>, and the instance value is <code>serverId:/diskPath</code>. ● Monitoring object is a LUN, and the instance value is the LUN name. If no <i>INSTANCE_ID</i> is specified, all instances of the monitoring object are queried by default.
-m <i>METRIC <1-n></i> or --metric <i>METRIC <1-n></i>	Specifies the metric name. <ul style="list-style-type: none"> ● Monitoring object is <code>system</code>, and the values are <code>PATH</code>, <code>IOPS</code>, <code>BANDWIDTH</code>, <code>LATENCY</code>, and <code>CLOUD</code>.

- Monitoring object is pool, and the values are PATH, IPOS, BANDWIDTH and LATENCY.
 - Monitoring object is server, and the values are MEM, CPU, PATH, IOPS, BANDWIDTH, LATENCY, and CLOUD.
 - Monitoring object is disk (diskPath), and the value is PATH.
 - Monitoring object is LUN, and its values are IOPS, BANDWIDTH, LATENCY, and CLOUD.
- By default, all metric data of the monitoring object is queried.

Examples

- View the system's real-time performance data.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor monitor V -d system
```

No.	Time	IOPS (T/W/R)	Bandwidth (T/W/R)	Latency (T/W/R)	Path (Total/Used/Rate)	Quota (Total/Used/Rate)	Cloud Bandwidth (T/U/D)
1	2025-08-12 14:54:00	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	278.57 GiB/ 48.69 GiB/ 17.48%	278.57 GiB/ 19.85 GiB/ 7.12%	-/ -/ -

- View the pool's real-time performance data.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor monitor V -d pool
```

No.	Time	Pool Name	IOPS (T/W/R)	Bandwidth (T/W/R)	Latency (T/W/R)	Path (Total/Used/Rate)	Quota (Total/Used/Rate)
1	2025-08-12 14:55:20	default (*)	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	278.57 GiB/ 48.69 GiB/ 17.48%	278.57 GiB/ 19.84 GiB/ 7.12%

- View the server's real-time performance data.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor monitor V -d server
```

No.	Time	Server ID	CPU Rate	Memory (Total/Used/Rate)	IOPS (T/W/R)	Bandwidth (T/W/R)	Latency (T/W/R)	Path (Total/Used/Rate)	Quota (Total/Used/Rate)	Cloud Bandwidth (T/U/D)
1	2025-08-12 14:56:20	hblock_1	5.51%	15.12 GiB/ 14.89 GiB/ 98.46%	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	92.86 GiB/ 6.63 GiB/ 7.14%	92.86 GiB/ 6.61 GiB/ 7.12%	-/ -/ -
2	2025-08-12 14:56:20	hblock_2	3.73%	15.12 GiB/ 14.71 GiB/ 97.26%	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	92.86 GiB/ 35.44 GiB/ 38.16%	92.86 GiB/ 6.61 GiB/ 7.12%	-/ -/ -
3	2025-08-12 14:56:20	hblock_3	4.42%	15.12 GiB/ 14.82 GiB/ 97.97%	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	92.86 GiB/ 6.62 GiB/ 7.13%	92.86 GiB/ 6.62 GiB/ 7.12%	-/ -/ -

- View the disk path's real-time performance data.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor monitor V -d disk
```

No.	Time	DiskPath	Total Path Capacity	Used Path Capacity	Path Rate	Path Quota Capacity	Used Path Quota	Quota Rate
1	2026-01-20 13:40:00	hblock_1:/mnt/stor02	92.86 GiB	7.76 GiB	8.36%	90 GiB	1.03 GiB	1.14%
2	2026-01-20 13:40:00	hblock_1:/mnt/storage01	92.86 GiB	3.05 GiB	3.29%	92.86 GiB	32 KiB	0%
3	2026-01-20 13:40:00	hblock_2:/mnt/stor01	92.86 GiB	8.4 GiB	9.04%	90 GiB	1.03 GiB	1.14%
4	2026-01-20 13:40:00	hblock_3:/mnt/stor01	92.86 GiB	7.82 GiB	8.43%	90 GiB	1.03 GiB	1.14%

- View the LUN's real-time performance data.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor monitor V -d LUN
```

No.	Time	LUN Name	IOPS (T/W/R)	Bandwidth (T/W/R)	Latency (T/W/R)	Cloud Bandwidth (T/U/D)	Wait Upload
1	2025-08-12 15:01:20	Clone-2	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	-/ -/ -	-
2	2025-08-12 15:01:20	lun001	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	-/ -/ -	-

3	2025-08-12 15:01:20	lun001-c001	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	-/ -/ -	-
4	2025-08-12 15:01:20	lun01a	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	-/ -/ -	-
5	2025-08-12 15:01:20	lun01a-clone2	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	-/ -/ -	-
6	2025-08-12 15:01:20	lun04	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	-/ -/ -	-/ -/ -	-
7	2025-08-12 15:01:20	lunc	0/ 0/ 0	0 B/s/ 0 B/s/ 0 B/s	1.34 ms/ 1.34 ms/ -	0 B/s/ 0 B/s/ 0 B/s	537.08 MiB

Table: Description of real-time performance data

Item	Description
No.	Number.
Time	The time when querying the performance data.
IOPS(T/W/R)	The total number of reads and writes per second, the number of writes per second, and the number of reads per second.
Bandwidth(T/W/R)	Total bandwidth, write bandwidth, read bandwidth.
Latency(T/W/R)	Total latency, write latency, read latency.
Path(Total/Used/Rate)	The total capacity, used capacity, and average usage rate of the disk where the disk path resides.
Quota(Total/Used/Rate)	HBlock total capacity quota, used capacity quota, capacity quota usage rate.
Cloud Bandwidth(T/U/D)	Cloud bandwidth: total bandwidth, upload bandwidth, down bandwidth.
Pool Name	Pool name.
Server ID	Server ID.
CPU Rate	The CPU usage of the server.
Memory(Total/Used/Rate)	Total server memory, used memory, memory usage rate.
DiskPath	Disk path.
Total Path Capacity	The total capacity of the disk where the disk path resides.
Used Path Capacity	Used capacity of the disk where the disk path resides.
Path Quota Capacity	The capacity quota of HBlock.
Used Path Quota	Used capacity of HBlock.
Quota Rate	The capacity quota usage rate of HBlock.
LUN Name	LUN name.
Wait Upload	The amount of data to be uploaded to the cloud.

4.17.2 Export Performance Data

```
./stor monitor { E | export } [ { -d | --dimension } DIMENSION ] [ { -i | --instance }
INSTANCE_ID ] [ --start-time START_TIME ] [ --end-time END_TIME ] [ { -o | --out }
DIRECTORY ]
```

This command is used to export HBlock performance data.

Parameters

Parameter	Description
-d <i>DIMENSION</i> or --dimension <i>DIMENSION</i>	Specifies the monitoring object. Value: <ul style="list-style-type: none"> ● system. ● pool (only supported by cluster mode). ● server. ● disk (diskPath). ● LUN. The default value is system for standalone mode, the default value is pool for cluster mode.
-i <i>INSTANCE_ID</i> or --instance <i>INSTANCE_ID</i>	Specifies the unique identifier of the monitoring object instance: <ul style="list-style-type: none"> ● Monitoring object is system and there is no instance. ● Monitoring object is server, the item is required, and the instance value is the server ID. ● Monitoring object is pool, the item is required, and the instance value is the storage pool name. ● Monitoring object is disk (diskPath), the item is required, and the instance value is <i>serverId:diskPath</i>. ● Monitoring object is LUN, the item is required, and the instance value is the LUN name.
--start-time <i>START_TIME</i>	The starting time for performance data to be exported. The format is <i>yyyy-MM-dd HH:mm:ss</i> . If not specified, the default is the time point 2 hours before the end time of performance data export. Note: The start time must be earlier than the end time, and the start time cannot be one year earlier than the current server time.
--end-time <i>END_TIME</i>	End time of performance data to be exported. The format is <i>yyyy-MM-dd HH:mm:ss</i> , and the default value is the current time. Note: The start time must be earlier than the end time.

<p><code>-o <i>DIIRECTORY</i></code> or <code>--out <i>DIIRECTORY</i></code></p>	<p>Specifies the directory where the exported file is stored, as an absolute path. By default, it is stored in the HBlock installation directory of the requested server and is named as <code>monitordata_<i>instanceId</i>_yyyyMMddHHmmss.csv</code>, where:</p> <ul style="list-style-type: none">● <i>instanceId</i>: The instance of the monitoring object. If not specified, it will be empty.● <i>yyyyMMddHHmmss</i>: The time when the file was generated.
--	---

Examples

Export the performance data of LUN lun01a from 2026-01-10 14:00 to 2026-02-23 14:00 to `/home/stor/monitor`.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor monitor E -d LUN -i lun01a --start-time 2026-01-10 14:00:00 --end-time 2026-02-23 14:00:00 --out /mnt/stor/monitor
Exported successfully.
```

4.18 Alarm

Note: The limit of the data in 'Unresolved' status is 10,000. If the limit is reached, new alarms cannot be displayed or notified. At that time, please fix the related problem, or try to manually resolve the alarm.

4.18.1 Query Alarm Information

```
./stor alarm ls [-S { Resolved | Expired } [--alarm-severity ALARM_SEVERITY ]
[ --num NUMBER ]
```

The command is used to query resolved alarm or expired alarm information.

```
./stor alarm ls [ -S Unresolved] [--alarm-severity ALARM_SEVERITY ] [--mute-
status MUTE_STATUS ] [ --num NUMBER ]
```

The command is used to query unresolved alarm information.

```
./stor alarm ls [ { -n | --alarm } ALARM_ID ]
```

The command is used to query the specified unresolved alarm information.

Note: The alarm data is recorded based on the system time of server. Adjusting server time, or inconsistent time of servers, may cause inaccurate alarm data. But the user's business data will not be affected.

Parameters

Parameter	Description
-S Resolved	Specifies that the alarm status is Resolved.
-S Expired	Specifies that the alarm status is Expired.
-S Unresolved	Specifies that the alarm status is Unresolved. If the alarm status is not specified, the default alarm status is Unresolved.
--alarm-severity <i>ALARM_SEVERITY</i>	Specifies the alarm level: <ul style="list-style-type: none"> ● Warning ● Major ● Critical By default, all levels of alarms are queried.
--num <i>NUMBER</i>	Specifies the number of the alarms to query. The value is an integer that ranges from 2 to 10000. If the number is not specified, a maximum of 10000 alarm records will be listed.
--mute-status <i>MUTE_STATUS</i>	Specifies the alarm mute status. <ul style="list-style-type: none"> ● Normal ● Muted The default value is Normal.

<p>-n <i>ALARM_ID</i> or --alarm <i>ALARM_ID</i></p>	<p>Specifies the alarm ID, and the alarm status is Unresolved.</p>
--	--

Examples

- Query alarm information of resolved alarms.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor alarm ls -S Resolved
```

Alarm ID	Instance ID	Alarm Rule	Severity	Alarm Time	Alarm Value	Resolve Time	Resolve Value	Duration	Resolve Type
4jlSnwld	hblock_3:/mnt/stor	CapacityQuotaUsageExceedsThreshold	Warning	2024-01-08 11:21:11	93.69%	2024-01-08 11:31:14	5.06%	10m 3s	Auto
4jlSnJ9h	hblock_3	CapacityQuotaUsageExceedsThreshold	Warning	2024-01-08 11:21:11	93.69%	2024-01-08 11:31:14	5.06%	10m 3s	Auto
4jlSmXuZ	hblock_3:/mnt/stor	FaultDomainWarning	Warning	2024-01-08 11:21:10	-	2024-01-08 11:31:14	-	10m 3s	Auto
4jlSmPPo	hblock_3:/mnt/stor	DataServiceHealthStatusWarning	Warning	2024-01-08 11:21:10	-	2024-01-08 11:31:13	-	10m 2s	Auto
4lZH7wX	hblock_1	ProtocolServiceAbnormal	Major	2024-01-05 17:16:37	-	2024-01-08 09:20:46	-	2d 16h 4m	Auto
4hDwJnaH	hblock_3	ProtocolServiceAbnormal	Major	2024-01-03 17:47:45	-	2024-01-03 17:53:47	-	6m 1s	Auto
4hD1BMWS	hblock_4	ProtocolServiceAbnormal	Major	2024-01-03 17:15:35	-	2024-01-03 17:53:46	-	38m 11s	Auto
4hD2AgSt	hblock_1:/mnt/stor	FaultDomainError	Major	2024-01-03 17:16:36	-	2024-01-03 17:17:37	-	1m 1s	Auto
4hD2zESY	hblock_1:/mnt/stor	DataServiceHealthStatusError	Major	2024-01-03 17:16:35	-	2024-01-03 17:17:36	-	1m 1s	Auto
4hCtBCK5	hblock_4:/mnt/storage01	DataServiceHealthStatusError	Major	2024-01-03 16:40:27	-	2024-01-03 16:41:27	-	1m 0s	Auto

- Query alarm information of unresolved alarms.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor alarm ls
```

Alarm ID	Instance ID	Alarm Rule	Severity	Alarm Time	Alarm Value	Current Value	Duration	Mute Status
4jlhtUel	hblock_4	ProtocolServiceAbnormal	Major	2024-01-08 10:43:03	-	-	20s	Normal

- Query the alarm information of unresolved alarms and the alarm level is Warning.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor alarm ls -S Unresolved --alarm-severity Warning
```

Alarm ID	Instance ID	Alarm Rule	Severity	Alarm Time	Alarm Value	Current Value	Duration	Mute Status
1kyvM3zJ	ehc2b6a9-f3fb-4098-a6b3-3652a5d76269	LicenseWillExpire	Warning	2022-08-16 17:26:02	-	-	16h 26m 41s	Normal

- Query the alarm whose alarm ID is 4jlhtUel.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor alarm ls -n 4jlhtUel
```

Alarm Id: 4jlhtUel
Instance Id: hblock_4
Instance Snapshot: hblock_4,ecs-9689-0915141,192.168.0.202
Severity: Major
Alarm Rule: ProtocolServiceAbnormal
Alarm Time: 2024-01-08 10:43:03
Alarm Value: -
Current Value: -
Duration: 2m 0s
Alarm Status: Unresolved
Mute Status: Normal
Mute Operations: -

Table: The description of alarm information

Item	Description
Alarm ID	Alarm ID.

Instance ID	Alarm instance ID.
Alarm Rule	Alarm rules. See the appendix Alarm List for detailed descriptions of alarm rules.
Severity	Alarm level: <ul style="list-style-type: none"> ● Warning. ● Major. ● Critical.
Alarm Time	Alarm occurrence time.
Alarm Value	<p>Alarm value. Only an alarm corresponding to a numerical indicator will display values. The following alarm rules involve numerical indicators:</p> <ul style="list-style-type: none"> ● AlarmNumberApproachingLimit: The number of unresolved alarms is close to the upper limit. Displays (total number of unresolved alarms /upper limit of number of unresolved alarms) as percentage. The unit is %. Calculation formula: (total number of unresolved alarms /upper limit of number of unresolved alarms)*100%. ● ResourceUsageApproachingLimit: The resource usage is close to the upper limit, showing the license used capacity usage as a percentage, the unit is %. Calculation formula: (total local LUN capacity/capacity allowed by the license)*100%. ● CapacityQuotaUsageExceedsThreshold: The quota usage exceeds the threshold, showing the quota usage (Path_Cap_Quota_Rate) of the disk associated with the storage pool or disk path, in percentage (%). ● CapacityQuotaUsageApproachLimit: The quota is exhausted, showing the quota usage rate (Path_Cap_Quota_Rate) of the disk associated with the storage pool or disk path, in percentage (%). ● DiskUsageExceedsThreshold: The disk usage exceeds the threshold, showing the usage (Path_Rate) of the disk associated with the storage pool or disk path, in percentage (%). ● InsufficientSpaceonInstallationPath: The remaining space of the installation directory is insufficient, showing the remaining space of the file system of the disk where the directory is located, in GiB. ● InsufficientSpaceonMetaDir: The remaining space in the disk path of the base service is insufficient. It shows the remaining space of the file system on the disk where the disk path resides, in GiB.
Current Value	<p>Current numerical value Only an alarm corresponding to a numerical indicator will display values. The following alarm rules involve numerical indicators:</p> <ul style="list-style-type: none"> ● AlarmNumberApproachingLimit: The number of unresolved alarms is close to the upper limit. Displays (total number of unresolved alarms /upper limit of number of unresolved alarms) as percentage. The unit is %. Calculation formula: (total number of unresolved alarms /upper limit of number of unresolved alarms)*100%. ● ResourceUsageApproachingLimit: The resource usage is close to the upper limit, showing the license used capacity usage as a percentage, the

	<p>unit is %. Calculation formula: (total local LUN capacity/capacity allowed by the license)*100%.</p> <ul style="list-style-type: none"> ● CapacityQuotaUsageExceedsThreshold: The quota usage exceeds the threshold, showing the quota usage (Path_Cap_Quota_Rate) of the disk associated with the storage pool or disk path, in percentage (%). ● CapacityQuotaUsageApproachLimit: The quota is exhausted, showing the quota usage rate (Path_Cap_Quota_Rate) of the disk associated with the storage pool or disk path, in percentage (%). ● DiskUsageExceedsThreshold: The disk usage exceeds the threshold, showing the usage (Path_Rate) of the disk associated with the storage pool or disk path, in percentage (%). ● InsufficientSpaceonInstallationPath: The remaining space of the installation directory is insufficient, showing the remaining space of the file system of the disk where the directory is located, in GiB. ● InsufficientSpaceonMetaDir: The remaining space in the disk path of the base service is insufficient. It shows the remaining space of the file system on the disk where the disk path resides, in GiB.
Duration	Alarm duration
Mute Status	Mute status: <ul style="list-style-type: none"> ● Muted. ● Normal.
Resolve Time	Alarm resolve time.
Resolve Value	Alarm resolve value.
Resolve Type	Alarm resolve type: <ul style="list-style-type: none"> ● Auto: Automatically resolve. ● Manual: Manually resolve.
Instance Snapshot	Alarm instance snapshot refers to the detailed information of the alarm instance when the alarm occurs.
Alarm Status	Alarm status: <ul style="list-style-type: none"> ● Resolved: ● Expired. ● Unresolved.
Mute Operations	Mute records include: <ul style="list-style-type: none"> ● operTime: Operation Time. ● operType: Mute operation type: <ul style="list-style-type: none"> ■ Mute. ■ AutoUnmute. ■ ManualUnmute. ● reason: Reason for muting/unmuting. ● dueTime: Mute due time.

4.18.2 Export Alarms

```
./stor alarm { export | E } [ { -S | --status } STATUS ] [ { -o | --out } DIRECTORY ]
```

This command is used to export alarms.

Parameters

Parameter	Description
-s <i>STATUS</i> or --status <i>STATUS</i>	Alarm status: <ul style="list-style-type: none"> ● Unresolved. ● Resolved: ● Expired. The default value is Unresolved.
-o <i>DIRECTORY</i> or --out <i>DIRECTORY</i>	The directory where exported alarms are stored is an absolute path. By default, it is stored in the HBlock installation directory of the requested server and is named <code>alarm_Status_yyyyMMddHHmmss.csv</code> , where: <ul style="list-style-type: none"> ● Status: Alarm status. ● <code>yyyymmddHHmmss</code>: The time when the file was generated.

Examples

Export the unresolved alarm records to `/home/stor/alarms`.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor alarm E -o /home/stor/alarms
Exported successfully.
```

4.18.3 Manually Resolve the Alarm

```
./stor alarm { resolve | R } { -n | --alarm } ALARM_ID
```

This command is used to manually resolve the alarm.

Note: When manually resolving the alarm, you need to enter the reason for resolving the alarm. The reason entered cannot exceed 50 characters. If the problem indicated by the alarm is not resolved, the system will send the alarm again.

Parameters

Parameter	Description
<code>-n ALARM_ID</code> or <code>--alarm ALARM_ID</code>	Alarm ID.

Examples

Manually resolve alarm 10DR4l86.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor alarm R -n 10DR4l86
Enter the troubleshooting instruction of this alarm record, limited to 50 characters:
Restart server.
Successfully set the alarm status to 'Resolved', make sure the related problem is fixed,
otherwise it will alarm again.
```

4.18.4 Mute Alarm

```
./stor alarm { M | mute } { -n | --alarm } ALARM_ID --due-time DUE_TIME
```

This command is used to mute the alarm.

Note: When a mute alarm occurs, you need to enter the reason for the mute alarm. The reason cannot exceed 50 characters.

Parameters

Parameter	Description
<code>-n ALARM_ID</code> or <code>--alarm ALARM_ID</code>	Alarm ID.
<code>--due-time DUE_TIME</code>	Mute due time. The format is yyyy-MM-dd HH:mm:ss.

Examples

Mute alarm 1kUyXPFy.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor alarm M -n 1kUyXPFy --due-time 2022-08-17 19:00:00
Enter the reason for this operation, limited to 50 characters:
I will apply new license.
Muted successfully.
```

4.18.5 Unmute Alarm

```
./stor alarm { UM | unmute } { -n | --alarm } ALARM_ID
```

This command is used to unmute the alarm.

Note: When unmuting the alarm, you need to enter the reason for unmuting the alarm and the due time for muting it. The reason entered cannot exceed 50 characters.

Parameters

Parameter	Description
<code>-n ALARM_ID</code> or <code>--alarm ALARM_ID</code>	Alarm ID.

Examples

Unmute alarm 1kUyXPFy mute.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor alarm unmute -n 1kUyXPFy
Enter the reason for this operation, limited to 50 characters:
No new license.
Unmuted successfully.
```

4.19 Event and Log Management

See the appendix **User Event List** and **System Event List** for event name descriptions.

Note: The system can retain events for 6 months.

4.19.1 Query HBlock Events

```
./stor event ls [ --type TYPE ] [ --num NUMBER ]
```

This command is used to query HBlock event information.

Parameters

Parameter	Description
<code>--type TYPE</code>	Specifies the type of HBlock event to query. Value: <ul style="list-style-type: none"> ● system: System Event. ● user: user event. The default value is user.
<code>--num NUMBER</code>	Specifies the number of queries for latest events. The value is an integer that ranges from 1 to 1000. When the value is 1, the detailed information of the latest event is displayed. When the value is greater than 1, the main information of the latest events is displayed.

Examples

- Query the latest user event information.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor event ls --num 1
event ID: f0ffa38c-7a8e-4253-9352-b19eb4af906c
module: Server
name: RemoveServer
event time: 2024-01-23 14:07:06
request ID: 1db0b648560144849ce63cfc59453d82
requester IP: 192.168.0.110
status code: 202
error code: -
error message: -
details: {method:DELETE,
URL:/rest/v1/system/server/hblock_4?force=false&deleteLocalData=false}
```

Table 1: The description of user event

Item	Description
event ID	User event ID.

moudle	User event module.
name	User event name. Note: For a description of the user event name, see User Event List .
event time	The time when HBlock received the user event request.
requester ID	Request ID of user event.
requester IP	The source IP address of the request.
status code	Response status code.
error code	Error Code. -: Indicates no error code.
error message	Error Message. -: Indicates no error information.
details	User event details.

- Query last five user events.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor event ls --type user --num 5
```

Event Time	Event ID	Requester IP	Module	Event Name	Status	Error Code
2024-01-23 14:07:06	f0ffa38c-7a8e-4253-9352-b19eb4af906c	192.168.0.110	Server	RemoveServer	202	-
2024-01-23 13:55:09	1f312b19-49e5-440f-a132-e1fe2cf99385	1.202.233.200	System	Login	200	-
2024-01-23 11:10:46	26f47555-b9e6-4a76-bdfe-cbd4a459b4aa	192.168.0.110	Server	SetPath	204	-
2024-01-23 11:07:09	a97ae9e8-9d92-429f-a947-c36ccf43fac8	192.168.0.110	Server	AddPath	200	-
2024-01-23 11:01:08	69f5f9f0-57a7-487d-b495-0391d4da8856	192.168.0.110	Server	SetServer	204	-

Table 2: The description of user event

Item	Description
Event Time	The time when HBlock received the user event request.
Event ID	User event ID.
Requester IP	The source IP address of the request.
Moudle	User event module.
Event Name	User event name. Note: See the appendix User Event List for the user event name description.
Status	Response status code.
Error Code	Error Code.

- Query the latest system event.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor event ls --type system --num 1
```

```
event ID: 0422e451-e4bf-4c28-b7eb-76c64ade653e
module: target
instance ID: targetc,iqn.2012-08.cn.ctyunapi.oos:targetc.10
event time: 2025-08-12 14:17:52
name: TargetMigrated
details: {"targetIqn":"iqn.2012-08.cn.ctyunapi.oos:targetc.10","ips":"192.168.0.67:3260"}
```

Table 3: The description of system event

Item	Description
event ID	System event ID.
moudle	System event module.
instance ID	Instance ID targeted by the system event. -: Indicates no instance ID.
event time	The time the system event occurred.
name	System event name. Note: For a detailed description of the system event name, see System Event List .
details	System event details. -: Indicates no event details.

- Query last five system events.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor event ls --type system --num 5
```

Event Time	Event ID	Module	Event Name	Instance ID	Details
2022-08-01 14:57:14	ad9a053d-fa77-4494-b0c5-ad3c2131cf2a	Disk	PathAdded	hblock_1:/mnt/storage03	-
2022-08-01 09:51:24	fad76652-fc51-4384-8c83-e0fdd2ad0ecb	Disk	PathAdded	hblock_1:/mnt/storage02	-
2022-07-28 18:18:05	2f131000-7593-4b2f-a3ed-0b924e285e6c	Server	ServiceAvailable	hblock_1	{"message": "Service up","service": "ws"}
2022-07-28 18:16:01	1eb1d604-e2eb-4018-bcfc-385e9c28bfff7	Server	ServiceAvailable	hblock_1	{"message": "Service up","service": "ms"}

Table 4: The description of system event

Item	Description
Event Time	The time the system event occurred.
Event ID	System event ID.
Moudle	The module to which the event belongs.
Event Name	System event name. Note: For event name description, see System Event List .
Instance ID	Instance ID targeted by the system event. -: Indicates no instance ID.
Details	System event details. -: Indicates no system event details.

4.19.2 Export HBlock Events

```
./stor event ls [ --type TYPE ] [ --num NUMBER ] { -o | --out } DIRECTORY
```

This command is used to export HBlock events.

Parameters

Parameter	Description
<code>--type TYPE</code>	Specifies the type of exported HBlock event. Value: <ul style="list-style-type: none"> ● system: System Event. ● user: User event. The default value is user.
<code>--num NUMBER</code>	Specifies the number of latest HBlock events to export. The value is an integer that ranges from 1 to 1000. Note: If not specified, all HBlock events recorded by the system will be exported by default, and the system will record at least six months of user event information.
<code>-o DIRECTORY</code> or <code>--out DIRECTORY</code>	Export HBlock events to the specified folder.

Examples

- Export the latest 1000 user events to the folder /home/stor/events.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor event ls --num 1000 -o /home/stor/events
The output is in /home/stor/events/events_20220803172056.csv.
```

- Export the latest 1000 system events to the folder /home/stor/events.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor event ls --type system --num 1000 -o /home/stor/events
The output is in /home/stor/events/events_20220803172118.csv.
```

4.19.3 Initiate HBlock Log Collection

```
./stor logcollect add [ --start-time START_TIME ] [ --end-time END_TIME ] [ { -o | --out } DIIRECTORY ] [ --server SERVER_ID <1-n> ] [ --type LOG_TYPE <1-n> ]
```

This command is used to initiate HBlock log collection.

Note: The number of processes for HBlock log collection cannot exceed 10.

Parameters

Parameter	Description
--start-time <i>START_TIME</i>	The start time of log to be collected. The format is <i>yyyy-MM-dd HH:mm:ss</i> . The default value is the time point two hours before the end time of HBlock log collection. Note: <ul style="list-style-type: none"> ● The start time of the HBlock log collection must be earlier than the end time of the HBlock log collection. ● If the start time of the HBlock log collection is earlier than the HBlock initialization time, the HBlock initialization time is the log collection start time. ● If the system time is modified, the last modification time of the log file may fluctuate or even be incorrect.
--end-time <i>END_TIME</i>	The end time of log to be collected. The format is <i>yyyy-MM-dd HH:mm:ss</i> , and the default value is the current time. Note: <ul style="list-style-type: none"> ● The start time of the HBlock log collection must be earlier than the end time of the HBlock log collection. ● The end time of the HBlock log collection must be later than the HBlock initialization time. ● When the last modification time of the log file is greater than the end time, if there are multiple logs of the same type that are greater than the end time, the log file with the smallest HBlock log collection time value will be used.
-o <i>DIIRECTORY</i> or --out <i>DIIRECTORY</i>	The directory where HBlock logs are stored after collection. It is an absolute path. By default, it is stored in the HBlock installation directory of the requested server, named as <i>collected_logs/hblock_logs_id_yyyyMMddHHmms_yyyyMMddHHmms.zip</i> . <ul style="list-style-type: none"> ● <i>id</i>: Unique identifier for this log request. ● <i>yyyyMMddHHmms</i>: Start and end time of log collection, in UTC+0 time.
--server <i>SERVER_ID <1-n></i>	Server ID for log collection. HBlock logs of all servers are collected by default. Logs from one to n servers can be collected, where n is the number of servers in the cluster.

<code>--type LOG_TYPE</code> &<1-n>	HBlock log type collected: <ul style="list-style-type: none">● Config: Configuration-related logs.● System: System-related logs.● Data: Logs related to data processing (only supported by cluster mode).● Coordination: Logs related to internal coordination services (only supported by cluster mode). All types of logs are collected by default.
--	--

Examples

- Collect all types of logs from all servers within the latest two hours

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor logcollect add
Start collecting logs with ID f5dc5f0599a74be3af03441726707f01. After the logs are collected, they will be stored as a zip file in the /mnt/storage01/CTYUN_HBlock_Plus_3.9.0/collected_logs directory. You can check status with "./stor logcollect ls".
```

- Collect logs and store the log files in the path /mnt/storage02.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor logcollect add --start-time 2022-08-05 10:00:00 --end-time 2022-08-06 00:00:00 -o /mnt/storage02
Start collecting logs with ID a745d931721f4ec48bcc389b00109a39. After the logs are collected, they will be stored as a zip file in the /mnt/storage02 directory. You can check status with "./stor logcollect ls".
```


Create Time	The creation time of the log collection request.
File	Log file storage path.
Size	The size of the collection log files.
Type	The type of log collection.
Servers	Server ID for log collection.
Status	Log collection status: <ul style="list-style-type: none">● Processing● Succeeded● PartiallySucceeded● Failed
Message	If log collection fails or part of the collection succeeds, the reason will be displayed. -: Indicates no information is returned.

4.19.5 Delete Collection Logs

```
./stor logcollect rm [ --id LOG_ID ]
```

This command is used to delete the collection logs.

Parameters

Parameter	Description
<code>--id LOG_ID</code>	Log ID. If not specified, all logs collected by HBlock will be deleted.

Examples

Delete the log with log ID a2682b3a89b94c68924bd7e46703658f.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor logcollect rm --id  
a2682b3a89b94c68924bd7e46703658f  
Removed successfully.
```

4.20 HBlock System Settings

4.20.1 Change Administrator Password

```
./stor config set { -i | --item } auth {-p | --password} PASSWORD
```

This command is used to change the administrator password.

Parameters

Parameter	Description
<code>-i auth</code> or <code>--item auth</code>	Change administrator password.
<code>-p PASSWORD</code> or <code>--password PASSWORD</code>	Specifies the administrator password. The value is a string of 8 to 16 case-sensitive characters: <ul style="list-style-type: none"> ● The password must contain at least three of the following: lowercase letters, uppercase letters, digits, or special characters. The special characters only include ~ ! @ # \$ % ^ & * () _ + [] { } ; : , . / < > ?. ● The password cannot contain t 3 consecutive repeating characters, 3 consecutive or in-reverse order of digits or letters (case-insensitive), 3 consecutive or in-reverse order of keyboard sequences (case-insensitive).

Examples

Change administrator password.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config set -i auth -p *****
Set successfully.
```

4.20.2 Email Settings

4.20.2.1 Set Email Configuration

```
./stor config set { -i | --item } email [ { -H | --smtp-host } SMTP_HOST ] [ { -L | --ssl } SSL_STATUS ] [ { -P | --smtp-port } SMTP_PORT ] [ { -R | --receiver-email } RECEIVER_EMAIL &<1-n> ] [ { -S | --sender-email } SENDER_EMAIL { -p | --password } PASSWORD ] { -s | --status } STATUS [ -T | --test-email ]
```

This command is used to set the email notification function.

Note: If the email server address is IPv6 address, to ensure successful email sending, it is recommended that each HBlock server has an IPv6 address that can connect to the email server.

Parameters

Parameter	Description
-i email or --item email	Sets the email notification function.
-H SMTP_HOST or --smtp-host SMTP_HOST	Specifies the SMTP host. The value is the domain name or IP of SMTP server.
-L SSL_STATUS or --ssl SSL_STATUS	Whether to enable the SSL function: <ul style="list-style-type: none"> ● Enabled (on): Enable SSL function. ● Disabled (off): Disable SSL function. The default value is Disabled (off).
-P SMTP_PORT or --smtp-port SMTP_PORT	Specifies the SMTP port. The value is an integer that ranges from 1 to 65535. If the SSL function is enabled, the default value is 465; If the SSL function is disabled, the default value is 25.
-R RECEIVER_EMAIL &<1-n> or --receiver-email RECEIVER_EMAIL &<1-n>	Specifies the receiver email. Multiple emails can be specified at one time, separated by comma (,). The email format is <i>local-part@domain</i> : <ul style="list-style-type: none"> ● <i>local-part</i>: The value is a string of 1 to 64 characters. It can contain lowercase letters, uppercase letters, digits, or special characters (! # \$ % & * + - / = ? ^ _ ` { } ~ .). A dot (.) cannot be used as a starting or ending character, nor can it appear consecutively. ● <i>domain</i>: The value is a string of 1 to 255 characters separated by dots (.). Each string separated by a dot (.) needs to meet the following requirements: <ul style="list-style-type: none"> ■ A string of 1 to 63 characters.

	<ul style="list-style-type: none"> ■ It can contain lowercase letters, uppercase letters, digits, or hyphens (-). ■ The top-level domain name cannot consist of pure digits. ■ The hyphen (-) cannot be used as the first and last characters.
<p>-S <i>SENDER_EMAIL</i> or --sender-email <i>SENDER_EMAIL</i></p>	<p>Specifies the sender email. The email format is <i>local-part@domain</i>:</p> <ul style="list-style-type: none"> ● <i>local-part</i>: The value is a string of 1 to 64 characters. It can contain lowercase letters, uppercase letters, digits, or special characters (!#\$%&*+ - / = ? ^ _ ` { } ~ .). A dot (.) cannot be used as a starting or ending character, nor can it appear consecutively. ● <i>domain</i>: The value is a string of 1 to 255 characters separated by dots (.). Each string separated by a dot (.) needs to meet the following requirements: <ul style="list-style-type: none"> ■ A string of 1 to 63 characters. ■ It can contain lowercase letters, uppercase letters, digits, or hyphens (-). ■ The top-level domain name cannot consist of pure digits. ■ The hyphen (-) cannot be used as the first and last characters.
<p>-p <i>PASSWORD</i> or --password <i>PASSWORD</i></p>	<p>Specifies authorization code of the sender email. Note: Email authorization code is a special password issued by the mailbox and used for third-party client login.</p>
<p>-s <i>STATUS</i> or --status <i>STATUS</i></p>	<p>Whether to enable the email notification function:</p> <ul style="list-style-type: none"> ● Enabled (on): Turn on the email notification function. ● Disabled (off): Disable the email notification function.
<p>-T or --test-email</p>	<p>Enables sending test emails.</p>

Examples

Set the email notification function.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config set -i email -H
smtp.chinatelecom.cn -P 475 -R test-1@chinatelecom.cn,test.2@chinatelecom.cn -S test-
2@chinatelecom.cn -p ***** -s on
Set successfully.
```

4.20.2.2 Delete Email Configuration

```
./stor config rm { -i | --item } email
```

This command is used to delete email configuration.

Examples

Delete email configuration information.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config rm -i email  
Removed successfully.
```

4.20.2.3 Query Email Configuration

```
./stor config ls { -i | --item } email
```

This command is used to query email configuration information.

Examples

Query email configuration information.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config ls -i email
Status: Enabled
SMTP server: smtp.ctyun.cn
SMTP port: 475
Sender Email: test@ctyun.cn
Enable SSL (Yes/No): No
Receiver Email: test1@chinatelecom.cn
```

Table: The description of email configuration

Item	Description
Status	Whether to enable the email notification function: <ul style="list-style-type: none">● Enabled: Eable email notification function.● Disabled: Disable the email notification function.
SMTP server	SMTP server.
SMTP port	SMPT port number.
Sender Email	Outbox.
Enable SSL(Yes/No)	Whether to enable the SSL function: <ul style="list-style-type: none">● Yes: Enable SSL function.● No: Disable SSL function.
Receiver Email	Inbox.

4.20.3 Remote Assistance Settings

4.20.3.1 Set Remote Assistance Configuration

```
./stor config set { -i | --item } remote {-H | --host} HOST {-P | --port} PORT { -s | --status } STATUS [ { -n | --server } SERVER_ID ]
```

This command is used to set remote assistance configuration.

Parameters

Parameter	Description
-i remote or --item remote	Sets the remote assistance configuration.
-H HOST or --host HOST	Specifies the remote assistance host. Please contact the software vendor to get the HOST info.
-P PORT or --port PORT	Specifies the remote assistance port. The value is an integer that ranges from 1 to 65535. The default value is 18100. Please contact the software vendor to get the PORT.
-s STATUS or --status STATUS	Whether to enable the remote assistance function: <ul style="list-style-type: none"> ● Enabled (on): Enable the remote assistance function. ● Disabled (off): Disable the remote assistance function.
-n SERVER_ID or --server SERVER_ID	Server ID to configure remote assistance. If the server ID is not specified, remote assistance is configured for the server executing the command.

Examples

Set remote assistance configuration.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config set -i remote -H 120.71.10.106 -P 18100 -s on
Remote access to 120.71.10.106:18100 has been enabled, and the remote access code is: 763589. Please tell the HBlock assistant the code or the HBlock ID (run 'stor info --stor-id' to get it).
```

4.20.3.2 Delete Remote Assistance Configuration

```
./stor config rm { -i | --item } remote [ { -n | --server } SERVER_ID ]
```

This command is used to delete remote assistance configuration.

Parameters

Parameter	Description
<code>-n SERVER_ID</code> or <code>--server SERVER_ID</code>	Server ID to delete remote assistance from. If no server ID is specified, Remote Assistance is removed for the server on which the command was executed.

Examples

Delete remote assistance configuration.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config rm -i remote  
Removed successfully.
```

4.20.3.3 Query Remote Assistance Configuration

`./stor config ls { -i | --item } remote`

This command is used to query remote assistance configuration information.

Examples

Query remote assistance configuration information.

```
[root@hbblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config ls -i remote
+-----+-----+-----+-----+-----+-----+-----+
| No. | Server Name | Server ID | Public Address | Cluster Address | Remote Access Host and Port | Remote Access Code |
+-----+-----+-----+-----+-----+-----+-----+
| 1 | server3 | hbblock_3 | 192.168.0.31 | 192.168.0.31 | 120.71.10.106:18100 | 494660 |
+-----+-----+-----+-----+-----+-----+-----+
```

Table: Description of remote assistance configuration

Item	Description
No.	Number.
Server Name	Service name.
Server ID	Server ID.
Public Address	Public network IP.
Cluster Address	Cluster network IP.
Remote Access Host and Port	Remote assistance host and port.
Remote Access Code	Remote assistance code.

4.20.4 Pushgateway Monitoring Configuration

4.20.4.1 Add Pushgateway Monitoring Configuration

```
./stor config add { -i | --item } monitor --pushgateway PUSHGATEWAY [ --pushgateway-  
labels LABEL &<1-n> ] [ { -n | --server } SERVER_ID &<1-n> ]
```

This command is used to add Pushgateway monitoring configuration.

Note: When adding the monitoring configuration for Pushgateway, the following labels are included by default: job, cluster ID, cluster name, ServerIP, and hostname. Users are allowed to modify or delete the cluster ID, cluster name, ServerIP, and hostname, but it is not recommended to do so. If a user-defined label has the same name as a default label, the value of the label will be determined by the user's settings.

Parameters

Parameter	Description
<code>--pushgateway PUSHGATEWAY</code>	Specifies the address and port of Pushgateway, the format is <i>IPv4:port</i> , <i>[IPv6]:port</i> or <i>domain-name:port</i> .
<code>--pushgateway-labels LABEL &<1-n></code>	Specifies the key-value pairs which assigned to the data in the Pushgateway, the format is <i>KEY:VALUE</i> . Multiple key-value pairs can be added at one time, separated by comma (,).
<code>-n SERVER_ID&<1-n></code> or <code>--server SERVER_ID &<1-n></code>	Specifies the server ID whose monitor configuration is to be added. Multiple server IDs can be added at one time, separated by comma (,). By default, all HBlock servers are configured with Pushgateway monitoring.

Examples

Add Pushgateway monitoring configuration.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config add -i monitor --  
pushgateway 192.168.0.1:9091 --pushgateway-labels agent:agentname,idc:idcname -n  
hblock_1,hblock_2  
Added pushgateway 192.168.0.1:9091 with pushgateway labels agent:agentname,idc:idcname on  
server hblock_1,hblock_2 successfully.
```

4.20.4.2 Modify Pushgateway Monitoring Configuration

```
./stor config set { -i | --item } monitor --pushgateway PUSHGATEWAY [ --pushgateway-
labels LABEL &<1-n> ] [ --all ] [ --pushgateway-timeout PUSHGATEWAY_TIMEOUT ] [ --
collect-interval COLLECT_INTERVAL ] [ --collect-metric COLLECT_METRIC [ --collect-metric-
items ITEM &<1-n> ] ] [ { -n | --server } SERVER_ID &<1-n> ]
```

This command is used to modify Pushgateway monitoring configuration.

Note: It is not recommended to modify the default labels carried: job, cluster ID, cluster name, ServerIP, and hostname.

Parameters

Parameter	Description
<code>--pushgateway PUSHGATEWAY</code>	Specifies the address and port of Pushgateway to be modified, the format is <i>IPv4:port</i> , <i>[IPv6]:port</i> or <i>domain-name:port</i> .
<code>--pushgateway-labels LABEL &<1-n></code>	Specifies the key-value pairs which assigned to the data in the Pushgateway, the format is <i>KEY:VALUE</i> . Multiple key-value pairs can be modified at one time, separated by comma (,).
<code>--all</code>	All labels/items configurations will be replaced with the specified value of labels/metrics. If <code>--pushgateway-labels LABEL</code> or <code>--collect-metric-items ITEM</code> is not specified, this parameter does not take effect.
<code>--pushgateway-timeout PUSHGATEWAY_TIMEOUT</code>	Specifies the timeout for sending data to Pushgateway. Value: If the unit is s (seconds), the value is an integer that ranges from 1 to 1000. If the unit is ms (milliseconds), the value is an integer that ranges from 1000 to 1000000. The default unit is s (seconds).
<code>--collect-interval COLLECT_INTERVAL</code>	Specifies the collection interval on the server. Value: If the unit is s (seconds), the value is an integer that ranges from 2 to 1000, or a negative integer. If the unit is ms (milliseconds), the value is an integer that ranges from 2000 to 1000000, or a negative integer. The default unit is s (seconds). Note: <ul style="list-style-type: none"> When specified together with <code>--collect-metric COLLECT_METRIC</code>,

	<p>the value can be a negative integer, indicating that the specified monitoring metric will not be collected.</p> <ul style="list-style-type: none"> ● If --collect-metric <i>COLLECT_METRIC</i> is not specified, it indicates setting the collection interval for all monitoring metrics, and the value can only be a positive integer within the allowed range.
--collect-metric <i>COLLECT_METRIC</i>	<p>Specifies the monitoring metrics modified. The value is server, fileSystem, interface, load, disk, tcp or os. All metrics will be collected by default.</p>
--collect-metric-items <i>ITEM</i> &<1- <i>n</i> >	<p>Specifies the monitoring metric items to be modified. Only the monitoring indicators disk, fileSystem, and interface can specify metric items.</p> <ul style="list-style-type: none"> ● If monitoring metric is specified and no metric item is specified, all metric items of monitoring indicator will be modified by default. ● If monitoring metric is not specified, this parameter cannot be specified.
-n <i>SERVER_ID</i> &<1- <i>n</i> > or --server <i>SERVER_ID</i> &<1- <i>n</i> >	<p>Specifies the HBlock server ID to modify the Pushgateway monitoring configuration. Multiple server IDs can be specified at one time, separated by comma (,). By default, modify the Pushgateway monitoring configuration for all HBlock servers.</p>

Examples

Modify Pushgateway monitoring configuration.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config set -i monitor --
pushgateway 192.168.0.1:9091 --pushgateway-labels agent:agentname -all -n hblock_1
Set all pushgateway labels {agent=agentname} of pushgateway 192.168.0.1:9091 on server
hblock_1 successfully.
```

4.20.4.3 Delete Pushgateway Monitoring Configuration

```
./stor config rm { -i | --item } monitor --pushgateway PUSHGATEWAY [ --pushgateway-  
labels LABEL &<1-n> ] [ --collect-metric COLLECT_METRIC [ --collect-metric-items ITEM  
&<1-n> ] ] [ { -n | --server } SERVER_ID &<1-n> ]
```

This command is used to delete Pushgateway monitoring configuration.

Note: Deleting the "labels that metrics must have" may pose the risk of making the related monitoring data unrecognizable.

Parameters

Parameter	Description
<code>--pushgateway PUSHGATEWAY</code>	Specifies the address and port of Pushgateway to be deleted, the format is <i>IPv4:port</i> , <i>[IPv6]:port</i> or <i>domain-name:port</i> .
<code>--pushgateway-labels LABEL &<1-n></code>	Specifies the keys of Pushgateway labels. Multiple key-value pairs can be deleted at one time, separated by comma (,). Value: The value is the value of the <i>key</i> .
<code>--collect-metric COLLECT_METRIC &<1-n></code>	Specifies the monitoring metrics to be deleted. The value is <i>server</i> , <i>fileSystem</i> , <i>interface</i> , <i>load</i> , <i>disk</i> , <i>tcp</i> or <i>os</i> . All metrics will be deleted by default.
<code>--collect-metric-items ITEM</code>	Specifies the monitoring metric items to be deleted. Multiple metric items can be deleted at one time, separated by comma (,).
<code>-n SERVER_ID &<1-n></code> or <code>--server SERVER_ID &<1-n></code>	Specifies the HBlock server ID to delete the Pushgateway monitoring configuration. Multiple server IDs can be specified at one time, separated by comma (,). By default, delete the Pushgateway monitoring configuration for all HBlock servers.

Examples

Delete Pushgateway monitoring configuration.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config rm -i monitor --  
pushgateway 192.168.0.1:9091  
Removed pushgateway 192.168.0.1:9091 successfully.
```

4.20.4.4 Query Pushgateway Monitoring Configuration

```
./stor config ls { -i | --item } monitor [ { -n | --server } SERVER_ID &<1-n> ]
```

This command is used to query Pushgateway monitoring configuration information.

Parameters

Parameter	Description
<code>-n SERVER_ID &<1-n></code> or <code>--server SERVER_ID &<1-n></code>	Specifies the HBlock server ID to query the Pushgateway monitoring configuration. Multiple server IDs can be specified at one time, separated by comma (.). By default, query the Pushgateway monitoring configuration of all HBlock servers.

Examples

Query Pushgateway monitoring configuration information for all HBlock servers.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config ls -i monitor
{
  "monitors" : [ {
    "serverId" : "hblock_1",
    "send" : {
      "interval" : 20000,
      "timeout" : 2000,
      "pushgatewayTimeout" : 10000,
      "pushgateways" : [ {
        "address" : "192.168.0.1",
        "port" : 9091,
        "labels" : {
          "agent" : "agentname",
          "idc" : "idcname",
          "hostname" : "hblockserver",
          "serverIp" : "192.168.0.110",
          "clusterId" : "B8BFE310-8DF6-429E-8C8A-ACD22310AC08",
          "clusterName" : "stor2",
          "job" : "hblock"
        }
      }
    ]
  },
  "collect" : {
    "interval" : 14000,
    "metrics" : [ {
      "name" : "storFs",
      "interval" : 20000,
      "items" : [ "ds-1", "ds-2", "mdm" ]
    } ]
  }
}
```

```
}
}, {
  "serverId" : "hblock_3",
  "send" : {
    "interval" : 20000,
    "timeout" : 2000,
    "pushgatewayTimeout" : 10000,
    "pushgateways" : [ ]
  },
  "collect" : {
    "interval" : 14000,
    "metrics" : [ {
      "name" : "storFs",
      "interval" : 20000,
      "items" : [ "ds-1" ]
    } ]
  }
}, {
  "serverId" : "hblock_2",
  "send" : {
    "interval" : 20000,
    "timeout" : 2000,
    "pushgatewayTimeout" : 10000,
    "pushgateways" : [ {
      "address" : "192.168.0.1",
      "port" : 9091,
      "labels" : {
        "agent" : "agentname",
        "idc" : "idcname",
        "hostname" : "pm-006",
        "serverIp" : "192.168.0.192",
        "clusterId" : "B8BFE310-8DF6-429E-8C8A-ACD22310AC08",
        "clusterName" : "stor2",
        "job" : "hblock"
      }
    } ]
  },
  "collect" : {
    "interval" : 14000,
    "metrics" : [ {
      "name" : "storFs",
      "interval" : 20000,
      "items" : [ "ds-1", "mdm" ]
    } ]
  }
} ]
}
```

Query the pushgateway monitoring configuration of servers hblock_1 and hblock_2.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor config ls -i monitor -n
hblock_1,hblock_2
{
  "monitors" : [ {
    "serverId" : "hblock_1",
    "send" : {
      "interval" : 20000,
      "timeout" : 2000,
      "pushgatewayTimeout" : 10000,
      "pushgateways" : [ {
        "address" : "192.168.0.1",
        "port" : 9091,
        "labels" : {
          "agent" : "agentname",
          "idc" : "idcname",
          "hostname" : "hblockserver",
          "serverIp" : "192.168.0.110",
          "clusterId" : "B8BFE310-8DF6-429E-8C8A-ACD22310AC08",
          "clusterName" : "stor2",
          "job" : "hblock"
        }
      }
    ]
  },
  "collect" : {
    "interval" : 14000,
    "metrics" : [ {
      "name" : "storFs",
      "interval" : 20000,
      "items" : [ "ds-1", "ds-2", "mdm" ]
    } ]
  }
}, {
  "serverId" : "hblock_2",
  "send" : {
    "interval" : 20000,
    "timeout" : 2000,
    "pushgatewayTimeout" : 10000,
    "pushgateways" : [ {
      "address" : "192.168.0.1",
      "port" : 9091,
      "labels" : {
        "agent" : "agentname",
        "idc" : "idcname",
        "hostname" : "pm-006",
        "serverIp" : "192.168.0.192",
        "clusterId" : "B8BFE310-8DF6-429E-8C8A-ACD22310AC08",
        "clusterName" : "stor2",
        "job" : "hblock"
      }
    }
  ]
},
},
```

```

"collect" : {
  "interval" : 14000,
  "metrics" : [ {
    "name" : "storFs",
    "interval" : 20000,
    "items" : [ "ds-1", "mdm" ]
  } ]
}
} ]
}

```

Description of pushgateway monitoring configuration information

Item	Description
serverId	HBlock server ID.
send.pushgatewayTimeout	The timeout for sending data to pushgateways. Timeout is considered a push failure.
send.pushgateways.adress	The IP or domain name of pushgateway.
send.pushgateways.labels	The labels corresponding to pushgateway.
collect.interval	The collection interval for hardware indicators.
collect.metrics.name	The hardware indicator category name of the collection interval or collection object.
collect.metrics.interval	The collection interval for specified metric items.
collect.metrics.items	The specified metric items of hardware indicators.

4.20.5 Authentication Mode

4.20.5.1 Set the Authentication Mode

```
./stor config set { -i | --item } parameter --auth-mode AUTH_MODE
```

This command is used to set the HBlock authentication mode.

Parameters

Parameter	Description
<code>-i parameter</code> or <code>--item parameter</code>	Sets HBlock system parameters.
<code>--auth-mode AUTH_MODE</code>	<p>Sets the HBlock authentication mode.</p> <p>Value:</p> <ul style="list-style-type: none"> ● strict: Uses the latest authentication method. ● compatible: Compatible with historical authentication methods. <p>For HBlock upgraded from versions below 3.8, the authentication mode is compatible; for versions 3.8 and above, the default authentication mode after installation is strict.</p> <p>Note:</p> <ul style="list-style-type: none"> ● Authentication method before version 3.8: SHA256(Date header value + "\n" + SHA256(password)) ● Authentication method for version 3.8 and later: HMAC-SHA256(SHA256(password), Date header value + "\n" + SHA256(password))

Examples

Set the HBlock authentication mode to **strict**.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor config set -i parameter --auth-mode strict
Set successfully.
```

4.20.5.2 Query the Authentication Mode

```
./stor config ls { -i | --item } parameter
```

This command is used to query the HBlock authentication mode.

Parameters

Parameter	Description
<code>-i parameter</code> or <code>--item parameter</code>	Specifies to query HBlock system parameters.

Examples

Query the HBlock authentication mode.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor config ls -i parameter
{"system":{"authMode":"compatible"}}
```

Description of authentication mode query information

Item	Description
authMode	The HBlock authentication mode: <ul style="list-style-type: none">● <code>strict</code>: Uses the latest authentication method.● <code>compatible</code>: Compatible with historical authentication methods.

4.20.6 Enable Pro Trial (Free Edition)

```
./stor config set { -i | --item } protrial [ --confirm ]
```

This command is used to enable the Pro Trial for the free edition.

After enabling the Pro Trial, some advanced features of the commercial edition can be trialed.

For differences between the commercial edition and the free edition, see **Commercial Edition and Free Edition**.

Note:

- The free edition includes 2 years of upgrade service; upgrades will not be supported after expiration. If the user has never enabled the Pro Trial, it can still be enabled normally after 2 years.
- For the free edition, the Pro Trial can only be enabled once. After enabling the Pro Trial, advanced features can be trialed for 1 month; after 1 month, advanced features will be disabled. When advanced features are disabled, if you wish to use the commercial edition, please contact the software vendor to obtain a formal software license and load it.

Parameters

Parameter	Description
<code>-i protrial</code> or <code>--item protrial</code>	Enable the Pro Trial.
<code>--confirm</code>	Confirm to enable the Pro Trial.

Examples

Enable the Pro Trial for the free edition.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor config set -i protrial
You are about to enable the Pro trial for the free edition. Please note that this action
cannot be undone.
Once activated, you will have free access to all professional features for 30 days. After
the trial expires, access will be automatically revoked. Please also note that this is
your one and only trial opportunity.
Are you sure you want to proceed? [Yes/No]
y
The Pro trial (started on 2026-01-12 11:29:43) ends on 2026-02-11 11:29:43. This is your
one and only trial opportunity.
```

4.21 Tune the Performance Parameters of HBlock (Cluster Mode)

4.21.1 Tune the Performance Parameters of HBlock

```
./stor tuning set [ --perf-scen PERF_SCEN ] [ --ha-sens HA_SENS ]
```

This command is used to optimize performance based on the hardware environment and business scenario.

Parameters

Parameter	Description
<code>--perf-scen <i>PERF_SCEN</i></code>	Specifies the scenario for performance tuning. Value: <ul style="list-style-type: none"> ● ICv1: Intelligent Computing Scenario. ● VMv1: Virtualization Scenario. ● default: The default type of the system.
<code>--ha-sens <i>HA_SENS</i></code>	Sets performance tuning for scenario sensitive types. Value: <ul style="list-style-type: none"> ● High: High-Sensitivity Scenario. ● Middle: Medium-Sensitivity Scenario. ● Low: Low-Sensitivity Scenario. The default value is Middle.

Examples

HBlock performance tuning.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor tuning set --perf-scen ICv1 --ha-sens High
Start adjusting parameters for optimization. You can use 'stor tuning ls' to check whether it is completed.
```

4.21.2 View the Performance Tuning Configuration

`./stor tuning ls`

This command is used to view the performance tuning configuration.

Examples

View the performance tuning configuration.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor tuning ls
The current optimization is adapted to the following scenarios.
Scenario for performance optimization: ICv1
Sensitivity for high availability: High
```

Table1. Description of performance tuning configuration information

Item	Description
Scenario for performance optimization	The scenario for performance tuning: <ul style="list-style-type: none"> ● ICv1: Intelligent Computing Scenario. ● VMv1: Virtualization Scenario. ● default: The default type of the system.
Sensitivity for high availability	The performance tuning for scenario sensitive types: <ul style="list-style-type: none"> ● High: High-Sensitivity Scenario. ● Middle: Medium-Sensitivity Scenario. ● Low: Low-Sensitivity Scenario. ● custom: The scenario set by the user. ● Unknown: The scenario is unknown.

4.22 Stop HBlock on the Server

```
./stor stop [ --scope SCOPE &<1-n> ] [ -f | --force ]
```

This command is used to stop HBlock services on the current server.

Parameters

Parameter	Description
<code>--scope SCOPE &<1-n></code>	<p>The scope of HBlock services to stop. Multiple services can be specified at one time, separated by comma (,).</p> <p>Value:</p> <ul style="list-style-type: none"> ● <code>All</code>: Stop all services on the server. ● <code>Min</code>: Stop only the ms and ps services on the server. ● <code>mdm</code>: Metadata management service (only supported by cluster mode). ● <code>fc</code>: Failover control service (only supported by cluster mode). ● <code>ls</code>: Ledger service (only supported by cluster mode). ● <code>ds</code>: Data service (only supported by cluster mode). ● <code>ds-x</code>: Specific data service, manage a particular disk path (only supported by cluster mode). ● <code>cs</code>: Coordination service (only supported by cluster mode). ● <code>ms</code>: Management service. ● <code>ws</code>: Watchdog service. ● <code>ps</code>: Processor service. ● <code>ag</code>: Aggregator service. ● <code>ua</code>: Upgrade agent service. <p>The default value is <code>Min</code>.</p> <p>Note: For details on the specific functions of each service, see HBlock Service.</p>
<code>-f</code> or <code>--force</code>	<p>Forcibly stop HBlock services on the current server.</p> <p>Note: Forcibly stopping HBlock services may cause data loss.</p>

Examples

- Stop the HBlock ms and ps services on the current server.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor stop
Stopping HBlock...
Stop completely.
```

- Stop the HBlock ws and ag services on the current server.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor stop --scope ws,ag
```

```
Stopping HBlock...  
Stop completely.
```

- Forcibly stop the HBlock ms and ps services on the current server.

Note: Forcibly stopping HBlock services may cause data loss.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor stop --force  
Stopping HBlock...  
Stop completely.
```

4.23 Start HBlock on the Server

```
./stor start [ { -t | --timeout } TIMEOUT ]
```

This command is used to start HBlock services on the current server.

Parameters

Parameter	Description
<code>-t <i>TIMEOUT</i></code> or <code>--timeout <i>TIMEOUT</i></code>	Set the maximum waiting time for starting the HBlock service on the current server. If the HBlock service does not start successfully within the maximum time, a startup failure will be reported. The value is an integer that ranges from 1 to 2,147,483,647. The unit is second (s).

Examples

Start HBlock services on the current server.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor start
Starting HBlock...
Start completely.
```

4.24 Restart HBlock on the Server

```
./stor restart [ --scope SCOPE &<1-n> ] [ { -f | --force } ]
```

This command is used to restart HBlock services on the current server.

Parameters

Parameter	Description
<code>--scope SCOPE &<1-n></code>	<p>The scope of HBlock services to restart. Multiple services can be specified at one time, separated by comma (,).</p> <p>Value:</p> <ul style="list-style-type: none"> ● <code>All</code>: Restart all services on the server. ● <code>Min</code>: Restart only the <code>ms</code> and <code>ps</code> services on the server. ● <code>mdm</code>: Metadata management service (only supported by cluster mode). ● <code>fc</code>: Failover control service (only supported by cluster mode). ● <code>ls</code>: Ledger service (only supported by cluster mode). ● <code>ds</code>: Data service (only supported by cluster mode). ● <code>ds-x</code>: Specific data service, manage a particular disk path (only supported by cluster mode). ● <code>cs</code>: Coordination service (only supported by cluster mode). ● <code>ms</code>: Management service. ● <code>ws</code>: Watchdog service. ● <code>ps</code>: Processor service. ● <code>ag</code>: Aggregator service. ● <code>ua</code>: Upgrade agent service. <p>The default value is <code>Min</code>.</p> <p>Note: For details on the specific functions of each service, see HBlock Service.</p>
<code>-f</code> or <code>--force</code>	<p>Forcibly restart HBlock on the current server.</p> <p>Note: Forcibly restarting HBlock may cause data loss.</p>

Examples

- Restart the HBlock `ws` and `ag` services on the current server.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor restart --scope ws,ag
Stopping HBlock...
Stop completely.
Starting HBlock...
System has been started, no need for duplicate actions. You can use restart if necessary.
```

- Forcibly restart the HBlock `ms` and `ps` services on the current server.

Note: Forcibly restarting HBlock may cause data loss.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor restart --force
Stopping HBlock...
Stop completely.
Starting HBlock...
Start completely.
```

4.25 Uninstall HBlock

```
./stor uninstall [ -l | --del-local ] [ -c | --del-cloud ]
```

This command is used to uninstall HBlock.

Note: For the cluster mode, using this command will uninstall HBlock on all servers in the cluster.

Parameters

Parameter	Description
<code>-l</code> or <code>--del-local</code>	Uninstall HBlock and delete local HBlock data.
<code>-c</code> or <code>--del-cloud</code>	Uninstall HBlock and delete HBlock data from the cloud. Note: When uninstalling HBlock, if you choose to delete HBlock data from the cloud, the uninstallation process may not be successful due to various reasons, such as network issues, permission issues, etc. In this case, you can choose not to delete the data on the cloud and perform the HBlock uninstallation operation first. After successful uninstallation, you log in to the cloud service and manually deletes the HBlock data.

Examples

Uninstall HBlock.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor uninstall
All servers will be uninstalled. The data on the local or in the cloud will not be
deleted by default. If you want to delete, please use the -l or -c option. The deleted
data cannot be recovered, please exercise caution. Are you sure you want to uninstall
HBlock? [Yes/No]
[0]
y
Start uninstalling HBlock, please wait.
Processing...
Uninstalled successfully.
```

4.26 Query HBlock Version

```
./stor { --version | -v }
```

This command is used to query the HBlock version.

Examples

Query the HBlock Version.

```
[root@hblockserver CTYUN_HBlock_Plus_4.0.0_x64]# ./stor --version  
3.0.0
```

4.27 Upgrade HBlock

4.27.1 Upgrade HBlock

```
./stor upgrade { --filename | -f } file [ --archfile PACKAGE &<1-n> ] [ --key VALUE ]
```

This command is used to upgrade HBlock. For clusters, the upgrade only needs to be performed on one server.

During the upgrade process, the system will perform checks. If the conditions for upgrade are not met, the upgrade may fail. The reason for the upgrade failure can be viewed in the log file `upgrade.log` (Log path: `HBlock_installation_directory/logs/ops/upgrade.log`) on the server where the upgrade operation was executed. We recommend that you check the system before the upgrade to ensure that:

- All statuses are normal:
 - All HBlock servers are connected properly, no server is in the process of being deleted.
 - HBlock is in working and upgrading status.
 - If the software license is subscription, it must be within the validity period; if the license is perpetual, it must be within the maintenance period; if it is in trial mode, the trial must not have expired.
- Cluster mode: If the LUN's type is ActiveStandby, the active/standby connection is normal.
- No LUNs are in failed or in-progress status.
- The overall data redundancy of the system is not degraded, and the percentage of normal data is 100%. In addition, the number of available fault domains and the number of healthy domains are greater than the write requirements of all LUNs.
- The listening service (`stor:ua`) can be properly upgraded.
- The protocol parsing service (`stor:ps`) is normal.
- Cluster mode: Data service (`stor:ds-x`) is normal.
- Cluster mode: Base services are normal, including metadata management service (`stor:mdm`), log service (`stor:ls`), and coordination service (`stor:cs`).

Note: Before upgrading HBlock, ensure that the file system corresponding to the HBlock installation path of each server has at least 1 GiB of free space.

Parameters

Parameter	Description
<code>--filename file</code> or <code>-f file</code>	The specific path and file name of the HBlock installation package for the target version of the current server upgrade.

	<p>Note:</p> <ul style="list-style-type: none"> ● The installation package used for upgrading the current server must match the type of the installation package used by the server. ● The specific path of the installation package should not contain -, --, or spaces.
--archfile <i>PACKAGE</i> &<1- <i>n</i> >	<p>The paths and file names of the installation packages for other servers with different architectures in the cluster to be upgraded (only supported in the cluster mode). If there are multiple types of servers, multiple installation packages should be added, separated by commas.</p> <p>Note: This parameter should be specified only when there are servers with different architectures from the server executing the upgrade within the cluster. It is used to specify the HBlock installation package for the corresponding architecture server.</p>
--key <i>VALUE</i>	<p>When upgrading to the target version of HBlock, if the version requires importing parameters, this item must be set. The specific parameter settings should be based on the target version. If the target version requires multiple parameters, they can be added multiple times in the format "--key VALUE", where "key" is the parameter name and "VALUE" is the parameter value.</p>

Examples

Upgrade the HBlock service: Upgrade from HBlock 3.10.0 to HBlock 4.0.0.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor upgrade -f
/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64.zip
Starting upgrade, the current version is 3.10.0, the target version is 4.0.0
Start time: 2026-03-11 13:42:14
Server info:
hblock_1, 192.168.0.65, hblockserver
hblock_2, 192.168.0.64, k8s-master
hblock_3, 192.168.0.67, songt-0006

2026-03-11 13:42:14 CST+0800 [Step 1/5] Checking for system status...
2026-03-11 13:42:14 CST+0800 Checking system info
2026-03-11 13:42:14 CST+0800 Checking data status

2026-03-11 13:42:14 CST+0800 [Step 2/5] Uploading update files...
2026-03-11 13:42:18 CST+0800 Upload [100%] for hblock_1 ###[server 1/3]
2026-03-11 13:42:20 CST+0800 Upload [100%] for hblock_2 ###[server 2/3]
2026-03-11 13:42:23 CST+0800 Upload [100%] for hblock_3 ###[server 3/3]

2026-03-11 13:42:23 CST+0800 [Step 3/5] Upgrade preparing...
```

```
2026-03-11 13:42:59 CST+0800 [Step 4/5] Installing updates...
2026-03-11 13:42:59 CST+0800 Upgrading config - Update[√] ### [module 1/12]
2026-03-11 13:43:08 CST+0800 Upgrading cs for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [server
1/3] [module 2/12]
2026-03-11 13:43:52 CST+0800 Upgrading cs for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [server
2/3] [module 2/12]
2026-03-11 13:44:36 CST+0800 Upgrading cs for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [server
3/3] [module 2/12]
2026-03-11 13:45:20 CST+0800 Upgrading storfsBase - Update[√] ### [module 3/12]
2026-03-11 13:45:29 CST+0800 Upgrading ls for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [server
1/3] [module 4/12]
2026-03-11 13:45:51 CST+0800 Upgrading ls for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [server
2/3] [module 4/12]
2026-03-11 13:46:08 CST+0800 Upgrading ls for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [server
3/3] [module 4/12]
2026-03-11 13:46:27 CST+0800 Upgrading mdm for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [server
1/3] [module 5/12]
2026-03-11 13:46:38 CST+0800 Upgrading mdm for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [server
2/3] [module 5/12]
2026-03-11 13:46:53 CST+0800 Upgrading mdm for hblock_3 - Update[√] ### [server 3/3] [module 5/12]
2026-03-11 13:46:57 CST+0800 Upgrading ds-1 for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ###
[service 1/1] [server 1/3] [module 6/12]
2026-03-11 13:47:10 CST+0800 Upgrading ds-1 for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ###
[service 1/1] [server 2/3] [module 6/12]
2026-03-11 13:47:21 CST+0800 Upgrading ds-1 for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ###
[service 1/1] [server 3/3] [module 6/12]
2026-03-11 13:47:32 CST+0800 Upgrading rb for hblock_1 - Check[√], Stop[√], Update[√] ### [server 1/3]
[module 7/12]
2026-03-11 13:47:39 CST+0800 Upgrading rb for hblock_2 - Check[√], Stop[√], Update[√] ### [server 2/3]
[module 7/12]
2026-03-11 13:47:45 CST+0800 Upgrading rb for hblock_3 - Check[√], Stop[√], Update[√] ### [server 3/3]
[module 7/12]
2026-03-11 13:47:51 CST+0800 Upgrading ps for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [server
1/3] [module 8/12]
2026-03-11 13:48:00 CST+0800 Upgrading ps for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [server
2/3] [module 8/12]
2026-03-11 13:48:05 CST+0800 Upgrading ps for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [server
3/3] [module 8/12]
2026-03-11 13:48:11 CST+0800 Upgrading tool - Update[√] ### [module 9/12]
2026-03-11 13:48:17 CST+0800 Upgrading remote for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ###
[server 1/3] [module 10/12]
2026-03-11 13:48:20 CST+0800 Upgrading remote for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ###
[server 2/3] [module 10/12]
2026-03-11 13:48:23 CST+0800 Upgrading remote for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ###
[server 3/3] [module 10/12]
2026-03-11 13:48:25 CST+0800 Upgrading monitor for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ###
[server 1/3] [module 11/12]
2026-03-11 13:48:33 CST+0800 Upgrading monitor for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ###
[server 2/3] [module 11/12]
2026-03-11 13:48:40 CST+0800 Upgrading monitor for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ###
[server 3/3] [module 11/12]
2026-03-11 13:48:47 CST+0800 Upgrading ws for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [service
1/2] [server 1/3] [module 12/12]
2026-03-11 13:48:52 CST+0800 Upgrading ms for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [service
2/2] [server 1/3] [module 12/12]
2026-03-11 13:49:04 CST+0800 Upgrading ws for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [service
1/2] [server 2/3] [module 12/12]
```

```
2026-03-11 13:49:09 CST+0800 Upgrading ms for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [service
2/2] [server 2/3] [module 12/12]
2026-03-11 13:49:20 CST+0800 Upgrading ws for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [service
1/2] [server 3/3] [module 12/12]
2026-03-11 13:49:25 CST+0800 Upgrading ms for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [service
2/2] [server 3/3] [module 12/12]

2026-03-11 13:49:36 CST+0800 [Step 5/5] Updating system info...

2026-03-11 13:49:38 CST+0800 [Post Upgrade] Cleaning...
2026-03-11 13:49:38 CST+0800 Cleaning hblock_1 ###[server 1/3]
2026-03-11 13:49:38 CST+0800 Cleaning hblock_3 ###[server 2/3]
2026-03-11 13:49:38 CST+0800 Cleaning hblock_2 ###[server 3/3]

2026-03-11 13:49:38 CST+0800 [Upgrade Complete] The system has been successfully upgraded to 4.0.0 version
```

4.27.2 Query the Upgrade Status

`./stor upgrade status`

This command is used to query the upgrade status of HBlock.

Examples

Query the upgrade status.

```
[root@hblockserver CTYUN_HBlock_Plus_3.10.0_x64]# ./stor upgrade status
The last upgrade operation was successfully executed. The target version is 4.0.0.
End time: 2026-03-11 13:49:38
Server info:
hblock_1, 192.168.0.65, hblockserver
hblock_2, 192.168.0.64, k8s-master
hblock_3, 192.168.0.67, songt-0006

2026-03-11 13:42:14 CST+0800 [Step 1/5] Checking for system status...
2026-03-11 13:42:14 CST+0800 Checking system info
2026-03-11 13:42:14 CST+0800 Checking data status

2026-03-11 13:42:14 CST+0800 [Step 2/5] Uploading update files...
2026-03-11 13:42:18 CST+0800 Upload [100%] for hblock_1 ###[server 1/3]
2026-03-11 13:42:20 CST+0800 Upload [100%] for hblock_2 ###[server 2/3]
2026-03-11 13:42:23 CST+0800 Upload [100%] for hblock_3 ###[server 3/3]

2026-03-11 13:42:23 CST+0800 [Step 3/5] Upgrade preparing...

2026-03-11 13:42:59 CST+0800 [Step 4/5] Installing updates...
2026-03-11 13:42:59 CST+0800 Upgrading config - Update[√] ### [module 1/12]
2026-03-11 13:43:08 CST+0800 Upgrading cs for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [server 1/3] [module 2/12]
2026-03-11 13:43:52 CST+0800 Upgrading cs for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [server 2/3] [module 2/12]
2026-03-11 13:44:36 CST+0800 Upgrading cs for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [server 3/3] [module 2/12]
2026-03-11 13:45:20 CST+0800 Upgrading storfsBase - Update[√] ### [module 3/12]
2026-03-11 13:45:29 CST+0800 Upgrading ls for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [server 1/3] [module 4/12]
2026-03-11 13:45:51 CST+0800 Upgrading ls for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [server 2/3] [module 4/12]
2026-03-11 13:46:08 CST+0800 Upgrading ls for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [server 3/3] [module 4/12]
2026-03-11 13:46:27 CST+0800 Upgrading mdm for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [server 1/3] [module 5/12]
2026-03-11 13:46:38 CST+0800 Upgrading mdm for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [server 2/3] [module 5/12]
2026-03-11 13:46:53 CST+0800 Upgrading mdm for hblock_3 - Update[√] ### [server 3/3] [module 5/12]
2026-03-11 13:46:57 CST+0800 Upgrading ds-1 for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [service 1/1] [server 1/3] [module 6/12]
2026-03-11 13:47:10 CST+0800 Upgrading ds-1 for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [service 1/1] [server 2/3] [module 6/12]
2026-03-11 13:47:21 CST+0800 Upgrading ds-1 for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [service 1/1] [server 3/3] [module 6/12]
```

```
2026-03-11 13:47:32 CST+0800 Upgrading rb for hblock_1 - Check[√], Stop[√], Update[√] ### [server 1/3]
[module 7/12]
2026-03-11 13:47:39 CST+0800 Upgrading rb for hblock_2 - Check[√], Stop[√], Update[√] ### [server 2/3]
[module 7/12]
2026-03-11 13:47:45 CST+0800 Upgrading rb for hblock_3 - Check[√], Stop[√], Update[√] ### [server 3/3]
[module 7/12]
2026-03-11 13:47:51 CST+0800 Upgrading ps for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [server
1/3] [module 8/12]
2026-03-11 13:48:00 CST+0800 Upgrading ps for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [server
2/3] [module 8/12]
2026-03-11 13:48:05 CST+0800 Upgrading ps for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [server
3/3] [module 8/12]
2026-03-11 13:48:11 CST+0800 Upgrading tool - Update[√] ### [module 9/12]
2026-03-11 13:48:17 CST+0800 Upgrading remote for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ###
[server 1/3] [module 10/12]
2026-03-11 13:48:20 CST+0800 Upgrading remote for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ###
[server 2/3] [module 10/12]
2026-03-11 13:48:23 CST+0800 Upgrading remote for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ###
[server 3/3] [module 10/12]
2026-03-11 13:48:25 CST+0800 Upgrading monitor for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ###
[server 1/3] [module 11/12]
2026-03-11 13:48:33 CST+0800 Upgrading monitor for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ###
[server 2/3] [module 11/12]
2026-03-11 13:48:40 CST+0800 Upgrading monitor for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ###
[server 3/3] [module 11/12]
2026-03-11 13:48:47 CST+0800 Upgrading ws for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [service
1/2] [server 1/3] [module 12/12]
2026-03-11 13:48:52 CST+0800 Upgrading ms for hblock_1 - Check[√], Stop[√], Update[√], Start[√] ### [service
2/2] [server 1/3] [module 12/12]
2026-03-11 13:49:04 CST+0800 Upgrading ws for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [service
1/2] [server 2/3] [module 12/12]
2026-03-11 13:49:09 CST+0800 Upgrading ms for hblock_2 - Check[√], Stop[√], Update[√], Start[√] ### [service
2/2] [server 2/3] [module 12/12]
2026-03-11 13:49:20 CST+0800 Upgrading ws for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [service
1/2] [server 3/3] [module 12/12]
2026-03-11 13:49:25 CST+0800 Upgrading ms for hblock_3 - Check[√], Stop[√], Update[√], Start[√] ### [service
2/2] [server 3/3] [module 12/12]

2026-03-11 13:49:36 CST+0800 [Step 5/5] Updating system info...

2026-03-11 13:49:38 CST+0800 [Post Upgrade] Cleaning...
2026-03-11 13:49:38 CST+0800 Cleaning hblock_1 ###[server 1/3]
2026-03-11 13:49:38 CST+0800 Cleaning hblock_3 ###[server 2/3]
2026-03-11 13:49:38 CST+0800 Cleaning hblock_2 ###[server 3/3]

2026-03-11 13:49:38 CST+0800 [Upgrade Complete] The system has been successfully upgraded to 4.0.0 version
```

5 FAQs

5.1 Deployment Environmental Requirements

Q: What are the basic requirements for the infrastructure environment and network environment for deploying HBlock?

A: HBlock is required to be deployed on a server with a Linux operating system environment and must meet the minimum configuration requirements. Each server in HBlock can access other servers through the network. In addition, users need to ensure that internal applications can access the HBlock service.

Q: How many servers are needed to deploy HBlock in a cluster?

A: At least three servers are required. HBlock stores data replicas in multiple fault domains. In this way, when any node fails, applications can access data replicas in other fault domains to ensure reliability and at the same time, the system reconstructs data on the faulty server and corrupted disk.

Q: After deploying HBlock, if the customer's existing application system reads and writes data to HBlock, does it require port modification?

A: No. HBlock provides standard iSCSI ports and supports Windows and Linux Clients. Under normal circumstances, customer applications do not require port modification.

Q: After deploying HBlock, what will be the impact of changing the server IP address in the cluster?

A: After HBlock is deployed:

- If the IP address used for server communication in the HBlock cluster is changed, the HBlock service on the server with the changed IP address will be unavailable.
- If the IP address connecting the HBlock server and the client changes, the client will be unable to connect to the server.

Therefore, after deploying HBlock, the server IP address that has been used cannot be changed.

Q: Does deploying HBlock require configuring NTP time synchronization?

A: No. Data storage through HBlock does not rely on the system's internal clock, and there is no need to configure NTP in advance.

Q: Does deploying HBlock require the configuration of a virtual IP address (VIP)?

A: No. HBlock adopts a distributed multi-controller architecture and does not rely on the traditional VIP mode. It only needs to ensure that the client is connected to the server where the active and standby targets are located and enables multi-path I/O (MPIO) to achieve second-level failover.

5.2 Common Operations

Q: Can the HBlock service start automatically after the server is restarted?

A: The HBlock service can be configured to start automatically after a server reboot using one of the following two solutions.

Prerequisite: The installation directory and disk path of HBlock have already been set up for automatic mounting at startup.

The main differences between the two solutions are as follows:

Category	Solution 1	Solution 2
Status	The systemd status remains consistent with that of the management service (stor:ms) (running, inactive, etc., change with the process).	The systemd status is only set at system startup; it does not change with the process status of the management service (stor:ms) thereafter.
Configuration file	The configuration file includes PIDFile and sets RemainAfterExit=no .	The configuration file removes PIDFile and sets RemainAfterExit=yes .
Upgrade handling	<p>Before and after the upgrade, the following operations must be completed on the servers where HBlock is configured to start on boot using this solution:</p> <p>Before upgrade:</p> <ol style="list-style-type: none"> 1. Modify the configuration file /usr/lib/systemd/system/HBlock.service: Comment out the ExecStop and PIDFile configurations to remove them. 2. Run the command systemctl daemon-reload to make the configuration take effect. 3. Run the command systemctl stop HBlock. 4. Run the command systemctl status HBlock to confirm that the systemd status is inactive (dead) and ensure systemd has stopped. 5. Run the command ./stor server ls to confirm that the Status of the currently modified node is Connected. <p>After upgrade:</p> <ol style="list-style-type: none"> 1. Modify the configuration file /usr/lib/systemd/system/HBlock.service: Restore the ExecStop and PIDFile configurations. 	No operations are required before or after the upgrade.

	<ol style="list-style-type: none"> 2. Run the command systemctl daemon-reload to make the configuration take effect. 3. Kill the stor:ws process and then the stor:ms process in sequence. 4. Run the command systemctl start HBlock. 5. Run the command systemctl status HBlock to confirm: <ul style="list-style-type: none"> ● The systemd status is active (running); ● The Main PID information exists. ● The Main PID is equal to the current process ID (PID) of stor:ms. 6. Run the command ./stor server ls to confirm that the Status of the currently modified node is Connected. 	
--	---	--

Solution 1

1. Copy the systemd service unit file.

```
cp HBlock_installation_directory/apps/tool/systemd/HBlock.service /usr/lib/systemd/system/
```

For example:

```
[root@hblockserver]# cp /mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64/apps/tool/systemd/HBlock.service /usr/lib/systemd/system/
[root@hblockserver]# cat /usr/lib/systemd/system/HBlock.service
[Unit]
Description=HBlock - Storage Resource Reutilization System
Documentation=https://www.ctyun.cn/products/hblock/
After=network-online.target
Wants=network-online.target
Before=iscsi.service iscsid.service

[Service]
Type=forking
KillMode=none
User=root
Group=root
TimeoutStartSec=660
TimeoutStopSec=660
RemainAfterExit=no

ExecStart=/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64/stor start -t 600
ExecStop=/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64/stor stop
PIDFile=/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64/run/pid/ms.pid
```

```
Restart=on-failure
RestartSec=30

[Install]
WantedBy=multi-user.target
```

Note: '/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64' is the installation directory of HBlock. If the installation is performed by the root user, use the default value root for User and Group. If the installation is performed by another user, please modify User and Group to the corresponding names.

2. Reload the systemd configuration and enable the service:

```
systemctl daemon-reload
systemctl enable HBlock.service
systemctl start HBlock
```

Solution 2

1. Modify the file **HBlock.service** under the path

HBlock_installation_directory/apps/tool/systemd:

- Change **RemainAfterExit=no** to **RemainAfterExit=yes**
- Delete the **PIDFile** field

2. Copy the systemd service unit file.

```
cp HBlock_installation_directory/apps/tool/systemd/HBlock.service /usr/lib/systemd/system/
```

For example:

```
[root@hblockserver]# cp
/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64/apps/tool/systemd/HBlock.service
/usr/lib/systemd/system/
[root@hblockserver]# cat /usr/lib/systemd/system/HBlock.service
[Unit]
Description=HBlock - Storage Resource Reutilization System
Documentation=https://www.ctyun.cn/products/hblock/
After=network-online.target
Wants=network-online.target
Before=iscsi.service iscsid.service

[Service]
Type=forking
KillMode=none
User=root
Group=root
TimeoutStartSec=660
TimeoutStopSec=660
RemainAfterExit=yes

ExecStart=/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64/stor start -t 600
ExecStop=/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64/stor stop
```

```
Restart=on-failure
RestartSec=30

[Install]
WantedBy=multi-user.target
```

Note: '/mnt/storage01/CTYUN_HBlock_Plus_4.0.0_x64' is the installation directory of HBlock. If the installation is performed by the root user, use the default value root for User and Group. If the installation is performed by another user, please modify User and Group to the corresponding names.

3. Reload the systemd configuration and enable the service:

```
systemctl daemon-reload
systemctl enable HBlock.service
systemctl start HBlock
```

Q: How to restart the HBlock Cluster?

A: You need to manually restart the HBlock service of each server in the cluster. In order to ensure that the user's data is not affected, it is recommended to proceed with the following steps:

1. Stop the read-and-write operations of the client application.
2. Client disconnects from iSCSI:
If client is Windows operating system, first take the disk offline, and then disconnect iSCSI target.
If client is Linux operating system, execute the following command.

```
umount DIRECTORY_NAME_OR_PATH #Disconnect
iscsiadm -m node -T iSCSI_TARGET_IQN -p SERVER_IP -u
```

3. Execute the command **./stor stop** on each server where HBlock is located.
4. To restart the HBlock service, you can use one of the following methods:
Execute the command **./stor start** on each server where HBlock is located.
Directly restart all servers where HBlock is located: You can see **"Can the HBlock service start automatically after the server is restarted?"**, ensure that all services of HBlock are automatically restarted when the server is restarted.

Q: How to configure HBlock access permissions?

A: To enhance the security of HBlock, you can configure the firewall permissions to limit the access source IP of the iSCSI port (for example, the iSCSI port is 3260). Please refer to the following steps:

1. Turn on the firewall: `systemctl start firewalld`.

2. Configure the Allow IP
 - For IPv4 addresses: `firewall-cmd --permanent --add-rich-rule=rule family=ipv4 source address=IP port protocol=tcp port=3260 accept.`
 - For IPv6 addresses: `firewall-cmd --permanent --add-rich-rule="rule family=ipv6 source address=IP port protocol=tcp port=3260 accept".`
3. Restart the Firewall: **`firewall-cmd --reload.`**
4. Automatically Starts at Boot: `systemctl enable firewalld.service.`

Q: When configuring the Linux cluster mode, how to change the timeout period for the client to connect to HBlock?

A: You can change the timeout period between client and HBlock by modifying the `/etc/iscsi/iscsid.conf` and `/etc/multipath.conf` configuration files.

1. iSCSI: `/etc/iscsi/iscsid.conf`

- `node.conn[0].timeo.noop_out_interval`
The interval of iSCSI initiator sending **NOP-OUT** request to the target, the default value is 5, the unit is second.
- `node.conn[0].timeo.noop_out_timeout`
Timeout of iSCSI initiator receiving **NOP-IN** response after sending **NOP-OUT**, the default value is 5, the unit is second.
- `node.session.timeo.replacement_timeout`
After reaching `noop_out_timeout`, the timeout of iSCSI initiator waiting for iSCSI session rebuilding before returning I/O failure to SCSI or multipath layer.
The default value is 200, the unit is second. The optional value is as follows:
 - 0, return I/O failure immediately, no waiting for session rebuilding.
 - <0, wait all the time, until the session is rebuilt successfully, or iscsi session logout executed by users.
 - >0, return I/O failure after waiting for specified time.

Check `replacement_timeout`:

```
cat /sys/class/iscsi_session/sessionXX/recovery_tmo
```

Modify `replacement_timeout`:

```
echo X > /sys/class/iscsi_session/sessionXX/recovery_tmo
```

Note: If the iSCSI device is configured with multipath, the `/sys/class/iscsi_session/sessionXX/recovery_tmo` will be covered by the `fast_io_fail_tmo` in the `multipath.conf`.

2. multipath: `/etc/multipath.conf`

- `polling_interval`

Interval of path detection, the default value is 5, the unit is second. When the path is normal, the interval will be doubled every time, and added to $4 * \text{polling_interval}$.

The path detection mode depends on `path_checker`, generally, send command **TUR**, or **read**

Sector 0.

- `fast_io_fail_tmo`

After the SCSI transmission layer is faulty, the timeout of waiting for transmission layer rebuilding before returning the I/O failure to the multipath layer, the default value is 5, the unit is second. The optional value is as follows:

- `>=0`, timeout, will cover `/sys/class/iscsi_session/sessionXX/recovery_tmo`
- `off`, not modify `/sys/class/iscsi_session/sessionXX/recovery_tmo`

- `no_path_retry`

After a path fails, the number of I/O retry, the default value is 0. The optional value is as follows:

- `0/fail`: Immediately fail.
- `>0`: Retry number.
- `queue`: Retry all the time. Run command **multipath -ll** to view device information, it will display features `"1 queue_if_no_path"`.

If `no_path_retry=queue`, and the corresponding path fails, the SCSI command that has been processed on this path will be blocked, until the path recovers, that is, the upper layer application cannot aware of I/O failure. For application scenarios in which the upper layer is cluster, it may affect the failure detection and failover of cluster.

For the SCSI command that has been blocked, you can execute `dmsetup message mpathX 0 "fail_if_no_path"`, so that the SCSI command on this path returns failure immediately, to avoid the upper layer application from waiting indefinitely.

Q: What should I do if the client is disconnected from the HBlock server?

A: For Windows Clients, you need to go offline before disconnecting. Disconnect the backup connection first, and then disconnect the main connection, otherwise data may be lost.

For Linux Clients, you need to execute the command `sync` before disconnecting, otherwise data may be lost.

Q: If the iSCSI Initiator has already established a connection with an iSCSI target, HBlock will then create a new LUN and add it to the iSCSI target. How does the iSCSI Initiator discover a new LUN without disconnecting the existing connection?

A: Depending on the client, you can use the following methods to discover new LUNs:

- Windows: In Server Server Manager > File and Storage Services > Volumes > Disks, click Refresh to complete the addition of LUN.
- Linux: Before mounting the new LUN, you need to execute the following command on the Linux Client:

```
rescan-scsi-bus.sh # Before using this command, the system needs to install sg3_utils
```

Example: The client has connected to lun01, lun01, and lun01-a both correspond to target01, and the Linux Client needs to mount lun01-a.

- Server: Query the LUN.

```
[root@hblockserver CTYUN_HB1lock_Plus_3.10.0_x64]# ./stor lun ls
```

No.	LUN Name	Storage Mode	Capacity	Local Storage Class	Minimum Replicas	Status	Target	Snapshot Count	Snapshot Size	Is Clone
1.	lun01 (LUN 0)	Local	110 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.72:3260,Active) iqn.2012-08.cn.ctyunapi.oos:target01.2(192.168.0.209:3260,Standby)	2	471.36 MiB	
2.	lun01-a (LUN 0)	Local	110 GiB	EC 2+1+16 KiB	2	Normal	iqn.2012-08.cn.ctyunapi.oos:target01.1(192.168.0.72:3260,Active) iqn.2012-08.cn.ctyunapi.oos:target01.2(192.168.0.209:3260,Standby)	1	4 KiB	clone

- Client:

1. Before executing the command `rescan-scsi-bus.sh`:

```
[root@client ~]# ls SCSI
[8:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sda
[9:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sdb
[root@client ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 1000G 0 disk
├─mpathh 252:0 0 1000G 0 mpath
│ └─mpathh1 252:1 0 1000G 0 part /mnt/disk_mpathh
sdb 8:16 0 1000G 0 disk
├─mpathh 252:0 0 1000G 0 mpath
│ └─mpathh1 252:1 0 1000G 0 part /mnt/disk_mpathh
vda 253:0 0 40G 0 disk
├─vda1 253:1 0 4G 0 part [SWAP]
└─vda2 253:2 0 36G 0 part /
vdb 253:16 0 1T 0 disk
├─vdb1 253:17 0 1024G 0 part
vdc 253:32 0 1T 0 disk
├─vdc1 253:33 0 1024G 0 part /mnt/storage01
```

2. Execute `rescan-scsi-bus.sh` and view

```
[root@client ~]# ls SCSI
[8:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sda
[9:0:0:0] disk CTYUN iSCSI LUN Device 1.00 /dev/sdb
[root@client ~]# rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 8 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning for device 8 0 0 0 ...
OLD: Host: scsi8 Channel: 00 Id: 00 Lun: 00
Vendor: CTYUN Model: iSCSI LUN Device Rev: 1.00
Type: Direct-Access ANSI SCSI revision: 06
Scanning for device 8 0 0 1 ...
```

```

NEW: Host: scsi8 Channel: 00 Id: 00 Lun: 01
    Vendor: CTYUN    Model: iSCSI LUN Device Rev: 1.00
    Type:   Direct-Access          ANSI SCSI revision: 06
Scanning host 9 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning for device 9 0 0 0 ...
OLD: Host: scsi9 Channel: 00 Id: 00 Lun: 00
    Vendor: CTYUN    Model: iSCSI LUN Device Rev: 1.00
    Type:   Direct-Access          ANSI SCSI revision: 06
..... Scanning for device 9 0 0 1 ...
NEW: Host: scsi9 Channel: 00 Id: 00 Lun: 01
    Vendor: CTYUN    Model: iSCSI LUN Device Rev: 1.00
    Type:   Direct-Access          ANSI SCSI revision: 06
2 new or changed device(s) found.
    [8:0:0:1]
    [9:0:0:1]
0 remapped or resized device(s) found.
0 device(s) removed.
[root@client ~]# ls SCSI
[8:0:0:0]   disk   CTYUN    iSCSI LUN Device 1.00  /dev/sda
[8:0:0:1]   disk   CTYUN    iSCSI LUN Device 1.00  /dev/sdc
[9:0:0:0]   disk   CTYUN    iSCSI LUN Device 1.00  /dev/sdb
[9:0:0:1]   disk   CTYUN    iSCSI LUN Device 1.00  /dev/sdd
[root@client ~]# lsblk
NAME                MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda                  8:0      0 1000G 0 disk
├─mpathh             252:0     0 1000G 0 mpath
│ └─mpathh1          252:1     0 1000G 0 part /mnt/disk_mpathh
sdb                  8:16     0 1000G 0 disk
├─mpathh             252:0     0 1000G 0 mpath
│ └─mpathh1          252:1     0 1000G 0 part /mnt/disk_mpathh
sdc                  8:32     0  101G 0 disk
├─mpathi             252:2     0   101G 0 mpath
sdd                  8:48     0  101G 0 disk
├─mpathi             252:2     0   101G 0 mpath
vda                  253:0     0    40G 0 disk
├─vda1               253:1     0     4G 0 part [SWAP]
├─vda2               253:2     0    36G 0 part /
vdb                  253:16    0     1T 0 disk
├─vdb1               253:17    0 1024G 0 part
vdc                  253:32    0     1T 0 disk
├─vdc1               253:33    0 1024G 0 part /mnt/storage01

```

At this time, you can mount the newly created LUN according to the normal steps for mounting a LUN.

Q: How to ensure that the Linux Client can directly mount the LUN created by HBlock after restarting the server?

A: After mounting the LUN created by HBlock to the client, please refer to the following steps:

Note:

- If the `/etc/iscsi/iscsid.conf` file has `node.startup=manual` configured, you can use either of the following methods:
 - Modify the `/etc/iscsi/iscsid.conf` file to set `node.startup=automatic`. Then, according to the `/etc/multipath.conf` settings, proceed with the mounting steps.
 - For each target configuration, automatically log in upon client reboot by running the command: `'iscsiadm -m node -T Target-IQN -p SERVER_IP --op update -n node.startup -v automatic'`. Then, according to the `/etc/multipath.conf` settings, proceed with the mounting steps.
- If the `/etc/iscsi/iscsid.conf` file has `node.startup=automatic` configured, according to the `/etc/multipath.conf` settings, proceed with the mounting steps.

If the `/etc/multipath.conf` file has `'user_friendly_names yes'` configured, follow these steps:

1. On the client, use the command `'lsblk -f'` to view the file system information of the mounted device and find the UUID corresponding to the file system:

```
[root@client ~]# lsblk -f
NAME            FSTYPE          LABEL UUID                                MOUNTPOINT
sda             mpath_member
├─mpatha
│   └─mpatha1 ext4              7269eef6-e401-454a-acb4-503d33337f21 /mnt/disk_mpatha
sdb             mpath_member
├─mpatha
│   └─mpatha1 ext4              7269eef6-e401-454a-acb4-503d33337f21 /mnt/disk_mpatha
vda
├─vda1          swap            9e33bd6f-c68c-41c7-95c8-703f4fe8c3d4 [SWAP]
├─vda2          xfs             a83f4fdc-2ea1-4fec-a1e2-a42016ce0afe /
vdb
├─vdb1          ext4            74296a9e-8cfd-4708-89b1-08086f71175b
vdc
├─vdc1          ext4            a9fedea4-391e-4d2a-8824-c9a3a6853394
```

2. Add the LUN mounting information created by HBlock in the `/etc/fstab` file, and the LUN can be automatically mounted the next time you start the computer.

```
UUID=7269eef6-e401-454a-acb4-503d33337f21 /mnt/disk_mpatha ext4 defaults,_netdev 0 0
```

If the `/etc/multipath.conf` file has `'user_friendly_names no'` configured, follow these steps:

1. On the client, use the command `'lsblk -f'` to view the file system information of the mounted devices.

```
[root@client ~]# lsblk -f
NAME                FSTYPE      FSVER          LABEL          UUID                                 FSAVAIL FSUSE% MOUNTPOINTS
sda                  mpath_member
└─0x300000001e908caf ext4          1.0             9aa0f703-f0db-4db7-bda5-9989aaba0269 615.6G   0% /mnt/disk_lun06a
sdb                  mpath_member
└─0x300000001e908caf ext4          1.0             9aa0f703-f0db-4db7-bda5-9989aaba0269 615.6G   0% /mnt/disk_lun06a
sr0                  iso9660      Joliet Extension config-2 2025-01-08-11-16-53-00
vda
└─vda1               xfs          a74fe630-4089-4ae3-b9f8-95a915cf42d7 34.1G    15% /
vdb                  ext4          1.0             f7728367-9021-4225-9fe6-129c36a4766a 91.5G    1% /mnt/storage01
```

2. Add the LUN mounting information created by HBlock in the `/etc/fstab` file, and the LUN can be automatically mounted the next time you start the computer.

```
/dev/mapper/0x300000001e908caf /mnt/disk_lun06a ext4 defaults,_netdev 0 0
```

Q: How to set up the disk path used by HBlock to be automatically mounted on startup in the server?

A: After the server uses the `mount` command to mount the disk path, you can refer to the following steps to automatically mount it on startup:

1. Use the command `lsblk -f` on the client to query the file system information of the mounted device and find the UUID corresponding to the file system:

```
[root@server ~]# lsblk -f
NAME  FSTYPE LABEL UUID                                 MOUNTPOINT
vda
└─vda1 swap          9e33bd6f-c68c-41c7-95c8-703f4fe8c3d4 [SWAP]
└─vda2 xfs           a83f4fdc-2ea1-4fec-a1e2-a42016ce0afe /
vdb
└─vdb1 ext4          c62d513e-c3cf-4719-b15c-4366e4b52664
vdc
└─vdc1 ext4          1c47025a-6028-42ce-90aa-59d6f5106818 /mnt/storage01
```

2. Add the mounting disk path information in the `/etc/fstab` file, and the disk path will be automatically mounted the next time you start the computer.

```
UUID=1c47025a-6028-42ce-90aa-59d6f5106818 /mnt/storage01 ext4 defaults 1 1
```

Q: On the Linux Client, how to delete the iSCSI target that is not connected to the server?

A:

- HBlock is not uninstalled, and the Linux Client deletes the iSCSI target connected to the server:

1. Use the following command on the Linux Client to remove the disk disconnect:

```
umount DIRECTORY_NAME_OR_PATH
iscsiadm -m node -T iSCSI_TARGET_IQN -p SERVER_IP -u
```

2. Linux Clients use one of the following commands to delete an iSCSI target that is disconnected from the server:

- Delete all iSCSI targets not connected to the server.

```
iscsiadm -m node --op delete
```

- Delete the specified iSCSI target that is not connected to the server

```
iscsiadm -m node --targetname iSCSI_TARGET_IQN -p SERVER_IP:port -o delete
```

- After uninstalling HBlock and restarting the Linux Client, you can use the following command to delete the **iSCSI target** that is not connected to the server:

- Delete all iSCSI targets not connected to the server.

```
iscsiadm -m node --op delete
```

- Delete the specified iSCSI target that is not connected to the server.

```
iscsiadm -m node --targetname iSCSI_TARGET_IQN -p SERVER_IP:port -o delete
```

Q: On the Linux Client, how to confirm the correspondence between the drive letter and the HBlock LUN?

A: You can use the following command to query

```
sg_inq /dev/sdX # will display the LUN name in the output information  
udevadm info --query=all --path=/block/sdx  
#The output information will display the Target and LUN numbers.
```

OR

```
iscsiadm -m session -P 3 #Display Target portal and corresponding drive letter
```

OR

```
ll /dev/disk/by-path #List all disks
```

Q: How to adjust the information of the HBlock running log, such as the number of logs, log file size, etc.?

A: You can adjust related information by editing the parameters in the configuration file `log4j2.properties`. The configuration files corresponding to different types of running logs are as follows:

- System log and configuration log:
`HBlock_installation_path/apps/base_version/conf/log4j2.properties`.
- Data log: `HBlock_installation_path/apps/ds/etc/dp/log4j2.properties`.
- Coordination log: `HBlock_installation_path/apps/cs/conf/log4j2.properties`.

Note:

- If HBlock has been initialized, restart HBlock to take effect after adjusting parameters.
- After the log parameters are modified, the storage space occupied by the log file may increase. Please ensure that the disk where the HBlock installation path is located has enough space, otherwise it may affect service.

Q: Some services of HBlock failed to restart. How to troubleshoot?

A: Please check whether the port used by the service is occupied by other applications. The port values used by the HBlock service can be queried by using the command `./stor server ls --port`. Please prevent other services in the system from occupying these port numbers. And make sure that the Linux local ephemeral port (`ip_local_port_range`) range does not include the port number used by the HBlock service. Use the command `cat /proc/sys/net/ipv4/ip_local_port_range` to query the local temporary port range.

Q: After both the HBlock server and client are restarted, if the server starts with a delay compared to the client, causing the client to fail to automatically reconnect, how should this be handled?

A: This issue is caused by the HBlock server restarting later than the client. Try to ensure that the HBlock client restarts after the HBlock server. If this cannot be avoided, you can use the following methods to configure it:

- For newly created target: Modify the login timeout retry count (`node.session.initial_login_retry_max`) for `iscsiadm` in `/etc/iscsi/iscsid.conf`, for example:

```
node.session.initial_login_retry_max=172800
```

The default restart interval is 15 seconds, you can set the number of retries according to their actual needs.

- For discovered targets: Modify the login timeout retry count (`node.session.initial_login_retry_max`) in `/etc/iscsi/nodes/iqn.2012-08.cn.ctyunapi.oos/IP,PORT/default` where `IP,PORT` represent the target's IP address and iSCSI port number respectively.

```
node.session.initial_login_retry_max=172800
```

The default restart interval is 15 seconds, you can set the number of retries according to their actual needs.

Note: The paths for the above two configuration files may differ depending on the system.

Q: When a client has mounted an HBlock LUN and formatted it with the XFS file system, and a clone LUN of this LUN is created on the HBlock server side. If the client mounting the source LUN also attempts to mount the clone LUN but the mounting fails, how to resolve this issue?

A: There are two ways to resolve this issue

- Modify the UUID of the clone LUN on the client and then remount it.

```
xfs_admin -U generate /dev/sdX
mount /dev/sdX PATH
```

- Mount the clone LUN with the nouuid option to ignore UUID checks.

```
mount -o nouuid /dev/sdX PATH
```

Note: /dev/sdX is the device of the clone LUN, PATH is the disk path.

Example: sda is the source volume. sdj and sdo are the cloned volumes. Mounting them directly causes errors. sdo is mounted by changing the UUID, and sdj is mounted by ignoring the UUID check.

```
[root@client /]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda          8:0    0  200G  0 disk /mnt/sda
sdj          8:144  0   211G  0 disk
sdo          8:224  0   233G  0 disk
vda         253:0   0    40G  0 disk
└─vda1      253:1   0    40G  0 part /
vdb         253:16  0   100G  0 disk /mnt/local_stor01
vdc         253:32  0   100G  0 disk
vdd         253:48  0   100G  0 disk
[root@client /]# mount /dev/sdo /mnt/sdo
mount: wrong fs type, bad option, bad superblock on /dev/sdo,
       missing codepage or helper program, or other error

       In some cases useful info is found in syslog - try
       dmesg | tail or so.
[root@client /]# xfs_admin -U generate /dev/sdo
Clearing log and setting UUID
writing all SBs
new UUID = fe99e47e-8091-4a12-860c-f60e02b07b41
[root@client /]# mount /dev/sdo /mnt/sdo
[root@client /]# mount /dev/sdj /mnt/sdj
mount: wrong fs type, bad option, bad superblock on /dev/sdj,
       missing codepage or helper program, or other error

       In some cases useful info is found in syslog - try
       dmesg | tail or so.
[root@client /]# mount -o nouuid /dev/sdj /mnt/sdj
[root@client /]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda          8:0    0  200G  0 disk /mnt/sda
sdg          8:96   0   555G  0 disk /mnt/sdg
sdj          8:144  0   211G  0 disk /mnt/sdj
```

```
vda      253:0    0   40G  0 disk
└─vda1   253:1    0   40G  0 part /
vdb      253:16   0  100G  0 disk /mnt/local_stor01
vdc      253:32   0  100G  0 disk
vdd      253:48   0  100G  0 disk
```

5.3 Upgrade and Expand Capacity

Q: How to upgrade the HBlock software?

A: Under the premise of meeting the upgrade requirements, follow the steps in **Upgrade HBlock** to complete the upgrade operation. During the upgrade process, HBlock query operations can be performed, but HBlock modification and deletion operations cannot be performed, but it will not affect the customer's normal read and write business. Upgrading requires a period of time to execute. If the upgrade fails, please follow the prompts to resolve the issue as soon as possible and restart the upgrade request. It is not recommended to run servers of different versions for a long time. If you encounter any problems during the upgrade and need support, please contact eSurfing Cloud staff.

Q: What expansion methods does HBlock support?

A: HBlock supports expanding local storage space by expanding new storage servers and expanding hard disks in existing nodes. And the number of disks on each server is not required to be the same.

Q: When HBlock expands a node, does the configuration of the new node need to be the same as the original node?

A: It can be different. As long as the new expansion node meets the minimum configuration requirements. The CPU, memory, number of hard disks, and hard disk size of each storage server node can be different.

Q: To what extent can HBlock's automated deployment be achieved?

A: The operating system, virtual machine, network configuration, and other work on the server need to be installed with the help of third-party tools or completed manually. After completing the above operations, the HBlock software can be installed through the command line. You need to execute the install installation command on all servers to be installed, and then execute the initialization command `setup` on one of the servers.

5.4 Data Migration

Q: Does HBlock support data migration between traditional SAN storage products?

A: Data migration between HBlock and other SAN products can be completed through the virtualization platform. For non-virtualized scenarios, such as databases, special data migration tools can be used to migrate data between SAN and HBlock. In addition, the Linux kernel module **dm-clone** can be used to clone block devices, switch data from the original LUN to a new LUN, and realize data migration from the original SAN storage to HBlock. Windows operating systems can [Getting Started with the User State Migration Tool \(USMT\)](#) to implement data migration.

5.5 Data Security

Q: If the disk is damaged, will it cause data loss?

A: Since HBlock supports multiple replicas and erasure coding redundant data modes, data can be stored in different fault domains through sharding. When a disk is damaged, it can be retrieved from other disks in the same fault domain or in different fault domains. Automatically recover data on the Internet, greatly reducing the risk of data loss. However, because the data is stored in the user's local server, the user needs to ensure the stable operation of the local server and disk. When there are too many disk damages in the system, there is still a risk of data loss.

Q: If there is a power outage or the HBlock server is down during the transmission process, will it cause data loss?

A: HBlock supports the cluster mode and can be deployed on multiple servers. When the client correctly configures MPIO, when a single point of fault occurs, the service will be automatically switched to the backup node and data can be read or written normally.

Q: What safeguards are in place for HBlock's data security, integrity and reliability?

A: HBlock supports multi-replica and erasure code data redundancy policies, and uses timing transmission, double check, and data caching mechanisms to ensure data security and integrity as much as possible. At the same time, HBlock supports the cluster mode, which can quickly implement failover and ensure the continuity of the user's business as much as possible.

Q: What solutions are there for unexpected events such as power outages and disk damage?

A: If the standalone mode is deployed, you are advised to configure uninterruptible power supplies (UPS) and hard disks as RAID to ensure data reliability. When the cluster mode is deployed, if the configuration is correct, a single point of failure in the server, such as power failure or disk damage, will be automatically switched over.

Q: Can the servers in the cluster be down for a long time or not start the HBlock service?

A: It is not recommended that the server is down for a long time or does not start the HBlock server.

Because when the server is restarted, it may take a long time to recover, which may affect the business. Therefore, it is recommended that if a server in the cluster is down, restart or repair it as soon as possible and start the HBlock service.

5.6 Technical Support

Q: If you encounter problems after using HBlock, will someone come to solve it?

A: If you have any questions during installation and use, you can contact us through the following methods:

- eSurfing Cloud official website consultation hotline 400-810-9889.
- eSurfing Cloud official website work order management system.

If you need remote assistance:

- HBlock has been installed but not initialized: After contacting eSurfing Cloud staff, according to the HOST and PORT provided by the eSurfing Cloud staff (the default port number is 18100), use the command. `./stor config set { -i | --item } remote {-H | --host} HOST {-P | --port} PORT { -s | --status } STATUS` to turn on remote assistance. Then inform eSurfing Cloud staff of the remote assistance code returned by the command. At this time, eSurfing Cloud staff can provide remote assistance.
- Completed installation and initialization of HBlock: You can operate in the same manner as when HBlock has been installed but not initialized. You can also first obtain the HBlock ID through `./stor info --stor-id`, and send the question and HBlock ID to eSurfing Cloud staff. And according to the HOST and PORT provided by eSurfing Cloud staff (the default port number is 18100), use the command `./stor config set { -i / --item } remote {-H / --host} HOST {-P / --port} PORT { -s / --status } STATUS` turns on the remote assistance function. At this time, eSurfing Cloud staff can provide remote assistance.

Remote Assistance is disabled by default and users can enable or disable this feature at any time. Once enabled, eSurfing Cloud staff have the right to log in to the HBlock system in the user environment to diagnose problems. When an eSurfing Cloud staff member logs in for remote assistance, they have the permissions of the following two users: users who enable remote collaboration operations and users who install HBlock. Users can view all operations performed by staff during remote assistance through `logs/remotearchive/remotearchive.log`.

Note: If the user enables the remote assistance function, it means that the user trusts eSurfing Cloud staff and authorizes eSurfing Cloud staff to access all data in the HBlock system. The staff of eSurfing Cloud will try their best to diagnose the problem and ensure data security, but due to the complexity of the system environment, we are not responsible for any consequences caused by remote assistance.

5.7 Purchase Guide

Q: How to obtain a software license?

A: HBlock provides a 30-day trial period. After expiration, please contact the HBlock software supplier to obtain a software license. When obtaining, you need to provide the HBlock serial number (can be obtained through the command `./stor info --serial-id`). After obtaining the software license, execute `./stor license add { -k / --key } KEY` to load the software license.

5.8 Other FAQs

Q: Does HBlock support creating snapshots for LUNs?

A: HBlock currently does not support creating snapshots for LUNs, but users can create snapshots in the following ways:

- If the upper-layer application is virtualization or container technology, you can use the snapshot technology that comes with the hypervisor or container to create a snapshot.
- If HBlock directly connects to the operating system, you can use its own **snapshot and backup technology** for Windows, and for Linux, you can use **btrfs** to create snapshots. Compared with creating snapshots directly on block devices, the snapshot technology provided by the operating system provides better support for data integrity.

Q: Can virtualized applications be deployed?

A: Yes. Virtualization applications that support user deployment include: KVM and VMware.

Q: Can database applications be deployed?

A: Yes. Database applications that support deployment include: Oracle (except Oracle rac), MySQL, SQL Server, PostgreSQL, MongoDB, and DB2.

Q: How to quickly determine the network connectivity status of each server in the cluster?

A: In addition to using the traditional method of detecting the network connectivity status between servers, when ICMP is not disabled in the inter-cluster network, you can also view the recent communication status between servers by viewing network logs. Log address:
HBlock_installation_directory/logs/network/network.log.

Q: Is NAT access supported?

A: Yes. Usually, the iSCSI initiator accesses the HBlock target through the intranet. If Network Address Translation (NAT) is configured on the intranet router, the iSCSI initiator can connect to the server where the target is located through the external IP of the NAT, thereby providing iSCSI as a cloud service remotely through HBlock.

Q: Why is the disk path utilization not equal to disk path used capacity/disk path total capacity in the system? Why is the disk path quota utilization not equal to the disk path used quota/disk path quota?

A: The disk path utilization of the HBlock system and server refers to the average value of the usage rates of all disk paths. The disk path quota utilization of the HBlock system and server

refers to the average value of the quota usage rates of all disk paths. For details, see Path_Rate and Path_Cap_Quota_Rate in **Monitoring Metrics**.

Q: When exporting a backup file, if `--allow-modify` (allow the current backup task to overwrite an existing backup file with the same name) is specified, will it uniformly overwrite the original files?

A: The system decides whether to overwrite the original file on a case-by-case basis:

- If a same-named file exists but was not generated by an HBlock export task, it cannot be modified.
- If a file with the same name exists and it is an HBlock-exported backup:
 - If compression is enabled for this export, the backup file will be re-exported from the beginning and overwrite the existing file.
 - If compression is not enabled for this export:
 - ◆ If the existing backup file is compressed, it will be re-exported from the beginning and overwrite the existing file.
 - ◆ If the existing backup file is not compressed, the system will determine the point of interruption based on the size of the existing file, resume the export from that point, and continue writing to the same file.

Note: The system cannot determine whether the backup data in a same-named file corresponds to the same backup task as the current export. In other words, the data may belong to different LUNs or be exported from different snapshots, potentially resulting in corrupted backup data. Always ensure that the same backup task uses the same file name.

6 Troubleshooting

How to manually uninstall the HBlock server?

If the HBlock server fails and cannot be uninstalled automatically, you can uninstall HBlock through manual method. The steps are as follows:

1. Use the command `ps -aux/grep stor` to query the HBlock process.
2. According to the found process ID, use the command `kill Stor_process_ID` to delete the Stor process.
3. Use the command `rm -rf Stor_installation_directory` to delete the Stor related installation files.
4. Use the command `rm -rf Stor_disk_path` to delete the Stor data.

The client mounts the LUN created by HBlock through NFS. When reading and writing data, a file verification error occurs. How to deal with it?

A file verification error occurs, which may be related to NFS mount parameters or client implementation. Please try to refresh the NFS Client cache and then read the data again.

7 Appendix

7.1 HBlock Service

Service	Service Name	Effect
stor:mdm	Metadata management services (only supported by cluster mode)	Manage metadata for the entire system.
stor:fc	Failover control service (only supported by cluster mode)	Perform system health check and implement failover control.
stor:ls	Ledger service (only supported by cluster mode)	Provides log-based data synchronization function.
stor:ds-x	Data service (only supported by cluster mode)	Manage user's file data blocks.
stor:cs	Coordination service (only supported by cluster mode)	Monitor the status of each server and trigger notification events to ensure high availability of cluster services.
stor:ms	Management services	Process request information and maintain cluster operating status.
stor:ws	Watchdog service	Monitor the status of each service and be responsible for starting the service.
stor:ps	Processor service	Responsible for iSCSI protocol parsing and data storage.
stor:ag	Aggregator service	Responsible for acquiring performance data.
stor:ua	Upgrade agent service	Responsible for receiving upgrade requests and performing upgrade-related operations.

7.2 User Event List

Server

Event	Description
AddServer	Add server
RemoveServer	Recover server
SetServer	Set server properties
DeleteTargetPortalIP	Delete server target portal IP
RestartService	Restart service
AddPath	Add disk path
RemovePath	Remove disk path
StartService	Start service
StopService	Stop service
SetPath	Edit disk path
MigrateService	Migrate service

iSCSI target

Event	Description
CreateTarget	Create target
DeleteTarget	Delete target
SetTarget	Set target property
MigrateTarget	Migrate target
DeleteCHAP	Delete CHAP
DeleteConnection	Delete Connection
DeletetargetAllowlist	Delete target allow list
SettargetAllowlist	Set target allow list

LUN

Event	Description
CreateLUN	Create a LUN
DeleteLUN	Delete a LUN
SetLUN	Set LUN properties
ExpandLUN	Expand a LUN
SwitchLUN	LUN active/backup switching
RecoverLUN	Recover a LUN
ResumeLUNRecovery	Resume recovery of a previously failed LUN
SetBatchLUN	Set LUNs in batch

Event	Description
CreateCloneLUN	Create Clone LUN
FlattenCloneLUN	Flatten Clone LUN
LUNDataResidue	LUN Data Residue
WipeLUN	Wipe LUN
DeleteLUNXattr	Delete LUN xattr
SetLUNXattr	Set LUN xattr
SuspendLUN	Suspend LUN
ResumeLUN	Resume LUN

System

Event	Description
Login	Log in
SetMailConfig	Set up email notification
DeleteMailConfig	Delete email notification
SendTestMail	Send test email
SetRemoteAccess	Set up remote assistance
DeleteRemoteAccess	Delete remote assistance
ImportLicense	Import software license
SetPassword	Set password
StartLogCollect	Initiate a log collection request
DeleteLogCollect	Delete log collection request
SetAlarmMuteStatus	Edit alarm silent status
ManuallyResolveAlarm	Manually Resolve the Alarm
Setup	Initialize
AddMonitorConfig	Add monitor config
SetMonitorConfig	Set monitor config
DeleteMonitorConfig	Delete monitor config

Cluster Topology

Event	Description
SetNode	Modify node information
CreateNode	Create a node
DeleteNode	Delete a node

Storage Pool

Event	Description
CreateStoragePool	Create a storage pool
DeleteStoragePool	Delete a storage pool
SetStoragePool	Modify storage pool information
AddNodeToPool	Add a node to a pool
RemoveNodeFromPool	Remove a node from a pool

Upgrade

Event	Description
StartUpgrade	Start the upgrade

Snapshot

Event	Description
CreateSnapshot	Create Snapshot
SetSnapshot	Set Snapshot
DeleteSnapshot	Delete Snapshot
RollbackSnapshot	Rollback Snapshot
CreateConsistencySnapshot	Create Consistency Snapshot
SetConsistencySnapshot	Set Consistency Snapshot
RollbackConsistencySnapshot	Rollback Consistency Snapshot
DeleteConsistencySnapshot	Delete Consistency Snapshot

QoS Policy

Event	Description
CreateQoSPolicy	Create QoS Policy
SetQoSPolicy	Set QoS Policy
DeleteQoSPolicy	Delete QoS Policy
AssociateQoSPolicy	Associate QoS Policy
DisassociateQoSPolicy	Disassociate QoS Policy

Backup

Event	Description
ExportBackup	Export Backup
ImportBackup	Import Backup

7.3 System Event List

Server

Event	Description
ServiceUnavailable	Service unavailable
ServiceAvailable	Service available
ServerAdded	Server added
ServerRemoved	Server removed
ProtocolServiceAbnormal	Protocol parsing service exception
ProtocolServiceResumed	Protocol parsing service resumed
InsufficientSpaceonInstallationPath	The remaining space of the installation directory is insufficient.
SpaceonInstallationPath	The remaining space of the installation directory meets service running requirements
BaseServiceAbnormal	Base service exception
BaseServiceResumed	Base service resumed
ServiceMigrated	Service migration completed
ServiceMigrateAbnormal	Service migration exception
InsufficientSpaceonMetaDir	Insufficient space on the disk where the base service data directory is located
SufficientSpaceonMetaDir	Sufficient space on the disk where the base service data directory is located
InsufficientMemory	Insufficient memory
ServiceRestartRequired	Service restart required

Disk

Event	Description
DiskIOError	Disk I/O error
DiskIOResumed	Disk I/O resumed
DiskWriteSpeedTooSlow	Disk write speed is slow
DiskWriteSpeedResumed	Disk write speed returns to normal
PathAdded	Add disk path
PathRemoved	Remove disk path
CapacityQuotaUsageExceedsThreshold	Quota usage exceeds threshold
CapacityQuotaUsageBelowThreshold	Quota usage returns to normal
CapacityQuotaUsageApproachLimit	Quota exhausted
CapacityQuotaUsageBelowLimit	Quota usage rate is lower than the upper limit
DiskUsageExceedsThreshold	Disk usage exceeds threshold
DiskUsageBelowThreshold	Disk usage returns to normal
DiskPathHealthStatusWarning	Disk path health status warning
DiskPathHealthStatusError	Disk path health status error
DiskPathHealthStatusResumed	Disk path health status resumed

DataServiceHealthStatusWarning	Data service health status warning
DataServiceHealthStatusError	Data service health status error
DataServiceHealthStatusResumed	Data service health status resumed

LUN

Event	Description
ActiveStandbySwitched	LUN active/standby switching
InsufficientFDForLUNToWrite	The number of available fault domains does not meet the LUN write requirements
SufficientFDForLUNToWrite	Number of available fault domains meets LUN write requirements
LUNRecovered	LUN recovered
CannotConnectToCloud	Cannot connect to cloud
ConnectWithCloudResumed	Connect with cloud resumed
CloudAccountAbnormal	Cloud account abnormal
CloudAccountNormal	Cloud account normal
LUNCloudDataConflict	LUN cloud data conflict
LUNCloudDataConflictResolved	LUN cloud data conflict resolved
LUNCloudHeartbeatConflict	LUN cloud heartbeat conflict
LUNCloudHeartbeatNormal	LUN cloud heartbeat normal
LUNFlattened	LUN flattened
LUNResumed	LUN resumed
LUNDataLowRedundancy	LUN data low redundancy
LUNDataRedundancyResumed	LUN data redundancy resumed
LUNDataAccessResumed	LUN data access resumed
LUNDataAccessFailed	LUN data access failed

Target

Event	Description
InitiatorConnectionFailed	Initiator connection failed
InitiatorConnectionNormal	Initiator connection normal

System

Event	Description
ReachLicenseMaxCapacity	License capacity reached upper limit
LicenseMaintenanceExpired	License out of warranty
LicenseExpired	License expired
LicenseImported	License import
DataLowRedundancy	Data low redundancy
DataRedundancyResumed	Data redundancy resumed
DataResumed	Data resumed

DataBalanceStart	Data balancing begin
DataBalanceProgress	Data balancing process
DataBalanceFailed	Data balancing failed
DataBalanceEnd	Data balancing end
DataAccessFailed	Data cannot be accessed
DataAccessResumed	Data access resumed
ProTrialActive	Pro Trial active

Fault Domain Module

Event	Description
FaultDomainWarning	Fault domain status changes to warning
FaultDomainError	Fault domain status changes to Error
FaultDomainResumed	The fault domain status returns to normal

Storage Pool

Event	Description
CapacityQuotaUsageExceedsThreshold	Quota usage exceeds threshold
CapacityQuotaUsageBelowThreshold	Quota usage returns to normal
CapacityQuotaUsageApproachLimit	Quota exhausted
CapacityQuotaUsageBelowLimit	Quota is lower than the upper limit
DiskUsageExceedsThreshold	Disk usage exceeds threshold
DiskUsageBelowThreshold	Disk usage returns to normal
PoolDataLowRedundancy	Pool data low redundancy
PoolDataRedundancyResumed	Pool data redundancy resumed
PoolDataAccessResumed	Pool data access resumed
PoolDataAccessFailed	Pool data access failed

7.4 Monitoring Metrics

- **Data Granularity**

The granularities of monitoring metrics are "Fine" and "Coarse". The following table lists the meanings of the metrics.

Category	Data Granularity	Data Retention Period	Description
Fine	20 seconds	2 hours	Acquires real-time data every 20 seconds. One data point is generated, and each data point is retained for two hours.
	1 minute	6 hours	Based on data of 20s granularity, aggregates the data of 1-minute granularity and retains it for 6 hours.
	5 minutes	1 day	Based on data of 20s granularity, aggregates the data of 5-minute granularity and retains it for 1 day.
	1 hour	7 days	Based on data of 20s granularity, aggregates the data of 1-hour granularity and retains it for 7 days.
	1 day	1 year	Based on data of 20s granularity, aggregates the data of 1-day granularity and retains it for 1 year.
Coarse	5 minutes	2 hours	Based on data of 20s granularity, aggregates the data of 5-minute granularity and retains it for 2 hours.
	1 hour	1 day	Based on data of 20s granularity, aggregates the data of 1-hour granularity and retains it for 1 day.
	1 day	1 month	Based on data of 20s granularity, aggregates the data of 1-day granularity and retains it for 1 month.
	1 week	6 months	Based on data of 20s granularity, aggregates the data of 1-week granularity and retains it for 6 months.
	1 month	1 year	Based on data of 20s granularity, aggregates the data of 1-month granularity and retains it for 1 year.

- **Monitoring Metrics**

Monitoring Dimension	Monitoring Metric	Description	Unit	Data Granularity
system	IOPS	Total IOPS between the client and HBlock.	None	Fine
	R_IOPS	IOPS when the client reads data from HBlock.	None	Fine
	W_IOPS	IOPS when the client writes data to HBlock.	None	Fine

Bandwith	Total bandwidth between the client and HBlock.	B/s	Fine
R_Bandwith	Bandwidth when the client reads data from HBlock.	B/s	Fine
W_Bandwith	Bandwidth when the client writes data to HBlock.	B/s	Fine
Latency	Total latency between the client and HBlock. The average latency of read and write operations in the system within an acquisition cycle reflects the time that HBlock takes to process read and write requests.	ms	Fine
W_Latency	Latency when the client writes data to HBlock. The average latency of write operations in the system within an acquisition cycle reflects the time that HBlock takes to process write requests.	ms	Fine
R_Latency	Latency when the client reads data from HBlock. The average latency of read operations in the system within an acquisition cycle reflects the time that HBlock takes to process read requests.	ms	Fine
Path_Cap	Total capacity of disks where all HBlock disk paths reside in the system.	bytes	Coarse
Path_Used	Used capacity of disks where all HBlock disk paths reside in the system.	bytes	Coarse
Path_Rate	Average usage rate of disks where HBlock disk paths reside in the system, i.e., $\text{Path_Used}/\text{Path_Cap}$.	%	Coarse
Path_Cap_Quota	Available space for HBlock in the system, i.e., the total capacity quota allocated by the user to HBlock for all HBlock disk paths.	bytes	Coarse
Path_Cap_Quota_Used	Space occupied by HBlock data on disks where all HBlock disk paths reside in the system.	bytes	Coarse
Path_Cap_Quota_Rate	Average usage rate of HBlock disk path capacity quota in the system, i.e., $\text{Path_Cap_Quota_Used}/\text{Path_Cap_Quota}$.	%	Coarse
Cloud_Bandwidth	The total bandwidth between HBlock and the cloud.	B/s	Fine
Cloud_U_Bandwidth	Bandwidth when HBlock upload data to cloud.	B/s	Fine
Cloud_D_Bandwidth	Bandwidth when HBlock download data from cloud.	B/s	Fine

server	CPU_Rate	CPU usage rate of the server.	%	Fine
	Mem_Rate	Memory usage rate of the server.	%	Fine
	Mem_Total	Total memory of the server.	bytes	Fine
	Mem_Used	Memory usage of the server.	bytes	Fine
	IOPS	Total IOPS between the client and HBlock.	None	Fine
	R_IOPS	IOPS when the client reads data from HBlock.	None	Fine
	W_IOPS	IOPS when the client writes data to HBlock.	None	Fine
	Bandwith	Total bandwidth between the client and HBlock.	B/s	Fine
	R_Bandwith	Bandwidth when the client reads data from HBlock.	B/s	Fine
	W_Bandwith	Bandwidth when the client writes data to HBlock.	B/s	Fine
	Latency	Total latency between the client and HBlock. The average read-write latency of the server-associated LUNs during the data acquisition period.	ms	Fine
	W_Latency	Latency when the client writes data to HBlock. The average write latency of the server-associated LUNs during the acquisition period.	ms	Fine
	R_Latency	Latency when the client reads data from HBlock. The average read latency of the server-associated LUNs during the acquisition period.	ms	Fine
	Path_Cap	Total capacity of disks where all HBlock disk paths reside on the HBlock server.	bytes	Coarse
	Path_Used	Used capacity of disks where all HBlock disk paths reside on the HBlock server.	bytes	Coarse
	Path_Rate	Average usage rate of disks where HBlock disk paths reside on the HBlock server, i.e., Path_Used/Path_Cap.	%	Coarse
	Path_Cap_Quota	Available space for HBlock on the HBlock server, i.e., the total capacity quota allocated by the user to HBlock for all disk paths.	bytes	Coarse
	Path_Cap_Quota_Used	Space occupied by HBlock data on disks where all disk paths reside on the HBlock server.	bytes	Coarse
Path_Cap_Quota_Rate	Average usage rate of HBlock disk path capacity quota on the HBlock server, i.e., Path_Cap_Quota_Used/Path_Cap_Quota.	%	Coarse	
Cloud_Bandwidth	The total bandwidth between HBlock server and the cloud.	B/s	Fine	

	Cloud_U_Bandwidth	Bandwidth when HBlock server upload data to cloud.	B/s	Fine
	Cloud_D_Bandwidth	Bandwidth when HBlock server download data from cloud.	B/s	Fine
disk	Path_Cap	Total capacity of disks where HBlock disk paths reside.	bytes	Coarse
	Path_Used	Used capacity of disks where HBlock disk paths reside.	bytes	Coarse
	Path_Rate	Usage rate of disks where HBlock disk paths reside.	%	Coarse
	Path_Cap_Quota	Available space for HBlock, i.e., the capacity quota allocated by the user to HBlock.	bytes	Coarse
	Path_Cap_Quota_Used	Space occupied by HBlock data on disks where HBlock disk paths reside.	bytes	Coarse
	Path_Cap_Quota_Rate	Usage rate of HBlock disk path capacity quota, i.e., $\text{Path_Cap_Quota_Used}/\text{Path_Cap_Quota}$.	%	Coarse
LUN	IOPS	Total IOPS between the client and HBlock LUN.	None	Fine
	R_IOPS	IOPS when the client reads data from the HBlock LUN.	None	Fine
	W_IOPS	IOPS when the client writes data to the HBlock LUN.	None	Fine
	Bandwith	Total bandwidth between the client and the HBlock LUN.	B/s	Fine
	R_Bandwith	Bandwidth when the client reads data from the HBlock LUN.	B/s	Fine
	W_Bandwith	Bandwidth when the client writes data to the HBlock LUN.	B/s	Fine
	Latency	Total latency between the client and the HBlock LUN. The average latency of read and write operations of the LUN within an acquisition cycle reflects the time that the HBlock LUN takes to process read and write requests.	ms	Fine
	W_Latency	Latency when the client writes data to the HBlock LUN. The average latency of write operations of the LUN within an acquisition cycle reflects the time that the HBlock LUN takes to process write requests.	ms	Fine
R_Latency	Latency when the client reads data from the HBlock LUN. The average latency of read operations in the system within an	ms	Fine	

		acquisition cycle reflects the time that the HBlock LUN takes to process read requests.		
	Cloud_Bandwidth	The total bandwidth between HBlock LUN and the cloud.	B/s	Fine
	Cloud_U_Bandwidth	Bandwidth when HBlock LUN upload data to cloud.	B/s	Fine
	Cloud_D_Bandwidth	Bandwidth when HBlock LUN download data from cloud.	B/s	Fine
	Wait_Upload	The amount of data to be uploaded from HBlock LUN to the cloud.	bytes	Fine
pool (only supported by cluster mode)	IOPS	Total IOPS between the client and HBlock.	None	Fine
	R_IOPS	IOPS when the client reads data from the HBlock.	None	Fine
	W_IOPS	IOPS when the client writes data to HBlock.	None	Fine
	Bandwidth	Total bandwidth between the client and HBlock.	B/s	Fine
	R_Bandwidth	Bandwidth when the client reads data from HBlock.	B/s	Fine
	W_Bandwidth	Bandwidth when the client writes data to HBlock.	B/s	Fine
	Latency	Total latency between the client and HBlock. The average latency of read and write operations in the system within an acquisition cycle reflects the time that HBlock takes to process read and write requests.	ms	Fine
	W_Latency	Latency when the client writes data to HBlock. The average latency of write operations in the system within an acquisition cycle reflects the time that HBlock takes to process write requests.	ms	Fine
	R_Latency	Latency when the client reads data from HBlock. The average latency of read operations in the system within an acquisition cycle reflects the time that HBlock takes to process read requests.	ms	Fine
	Path_Cap	Total capacity of HBlock disk paths, i.e., total capacity of disks where all disk paths reside in the storage pool.	bytes	Rough
Path_Used	Used capacity of HBlock disk paths, i.e., used capacity of disks where all disk paths reside in the storage pool.	bytes	Rough	

	Path_Rate	Average usage rate of HBlock disk paths in the storage pool, i.e., Path_Used/Path_Cap.	%	Rough
	Path_Cap_Quota	Available space for HBlock, i.e., total capacity quota allocated by the user to HBlock on disks where all HBlock disk paths reside in the storage pool.	bytes	Rough
	Path_Cap_Quota_Used	Space occupied by HBlock data on disks where all HBlock disk paths reside in the storage pool.	bytes	Rough
	Path_Cap_Quota_Rate	Average usage rate of HBlock disk path capacity quota in the storage pool, i.e., Path_Cap_Quota_Used/Path_Cap_Quota.	%	Rough

7.5 Alarm List

Alarm Rule Name	Alarm Level	Alarm Condition	Automatic Resolution Condition	Alarm Expiration Condition	Allow Manual Resolving	Alert Email Sending Frequency
PathIOError	Major	The disk path status is error.	The disk path status is normal.	<ul style="list-style-type: none"> The disk path is removed. The disk path is removed from the storage pool. The server where the disk path located is removed. 	Yes	One time a day
DiskWriteSlow	Warning	The disk where the disk path resides is slow to write.	The disk where the disk path resides returns to normal.	<ul style="list-style-type: none"> The disk path is removed. The disk path is removed from the storage pool. The server where the disk path located is removed. 	Yes	One time a day
LicenseWillExpire	Warning	The current time (the system time of the server where the alarm module is located) is ≤ 15 days and > 0 days from the expiration time of the last imported license.	The current time (the system time of the server where the alarm module is located) is > 15 days from the expiration time of the alarm license.	<ul style="list-style-type: none"> The license has expired. A new license (license with a different ID) has been imported. 	Yes	One time a day
LicenseExpired	Critical	The current time (the system time of the server where the alarm module is located) is ≤ 0 days from the expiration time of the last imported license.	The current time (the system time of the server where the alarm module is located) is > 0 days from the expiration time of the alarm license.	A new license (license with a different ID) is imported.	Yes	Send one time
LicenseMaintenanceWillExpire	Warning	The current time (the system time of the server where the alarm module is located) is ≤ 15 days and > 0 days from the maintenance expiration time of the last imported license.	The current time (the system time of the server where the alarm module is located) is > 15 days from the maintenance expiration date of the alarm license.	<ul style="list-style-type: none"> The license has expired. A new license (license with a different ID) has been imported. 	Yes	One time a day
LicenseMaintenanceExpired	Warning	The current time (the system time of the server where the alarm module is located) is ≤ 0 days from the maintenance expiration time of the last imported license.	The current time (the system time of the server where the alarm module is located) is > 0 days from the maintenance expiration date of the alarm license.	A new license is imported.	Yes	Send one time
TrialVersionWillExpire	Warning	No valid license is currently imported, and the current time (the system time of the server where the alarm module is located) is ≤ 15 days and ≥ 0 days away from the trial period expiration time.	There are no conditions for resolving and can only be resolved manually.	A new license is imported.	Yes	One time a day
ResourceUsageApproachingLimit	Major	Total local LUN capacity $\geq 80\%$ of license capacity	Total local LUN capacity $< 75\%$ of license capacity	A new license is imported.	Yes	Send one time
AlarmNumberApproachingLimit	Critical	The number of unresolved alarms ≥ 8000	Number of alarms < 7500	None	Yes	One time a day
FailToSendAlarmEmail	Critical	Alarm email failed to be sent.	Alarm email sent successfully.	<ul style="list-style-type: none"> The email configuration is deleted. The email sending 	Yes	One time a day

				setting is disabled.		
CapacityQuotaUsageExceedsThreshold	Warning	Path_Cap_Quota_Rate of disks associated with disk paths in a storage pool is greater than or equal to 80%. Path_Cap_Quota_Rate of the disk associated with disk paths at the disk path level is greater than or equal to 80%. Note: If no capacity quota is set for disk paths, capacity quota = total disk capacity.	Path_Cap_Quota_Rate of disks associated with the disk paths in the storage pool is lower than 75%, or Path_Cap_Quota_Rate of the disks associated with disk paths at the disk path level is lower than 75%. Note: If no capacity quota is set for disk paths capacity quota = total disk capacity.	<ul style="list-style-type: none"> The name of a storage pool is changed. All disk paths in the storage pool are removed. 	Yes	One time a day
CapacityQuotaUsageApproachLimit	Critical	The total disk quota usage of disk paths in a base storage pool is greater than or equal to 95%.	The total disk quota usage of disk paths in a base storage pool is lower than 90%.	The storage pool name has been modified.	Yes	One time a day
CapacityQuotaUsageApproachLimit	Warning	The total disk quota usage of disk paths in a non-base storage pool is greater than or equal to 95%, or the disk quota usage of disk paths is greater than or equal to 95%.	The total disk quota usage of disk paths in a non-base storage pool is lower than 90%, or the disk quota usage of disk paths is lower than 90%.	<ul style="list-style-type: none"> The storage pool name has been modified. All disk paths in the storage pool are removed. 	Yes	One time a day
DiskUsageExceedsThreshold	Warning	Path_Rate of disks where disk paths are located in a storage pool is greater than or equal to 80%, or Path_Rate of disks where disk paths are located in is greater than or equal to 80%.	Path_Rate of the disks where disk paths are located in a storage pool is lower than 75%, or Path_Rate of disks where disk paths are located in is lower than 75%.	<ul style="list-style-type: none"> The storage pool name has been modified. All disk paths in the storage pool are removed. 	Yes	One time a day
InsufficientFDForLUNToWrite	Warning	The number of available fault domains and the number of health disk paths in the cache storage pool or storage pool where the LUN resides do not meet the minimum number of replicas of a LUN.	The number of available fault domains and the number of health disk paths in the warning alarm storage pool meets the minimum number of replicas of a LUN.	<ul style="list-style-type: none"> The LUN is deleted. The LUN deletion failed. The LUN recovered failed. The LUN is being suspended. The LUN is suspended. The LUN failed to be suspended. The name of the storage pool is changed. 	Yes	One time a day
DiskPathHealthStatusWarning	Warning	Disk path health status changes to Warning.	Disk path health status returns to normal.	<ul style="list-style-type: none"> Disk path removed. The disk path is removed from the storage pool. Disk path health status changes to "Error". 	Yes	One time a day
DiskPathHealthStatusError	Major	Disk path health status changes to Error.	Disk path health status returns to normal.	<ul style="list-style-type: none"> The disk path is removed. The disk path is removed from the storage pool. 	Yes	One time a day
DataServiceHealthStatusWarning	Warning	Data service health status changes to Warning.	Data service health status returns to normal.	<ul style="list-style-type: none"> The server is removed. 	Yes	One time a day

				<ul style="list-style-type: none"> ● The disk path is removed. ● The disk path is removed from the storage pool. ● The health status of the disk path changes to Error. ● The name of a storage pool is changed. 		
DataServiceHealthStatusError	Major	Data service health status changes to Error.	Data service health status returns to normal.	<ul style="list-style-type: none"> ● The server is removed. ● The disk path is removed. ● The disk path is removed from the storage pool. ● The name of a storage pool is changed. 	Yes	One time a day
ProtocolServiceAbnormal	Major	Protocol resolution service exception.	Protocol resolution service returns to normal.	<ul style="list-style-type: none"> ● Server removed. ● The target is deleted. ● The target is migrated. 	Yes	One time a day
FaultDomainWarning	Warning	Fault domain status changes to Warning.	Fault domain status returns to normal.	<ul style="list-style-type: none"> ● Server is removed. ● Disk path is removed. ● Fault domain health status changes to "Error". ● The name of the storage pool is changed or the name of any node in the full path name of the fault domain is changed. ● All disk paths in the fault domain are removed. 	Yes	One time a day
FaultDomainError	Major	Fault domain status changes to Error	Fault domain status returns to normal.	<ul style="list-style-type: none"> ● Server removed. ● Disk path removed. ● The disk path is removed from the storage pool. ● The name of the storage pool is changed or the name of any node in the full path name of the fault domain is changed. 	Yes	One time a day

				<ul style="list-style-type: none"> ● All disk paths in the fault domain are removed. 		
CannotConnectToCloud	Critical	Disconnected from the cloud for more than 10 minutes.	The LUN successfully reads data from the HBlock server to the cloud or writes data from the cloud once.	<ul style="list-style-type: none"> ● The LUN is deleted. ● The LUN deletion failed. ● The LUN recovered failed. ● The LUN failed to be suspended. ● The LUN is suspended. ● The sever has been removed. ● The target is migrated. 	Yes	One time a day
CloudAccountAbnormal	Critical	Failed to read data from the cloud or write data to the cloud: frozen overdue fees, frozen overdue fees, or illegally frozen fees.	Successfully read data from the cloud or successfully write data to the cloud.	<ul style="list-style-type: none"> ● The LUN is deleted. ● The LUN deletion failed. ● The LUN recovered failed. ● The LUN failed to be suspended. ● The LUN is suspended. 	Yes	One time a day
LUNCloudDataConflict	Critical	The cloud data corresponding to the LUN has a more recent version than the local version.	The cloud data version corresponding to the LUN is all older than the local version.	<ul style="list-style-type: none"> ● The LUN is deleted. ● The LUN deletion failed. ● The LUN recovered failed. ● The LUN failed to be suspended. ● The LUN is suspended. 	Yes	One time a day
LUNCloudHeartbeatConflict	Critical	A heartbeat from a non local cluster appears in the cloud of the LUN.	A heartbeat from the local cluster appears in the cloud of the LUN.	<ul style="list-style-type: none"> ● The LUN is deleted. ● The LUN deletion failed. ● The LUN recovered failed. ● The LUN failed to be suspended. ● The LUN is suspended. 	Yes	One time a day
LUNDataResidue	Warning	Forcefully delete the LUN. The LUN data is left over: <ul style="list-style-type: none"> ● Local leftover data: Due to a local disk failure, the LUN data cannot be synchronized for deletion. ● Cloud leftover data: When deleting the 	The alarm cannot be cleared automatically and can only be cleared manually.	<ul style="list-style-type: none"> ● Local data leftover: The data directory is removed from the machine. ● Cloud data leftover: It won't be cleared 	Yes	One time a day

		cloud data, the data cannot be synchronized for deletion because the cloud data storage location is inaccessible (including network connection failures, account abnormalities, etc.).		automatically. After cleaning up the residual cloud data, the alert can be cleared manually.		
InitiatorConnectionFailed	Warning	Due to the client's fault, HBlock was unable to receive the client's heartbeat, causing HBlock to believe that the connection with the client was disconnected and immediately raise an alarm. Except in cases where the client actively disconnects.	The client has successfully connected to target.	<ul style="list-style-type: none"> The target where the IQN of the alarm is located has been deleted. The target is migrated. 	Yes	One time a day
InsufficientSpaceonInstallationPath	Critical	The remaining space of the file system on the disk where the installation directory is located is less than or equal to 4 GiB.	The remaining space of the file system on the disk where the installation directory is located is greater than 5 GiB.	The server is removed.	Yes	One time a day
BaseServiceAbnormal	Critical	An alarm is generated when any of the following situations occur: <ul style="list-style-type: none"> The metadata management service is abnormal: There are two stor:mdm services in the cluster, and only either of them can fail. An alarm is generated when a fault occurs. The failover control service is abnormal: There are two stor:fc services in the cluster and only either of them can fail. An alarm is generated when a fault occurs. The log service is abnormal: There are three stor:ls services in the cluster, and only one of them can fail. An alarm is generated when a fault occurs. The coordination service is abnormal: There are three stor:cs services in the cluster, and only one of them can fail. An alarm is generated when a fault occurs. 	The matching alarm instance has been restored to the following levels: <ul style="list-style-type: none"> The metadata management service stor:mdm has been restored to normal on the alarming machine. The failover control service stor:fc has been restored to normal on the alarming machine. The log service stor:ls has been restored to normal on the alarming machine. The coordination service stor:cs has been restored to normal on the alarming machine. 	Base services on the alarming machine are migrated successfully.	Yes	One time a day
InsufficientSpaceonMetaDir	Critical	The remaining space of the file system on the disk where the base service data directory is located is less than or equal to 4GiB.	The remaining space of the file system on the disk where the base service data directory is located is greater than 5GiB.	The service begins to migrate.	Yes	One time a day

LUNDataLowRedundancy	Warning	LUN data low redundancy percentage (including cache pool and final pool data) > 10%, persisting for more than 10 minutes (low redundancy percentage remains below or equal to 10% within the 10-minute window).	UN data low redundancy percentage (including cache pool and final pool data) = 0%.	<ul style="list-style-type: none"> ● The LUN is deleted. ● The LUN deletion failed. 	Yes	One time a day
LUNDataCorrupted	Major	LUN data Error percentage (including cache pool and final pool data) > 0%.	LUN data Error percentage (including cache pool and final pool data) = 0%.	<ul style="list-style-type: none"> ● The LUN is deleted. ● The LUN deletion failed. 	Yes	One time a day
PoolDataLowRedundancy	Major	Base storage pool data low redundancy percentage > 10%, persisting for more than 10 minutes (low redundancy percentage remains below or equal to 10% within the 10-minute window).	Base storage pool data low redundancy percentage = 0%.	The storage pool name has been modified.	Yes	One time a day
PoolDataLowRedundancy	Warning	Non-base storage pool data low redundancy percentage > 10%, persisting for more than 10 minutes (low redundancy percentage remains below or equal to 10% within the 10-minute window).	Non-base storage pool data low redundancy percentage = 0%.	<ul style="list-style-type: none"> ● The storage pool name has been modified. ● The storage pool is deleted. 	Yes	One time a day
PoolDataCorrupted	Critical	Base storage pool data Error percentage > 0%.	Base storage pool data Error percentage = 0%.	The storage pool name has been modified.	Yes	One time a day
PoolDataCorrupted	Warning	Non-base storage pool data Error percentage > 0%.	Non-base storage pool data Error percentage = 0%.	<ul style="list-style-type: none"> ● The storage pool name has been modified. ● The storage pool is deleted. 	Yes	One time a day

7.6 OOS Endpoint and Region

The endpoints of Object Storage Network, Object Storage Network 2, Hongkong and other regions are different.

7.6.1 Object Storage Network

For Object Storage Network, the OOS Endpoint is oos-cn.ctyunapi.cn, region is cn.

Note: For Object Storage Network, if your data is stored in a certain region, it is recommended that you can directly use the region endpoint. The Endpoint list is as follows (the Endpoint list is only a description of the access information of the region Endpoint, and is not related to the region status):

Area	OOS API Endpoint	Region
ZhengZhou	oos-hazz.ctyunapi.cn	hazz
ShenYang	oos-lnsy.ctyunapi.cn	lnsy
ChengDu	oos-sccd.ctyunapi.cn	sccd
WuLuMuQi	oos-xjwlmq.ctyunapi.cn	xjwlmq
LanZhou	oos-gslz.ctyunapi.cn	gslz
QingDao	oos-sdqd.ctyunapi.cn	sdqd
GuiYang	oos-gzgy.ctyunapi.cn	gzgy
WuHan	oos-hbwh.ctyunapi.cn	hbwh
WuHu	oos-ahwh.ctyunapi.cn	ahwh
ShenZhen	oos-gdsz.ctyunapi.cn	gdsz
SuZhou	oos-jssz.ctyunapi.cn	jssz
SH2	oos-sh2.ctyunapi.cn	sh2

7.6.2 Object Storage Network 2

For Object Storage Network 2, the OOS Endpoint is oos-cn2.ctyunapi.cn, region is cn2.

Note: For Object Storage Network 2, if your data is stored in a certain region, it is recommended that you can directly use the region endpoint. The Endpoint list is as follows (the Endpoint list is only a description of the access information of the region Endpoint, and is not related to the region status):

Area	OOS API Endpoint	Region
NeiMeng1	oos-nm1.ctyunapi.cn	nm1
HangZhou1	oos-hz1.ctyunapi.cn	hz1

7.6.3 Hongkong Node

Hongkong nodes has two modes: highquality network and normal network. For Hongkong highquality network, OOS Endpoint is oos-cnkh-hqnet.ctyunapi.cn, region is cnhk-hqnet. For Hongkong normal network, OOS Endpoint is oos-cnkh-nqnet.ctyunapi.cn, region is cnhk-nqnet.

7.6.4 Other Region

Note: Other regions only support V2 signature, V4 signature is not supported.

Endpoint list for other regions is as follows.

region status):

Area	OOS API Endpoint	Region
BeiJing2	oos-bj2.ctyunapi.cn	bj2
NeiMeng2	oos-nm2.ctyunapi.cn	nm2
HangZhou	oos-hz.ctyunapi.cn	hz
JiangSu	oos-js.ctyunapi.cn	js
BeiJing	oos-hq-bj.ctyunapi.cn	hq-bj
ShangHai	oos-hq-sh.ctyunapi.cn	hq-sh

7.7 Operating System Monitoring Metrics Which HBlock can

Push

Category	Monitoring Metric	Description
server	hblock_cpu_seconds_user	User time.
	hblock_cpu_seconds_nice	Nice time.
	hblock_cpu_seconds_system	System time.
	hblock_cpu_seconds_idle	Idle time.
	hblock_cpu_seconds_iowait	I/O waiting time.
	hblock_cpu_seconds_irq	Hard interrupt time.
	hblock_cpu_seconds_softirq	Soft interrupt time.
	hblock_cpu_seconds_steal	The time spent forcibly waiting for another virtual CPU to finish processing.
	hblock_cpu_guest_seconds_user	The CPU time spent running a virtual machine. When the system runs a virtual machine in a virtualized environment, this field counts the CPU time used by the virtual machine.
	hblock_cpu_guest_seconds_nice	The CPU time spent running low priority virtual machines. Similar to the guest field, this field calculates the CPU time used to run low priority virtual machines.
	hblock_memory_MemTotal_bytes	Total memory sizes of the system.
	hblock_memory_MemFree_bytes	Unused memory size of the system.
	hblock_memory_MemAvailable_bytes	The available memory size of the system.
	Meminfo metrics for operating systems such as hblock_cemory_Buffers_bytes	Other meminfo indicators. After enabling pushgate function, you can view the specific indicator names on Prometheus.
hblock_memory_oom_kill	The number of oom killers occurring in the operating system.	
load	hblock_load1	The average load in the past minute.
	hblock_load5	The average load in the past 5 minutes.
	hblock_load15	The average load in the past 15 minutes.
interface	hblock_network_receive_bytes	The total number of bytes received by the interface.
	hblock_network_transmit_bytes	The total number of bytes sent by the interface.
	hblock_network_receive_packets	The number of packets received by the interface.
	hblock_network_transmit_packets	The number of packets sent by the interface.
	hblock_network_receive_errs	The number of erroneous packets that occurred during the interface reception process.
	hblock_network_transmit_errs	The number of erroneous packets that occurred during the interface sending process.

	hblock_network_receive_drop	The number of packets discarded during interface reception.
	hblock_network_transmit_drop	The number of packets discarded during the interface sending process.
	hblock_network_bandwidth	The interface bandwidth.
	hblock_network_status	Interface status, such as: <ul style="list-style-type: none"> ● down: 0 ● up: 1 ● unknown: 2 ● notpresent: 3 ● lowerlayerdown: 4 ● testing: 5 ● dormant: 6
	hblock_network_up_count	Interface up times.
	hblock_network_down_count	Interface down times.
tcp	hblock_netstat_tcp_RetransSegs	The number of TCP retransmitted packets.
	hblock_netstat_tcp_OutSegs	The number of TCP output packets.
	hblock_netstat_tcp_InSegs	The number of packets received by TCP.
	hblock_netstat_tcp_ActiveOpens	The current number of TCP connections in ActiveOpen status.
	hblock_netstat_tcp_CurrEstab	The current number of TCP connections in CurrEstab status.
	hblock_netstat_tcp_PassiveOpens	The number of TCP connections in the current PassiveOpens status.
	hblock_sockstat_tcp_mem	The number of TCP connections in the current mem status.
	hblock_sockstat_tcp_alloc	The current number of TCP connections in alloc status.
	hblock_sockstat_tcp_inuse	The current number of TCP connections in the inuse status.
	hblock_sockstat_tcp_orphan	The number of TCP connections in the current orphan status.
	hblock_sockstat_tcp_tw	The number of TCP connections in the current tw status.
disk	hblock_disk_read_bytes	The amount of data read from the disk.
	hblock_disk_written_bytes	The amount of data written to the disk.
	hblock_disk_reads_completed	The number of disk read requests.
	hblock_disk_read_time_seconds	Time of disk read request.
	hblock_disk_writes_completed	The number of disk write requests.
	hblock_disk_write_time_seconds	Time of disk write request.
	hblock_disk_io_time_seconds	The time for disk processing of I/O operations.
	hblock_disk_io_time_weighted_seconds	Weighted time for disk processing of I/O operations.

	hblock_disk_io_now	The actual number of I/O requests on the currently running disk.
fileSystem	hblock_fileSystem_size_bytes	Total file system capacity.
	hblock_fileSystem_free_bytes	Remaining capacity of file system.
	hblock_fileSystem_free_inode_count	Number of idle inodes.
	hblock_fileSystem_total_inode_count	Total number of inodes.
	hblock_fileSystem_readonly	Is the file system read-only.
	hblock_os_boot_time_seconds	The last startup time of the server.
OS	hblock_os_boot_time_seconds	The last boot time of the server.
Cloud	hblock_cloud_wait_upload_bytes	Data to be uploaded to the cloud.
	hblock_cloud_upload_size_bytes	Data already uploaded to the cloud.
	hblock_cloud_download_size_bytes	Data already downloaded from the cloud.
LUN	hblock_blockDevice_read_count	I/O read count (including successful and failed reads).
	hblock_blockDevice_read_fail_count	Number of I/O read failures.
	hblock_blockDevice_read_time_seconds	I/O read cumulative time (including successful and failed reads).
	hblock_blockDevice_read_fail_time_seconds	Cumulative time of I/O read failures.
	hblock_blockDevice_read_bytes	Size of successfully read I/O data.
	hblock_blockDevice_write_count	I/O write count (including successful and failed writes).
	hblock_blockDevice_write_fail_count	Number of I/O write failures.
	hblock_blockDevice_write_time_seconds	Cumulative time of I/O writes (including both successful and failed writes).
	hblock_blockDevice_write_fail_time_seconds	Cumulative time of I/O write failures.
	hblock_blockDevice_write_bytes	Size of successfully written I/O data.
diskpath	hblock_dataDir_total_bytes	The total capacity of the disk where the disk path resides.
	hblock_dataDir_avail_bytes	The available capacity of the disk where the disk path resides.
	hblock_dataDir_quota_total_bytes	The capacity quota of the disk where the disk path resides.
	hblock_dataDir_quota_used_bytes	The total amount of data that has been written to the disk path by HBlock.
status	hblock_status_system	System status: <ul style="list-style-type: none"> ● 0: Working ● 1: Upgrading ● 2: Uninstalling ● 3: Unknown
	hblock_status_data	System data health status, unit is %.
	hblock_status_license	Software license status:

		<ul style="list-style-type: none"> ● 0: Subscription license is effective. ● 1: Perpetual license is effective. ● 2: The license has not been imported. ● 3: Perpetual license has expired. ● 4: Subscription license has expired.
	hblock_status_licenseExpiredTime	License expiration time, Unix timestamp.
	hblock_status_poolDomain	Status of each fault domain in the storage pool: <ul style="list-style-type: none"> ● 0: Healthy ● 1: Warning ● 2: Error ● 3: Unknown
	hblock_status_baseService	Base service status: <ul style="list-style-type: none"> ● 0: Normal ● 1: Abnormal ● 2: Unknown
	hblock_status_dataService	Data service status: <ul style="list-style-type: none"> ● 0: Healthy ● 1: Warning ● 2: Error ● 3: Unknown

7.8 Cluster Topology File

A cluster topology displays the topological structure of all nodes within a cluster, including their names, statuses, and compositions. You can create a cluster topology file and import it during initialization. The topology file is a JSON file in the UTF-8 encoding format.

Note:

- When the parent node is root, the child node can be room, rack, or server.
- When the parent node is room, the child node can be rack or server.
- When the parent node is rack, the child node can only be server.
- When the parent node is server, the child node can only be path.

Parameter		Description	Required
name		The name of a root node. Type: String The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit. The default value is default.	No
type		Type of a root node. Type: String Value: root . The default value is root.	No
childNodes		A collection of child node information. Parameters vary with child node type. Type: Array	Yes
The child node type is room	name	The name of a child node. Type: String The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit.	Yes
	type	Type of child node. Type: String Value: room .	Yes
	description	Description of a child node. Type: String The value is a sting of 1 to 50 characters.	No
	childNodes	A collection of child node information. Parameters vary with child node type. A child node can be either rack or server.	No

		Type: Array	
The child node type is rack	name	The name of a child node. Type: String The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit.	Yes
	type	Type of child node. Type: String Value: rack .	Yes
	description	Description of a child node. Type: String The value is a sting of 1 to 50 characters.	No
	childNodes	A collection of child node information. The child node can only be server. Type: Array	No
The child node type is server	name	The name of a child node. Type: String The value is a string of 1 to 63 case-sensitive characters. It can contain letters, digits, dots (.), hyphens (-), and underscores (_). Only supports starting with a letter or a digit. The default value is server ID.	No
	type	Type of child node. Type: String Value: server , indicating server type.	Yes
	description	Description of a child node. Type: String The value is a sting of 1 to 50 characters.	No
	ip	IP address of the HBlock server. The value is IPv4 or [IPv6] address.	Yes
	apiPort	Manage API ports. Type: integer The value is an integer that ranges from 1 to 65535. The default value is 1443. The setting must be the same as the API port number set when HBlock is installed on the server.	No
	childNodes	A collection of child node information. The child node type is path.	Yes
The child node type is path	name	The name of a child node. Type: String Value: Specifies the disk path.	Yes
	type	Type of child node.	Yes

		Type: String Value: path.	
	capacityQuota	<p>Specifies the capacity quota of the disk path, that is, for each disk path added to server, the total amount of data that can be written by HBlock. Once the space used by HBlock reaches the quota, data writing is immediately blocked, and no more space is allowed to be used beyond the quota.</p> <p>Type: long integer Value: The value is smaller than the total capacity of the disk where the disk path resides. A negative integer means unlimited writing and 0 means write inhibit. Writes are not limited by default.</p> <p>Note: If the same disk path appears multiple times, the capacity quota of the first occurrence will be used.</p>	No

Example of a topology file:

```

{
  "name": "default",
  "childNodes": [
    {
      "name": "room1",
      "type": "room",
      "childNodes": [
        {
          "type": "server",
          "name": "server1",
          "ip": "192.168.0.1",
          "apiPort": 1443,
          "childNodes": [
            {
              "name": "/mnt/storage01",
              "type": "path",
              "capacityQuota": 96636764160
            },
            {
              "name": "/mnt/storage02",
              "type": "path"
            }
          ]
        },
        {
          "type": "server",
          "name": "server2",

```

```
"ip": "192.168.0.2",
"apiPort": 1443,
"childNodes": [
  {
    "name": "/mnt/storage01",
    "type": "path",
    "capacityQuota": 96636764160
  }
],
},
{
  "type": "server",
  "name": "server3",
  "ip": "192.168.0.3",
  "apiPort": 1443,
  "childNodes": [
    {
      "name": "/mnt/storage01",
      "type": "path"
    },
    {
      "name": "/mnt/storage02",
      "type": "path"
    }
  ]
}
]
}
]
}
```

7.9 iSCSI Target Allowlist File

The allowlist file contains both the iSCSI initiator allowlist and the target allowlist.

The allowlist file is a JSON file in the UTF-8 encoding format.

The format of the allowlist file is as follows:

```
{
  "initiator": [
    {
      "IPs": [ ip, ip, ... ],
      "names": [ name, name, ... ]
    },
    {
      "IPs": [ ip, ip, ... ],
      "names": [ name, name, ... ]
    },
    .....
  ],
  "target": [
    {
      "IPs": [ ip, ip, ... ],
      "NICs": [ NIC, NIC, ... ]
    },
    {
      "IPs": [ ip, ip, ... ],
      "NICs": [ NIC, NIC, ... ]
    },
    .....
  ]
}
```

Parameter

Parameter	Type	Description	Required
initiator	Array of initiator	Set the iSCSI initiator allowlist for every IQN under the target. You can define several allowlists, the relationship among them is logical OR . Within one allowlist you may simultaneously specify both IP address and initiator name; these two conditions are combined with logical AND . See Table 1 Parameter description of initiator for details.	No
target	Array of target	Set the target allowlist for every IQN under the target. You can define several allowlists, the relationship among them is logical OR . Within one allowlist you may simultaneously specify both IP address and NIC name; these two conditions are combined with logical AND .	No

		See Table 2 Parameter description of target for details.	
--	--	---	--

Table 1 Parameter description of initiator

Parameter	Type	Description	Required
IPs	Array of ip	Set the allowlist for iSCSI initiators based on their IP addresses. Value: IPv4, IPv6, or CIDR subnets. Multiple entries can be configured and must be separated by commas.	No
names	Array of name	Set the allowlist for iSCSI initiators based on their names. Value: The value is a string of 1 to 223 case-sensitive characters. It can only contain letters, digits, dots (.), colons (:), or hyphens (-). Wildcards * and ? are supported. Multiple entries can be configured and must be separated by commas.	No

Table 2 Parameter description of target

Parameter	Type	Description	Required
IPs	Array of ip	Set the allowlist for targets based on their IP addresses. Value: IPv4, IPv6, or CIDR subnets. Multiple entries can be configured and must be separated by commas. Note: 'localhost' addresses are not permitted.	No
NICs	Array of NIC	Set the allowlist for targets based on their NICs. Value: The value is a string of 1 to 100 characters. It can contain letters, digits, dots (.), hyphens (-), or underscores (_). Wildcards * and ? are supported. Multiple entries can be configured and must be separated by commas.	No