Bare Metal Server

# Best Practices

| | |
|---|---|
| **Issue** | 01 |
| **Date** | 2025-11-07 |

HUAWEI

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

## Trademarks and Permissions


and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

## Notice

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

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

# Contents



  

# 1 Monitoring

## 1.1 Overview

### Solution Introduction

After purchasing a BMS, you want to know its running status. Bare Metal Server (BMS) works with the Cloud Eye service to automatically collect monitoring metrics, such as the CPU, memory, disk, and network usage of a BMS. These metrics help you learn about the running status and performance of your BMS in time.

This document is prepared based on the BMS and Cloud Eye practices and provides guidance for you to configure server monitoring for BMSs.

### Constraints

- Agent can be installed only on BMSs running a 64-bit Linux OS.
- An agency must be configured for the BMS. For details, see **How Do I Configure an Agency?**
- Only AP-Bangkok (**ap-southeast-2**) and CN-Hong Kong (**ap-southeast-1**) are supported now.
- Private images do not support this function.

  **Table 1-1** lists the Linux images that support server monitoring.

  **Table 1-1** Linux images that support server monitoring

| OS Type (64-bit) | Version |
| --- | --- |
| SUSE | Enterprise11 SP4 |
| CentOS | 6.9, 7.2, 7.3, and 7.4 |

  

# 1.2 Installing and Configuring the Agent for an Existing BMS

## 1.2.1 Installing the Agent

This section describes how to install the Agent for an existing BMS. The procedure is as follows:

1. **Configure an Agency**: Use an agency to authorize the Agent installed on BMSs in the region.
2. **Adding the Resolved Domain Names**: Add the resolved domain names of regions to the **/etc/resolv.conf** file on the BMS.
3. **Configuring the Security Group**: Download the Telescope package, send metrics, and collect logs.
4. **Procedure**: Manually install the Agent on the BMS.

### Adding the Resolved Domain Names

1. Log in to the BMS as user **root**.
2. Enter **vi /etc/resolv.conf** to open the **/etc/resolv.conf** file.
3. Add **nameserver 100.125.1.250** and **nameserver 100.125.21.250** to the file, as shown in **Figure 1-1**.

   **Figure 1-1** Adding the resolved domain names

   ```
   # Dynamic resolv.conf(5) file for glibc resolver(3) generated by resolvconf(8)
   #     DO NOT EDIT THIS FILE BY HAND -- YOUR CHANGES WILL BE OVERWRITTEN
   nameserver 100.125.1.250
   nameserver 114.114.114.114
   nameserver 114.114.115.115
   search openstacklocal
   ~
   ```

   NOTE

   The values of **nameserver** vary depending on the region.

   CN North-Beijing1: 100.125.1.250 and 100.125.21.250

   CN North-Beijing4: 100.125.1.250 and 100.125.129.250

   CN East-Shanghai1: 100.125.1.250,100.125.64.250

   CN South-Guangzhou: 100.125.1.250 and 100.125.136.29

   CN-Hong Kong: 100.125.1.250, 100.125.3.250

   AP-Bangkok: 100.125.1.250,100.125.3.250

   LA-Santiago: 100.125.1.250

4. Press **Esc** and enter **:wq!** to save the configuration.

### Configuring the Security Group

1. On the page showing the BMS details, click the **Security Groups** tab.

  

2. Click to expand the security group details, showing the configured security group rules.
3. In the upper right corner of the rule list, click the security group ID to go to the **Security Groups** page.
4. In the **Operation** column, click **Manage Rule**. On the **Outbound Rules** tab page, click **Add Rule** to add a rule based on **Table 1-2**.

**Table 1-2** Security group rules

| Direction | Protocol | Port | Destination IP address | Description |
| --- | --- | --- | --- | --- |
| Outbound | TCP | 80 | 100.125.0.0/16 | Used to download the Agent installation package from the OBS bucket to the BMS and obtain the metadata and authentication information of the BMS. |
| Outbound | TCP and UDP | 53 | 100.125.0.0/16 | Used by DNS to resolve domain names, for example, resolve the OBS domain name when users are downloading the Agent installation package, and resolve the Cloud Eye endpoint when the Agent is sending monitoring data to Cloud Eye. |
| Outbound | TCP | 443 | 100.125.0.0/16 | Used to collect monitoring data that will be sent to Cloud Eye. |

## Procedure

1. Log in to the BMS as user **root**.
2. Run the following command to install the Agent:

   CN North-Beijing1:

   ```
   cd /usr/local && curl -k -O https://obs.cn-north-1.myhuaweicloud.com/uniagent-cn-north-1/script/agent_install.sh && bash agent_install.sh
   ```

   CN North-Beijing4:

   ```
   cd /usr/local && curl -k -O https://obs.cn-north-4.myhuaweicloud.com/uniagent-cn-north-4/script/agent_install.sh && bash agent_install.sh
   ```

   CN South-Guangzhou:

   ```
   cd /usr/local && curl -k -O https://obs.cn-south-1.myhuaweicloud.com/uniagent-cn-south-1/script/agent_install.sh && bash agent_install.sh
   ```

   CN East-Shanghai1:

   ```
   cd /usr/local && curl -k -O https://obs.cn-east-3.myhuaweicloud.com/uniagent-cn-east-3/script/agent_install.sh && bash agent_install.sh
   ```

   CN East-Shanghai2:

   ```
   cd /usr/local && curl -k -O https://obs.cn-east-2.myhuaweicloud.com/uniagent-cn-east-2/script/agent_install.sh && bash agent_install.sh
   ```

   CN Southwest-Guiyang1:

  

```
cd /usr/local && curl -k -O https://obs.cn-southwest-2.myhuaweicloud.com/uniagent-cn-southwest-2/script/agent_install.sh && bash agent_install.sh
```

CN-Hong Kong:

```
cd /usr/local && curl -k -O https://obs.ap-southeast-1.myhuaweicloud.com/uniagent-ap-southeast-1/script/agent_install.sh && bash agent_install.sh
```

AP-Bangkok:

```
cd /usr/local && curl -k -O https://obs.ap-southeast-2.myhuaweicloud.com/uniagent-ap-southeast-2/script/agent_install.sh && bash agent_install.sh
```

AP-Singapore:

```
cd /usr/local && curl -k -O https://obs.ap-southeast-3.myhuaweicloud.com/uniagent-ap-southeast-3/script/agent_install.sh && bash agent_install.sh
```

AP-Jakarta:

```
cd /usr/local && curl -k -O https://obs.ap-southeast-4.myhuaweicloud.com/uniagent-ap-southeast-4/script/agent_install.sh && bash agent_install.sh
```

AF-Johannesburg:

```
cd /usr/local && curl -k -O https://obs.af-south-1.myhuaweicloud.com/uniagent-af-south-1/script/agent_install.sh && bash agent_install.sh
```

LA-Santiago:

```
cd /usr/local && curl -k -O https://uniagent-la-south-2.obs.la-south-2.myhuaweicloud.com/script/agent_install.sh && bash agent_install.sh
```

LA-Sao Paulo1:

```
cd /usr/local && curl -k -O https://uniagent-sa-brazil-1.obs.sa-brazil-1.myhuaweicloud.com/script/agent_install.sh && bash agent_install.sh
```

LA-Mexico City1:

```
cd /usr/local && wget https://telescope-na-mexico-1.obs.na-mexico-1.myhuaweicloud.com/scripts/agentInstall.sh && chmod 755 agentInstall.sh && ./agentInstall.sh
```

LA-Mexico City2

```
cd /usr/local && curl -k -O https://uniagent-la-north-2.obs.la-north-2.myhuaweicloud.com/script/agent_install.sh && bash agent_install.sh
```

The Agent is installed successfully if the command output similar to the following figure is displayed.

  

**Figure 1-2** Successful installation

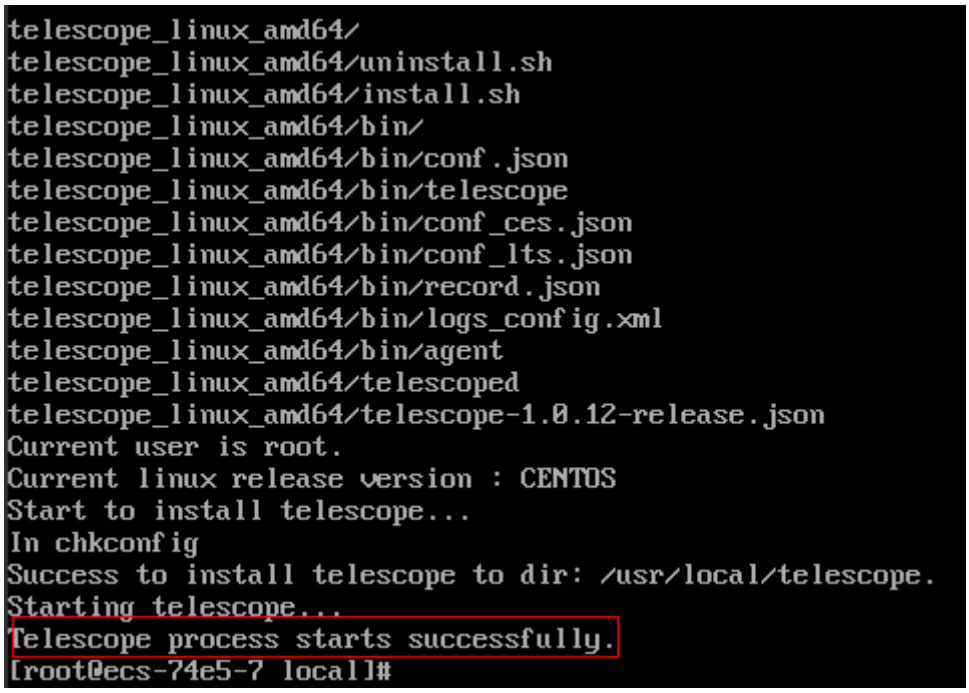


3. After the installation is complete, configure the Agent by referring to **(Optional) Manually Configuring the Agent (Linux)**.
4. Run the following command to delete the installation script:

   **if [[ -f /usr/local/uniagent/extension/install/telescope/bin/telescope ]]; then rm /usr/local/agent_install.sh; else rm /usr/local/agentInstall.sh; fi**

## 1.2.2 (Optional) Managing the Agent

This section guides you to manage the Agent. You can view, start, stop, and uninstall the Agent as needed.

NOTE

You need to view, start, stop, and uninstall the Agent as user **root**.

### Checking the Agent Status

Log in to the BMS and run the following command to check the Agent status:

**service telescoped status**

The Agent is running properly if the system displays the following information:

```
"Telescope process is running well."
```

### Starting the Agent

Run the following command to start the Agent:

**/usr/local/telescope/telescoped start**

  

### Restarting the Agent

Run the following command to restart the Agent:

**/usr/local/telescope/telescoped restart**

### Stopping the Agent

Run the following command to stop Agent:

**service telescoped stop**

NOTE

If the Telescope installation fails, you may fail to stop the Agent, and you can run the following command to stop the Agent again:

**/usr/local/telescope/telescoped stop**

### Uninstalling the Agent

You can manually uninstall the Agent. After the uninstallation, Cloud Eye does not collect the BMS monitoring data. If you need to use the Agent again, install it again. For details, see section **Installing the Agent**.

Run the following command to uninstall the Agent:

**/usr/local/telescope/uninstall.sh**

## 1.3 Monitoring Data

Log in to the management console and choose **Cloud Eye**. In the navigation pane on the left, choose **Server Monitoring** > **Bare Metal Server**. In the right pane, **Name/ID**, **Status**, and **Agent Status** of the BMS are displayed.

**Figure 1-3** Server monitoring

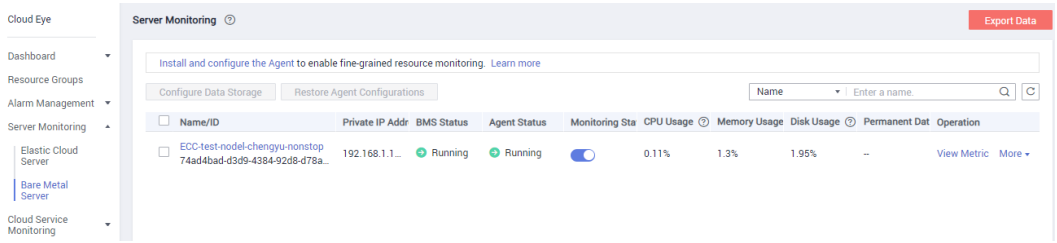


You can click **View Metric** in the **Operation** column to obtain the visualized monitoring graph of the BMS and view monitoring metrics of the BMS, such as the CPU usage, CPU load, and memory usage.

  

**Figure 1-4** Visualized monitoring graph

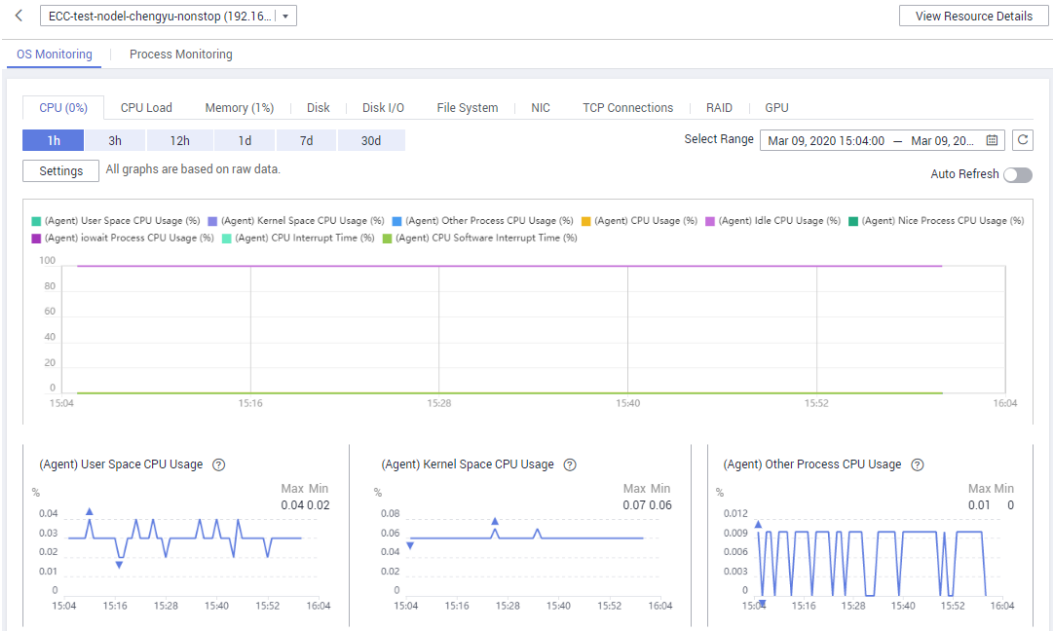


# 1.4 Monitored Metrics (with Agent Installed)

## Description

This section describes monitoring metrics reported by BMS to Cloud Eye as well as their namespaces and dimensions. You can use the management console or APIs provided by Cloud Eye to query the metrics of the monitored objects and alarms generated for BMS.

NOTE

After installing the Agent on a BMS, you can view its OS monitoring metrics. Monitoring data is collected at an interval of 1 minute.

## Namespace

SERVICE.BMS

## Metrics

Supported BMS **OS Monitoring** metrics include CPU metrics listed in **Table 1-3**, CPU load metrics listed in **Table 1-4**, memory metrics listed in **Table 1-5**, disk metrics listed in **Table 1-6**, disk I/O metrics listed in **Table 1-7**, file system metrics listed in **Table 1-8**, NIC metrics listed in **Table 1-9**, software RAID metrics listed in **Table 1-10**, and process metrics in **Table 1-11**.

NOTE

To monitor software RAID metrics, Agent 1.0.5 or later is required.

Currently, BMSs running the Windows OS cannot be monitored.

  

**Table 1-3** CPU metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| cpu_usage_idle | (Agent) Idle CPU Usage | Percentage of time that CPU is idle<br>Check the metric value changes in the **/proc/stat** file in a collection period.<br>Run the **top** command to check the **%Cpu(s) id** value.<br>Unit: percent | 0-100% | BMS | 1 minute |
| cpu_usage_other | (Agent) Other Process CPU Usage | Percentage of time that the CPU is used by other processes<br>Formula:<br>**Other Process CPU Usage** = **1**- **Idle CPU Usage** (%) - **Kernel Space CPU Usage** (%) - **User Space CPU Usage** (%)<br>Unit: percent | 0-100% | BMS | 1 minute |
| cpu_usage_system | (Agent) Kernel Space CPU Usage | Percentage of time that the CPU is used by kernel space<br>Check the metric value changes in the **/proc/stat** file in a collection period.<br>Run the **top** command to check the **%Cpu(s) sy** value.<br>Unit: percent | 0-100% | BMS | 1 minute |
| cpu_usage_user | (Agent) User Space CPU Usage | Percentage of time that the CPU is used by user space<br>Check the metric value changes in the **/proc/stat** file in a collection period.<br>Run the **top** command to check the **%Cpu(s) us** value.<br>Unit: percent | 0-100% | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| cpu_usage | (Agent) CPU Usage | CPU usage of the monitored object<br>Check the metric value changes in the **/proc/stat** file in a collection period.<br>Run the **top** command to check the **%Cpu(s)** value.<br>Unit: percent | 0-100% | BMS | 1 minute |
| cpu_usage_nice | (Agent) Nice Process CPU Usage | Percentage of time that the CPU is used by the Nice process<br>Check the metric value changes in the **/proc/stat** file in a collection period. Run the **top** command to check the **%Cpu(s) ni** value.<br>Unit: percent | 0-100% | BMS | 1 minute |
| cpu_usage_iowait | (Agent) iowait Process CPU Usage | Percentage of time during which the CPU is waiting for I/O operations to complete<br>Check the metric value changes in the **/proc/stat** file in a collection period.<br>Run the **top** command to check the **%Cpu(s) wa** value.<br>Unit: percent | 0-100% | BMS | 1 minute |
| cpu_usage_irq | (Agent) CPU Interrupt Time | Percentage of time that the CPU is servicing interrupts<br>Check the metric value changes in the **/proc/stat** file in a collection period.<br>Run the **top** command to check the **%Cpu(s) hi** value.<br>Unit: percent | 0-100% | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| cpu_usage_softirq | (Agent) CPU Software Interrupt Time | Percentage of time that the CPU is servicing software interrupts<br>Check the metric value changes in the **/proc/stat** file in a collection period.<br>Run the **top** command to check the **%Cpu(s) si** value.<br>Unit: percent | 0-100% | BMS | 1 minute |

**Table 1-4** CPU load metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| load_average1 | (Agent) 1-Minute Load Average | CPU load averaged from the last 1 minute<br>Obtain its value by dividing the **load1/** value in **/proc/loadavg** by the number of logical CPUs.<br>Run the **top** command to check the **load1** value. | ≥ 0 | BMS | 1 minute |
| load_average5 | (Agent) 5-Minute Load Average | CPU load averaged from the last 5 minutes<br>Obtain its value by dividing the **load5/** value in **/proc/loadavg** by the number of logical CPUs.<br>Run the **top** command to check the **load5** value in the **/proc/loadavg** file. | ≥ 0 | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| load_average15 | (Agent) 15-Minute Load Average | CPU load averaged from the last 15 minutes<br>Obtain its value by dividing the **load15/** value in **/proc/loadavg** by the number of logical CPUs.<br>Run the **top** command to check the **load15** value in the **/proc/loadavg** file. | ≥ 0 | BMS | 1 minute |

**Table 1-5** Memory metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mem_available | (Agent) Available Memory | Available memory size of the monitored object<br>Obtain the **MemAvailable** value by checking the file **/proc/meminfo**. If it is not displayed in the file:<br>**MemAvailable** = **MemFree** + **Buffers** + **Cached**<br>Unit: GB | ≥ 0 GB | BMS | 1 minute |
| mem_usedPercent | (Agent) Memory Usage | Memory usage of the monitored object<br>Obtain its value by checking the file **/proc/meminfo**. **Memory Usage** = (**MemTotal** - **MemAvailable**)/ **MemTotal**<br>Unit: percent | 0-100% | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mem_free | (Agent) Idle Memory | Amount of memory that is not being used<br>Obtain its value by checking the file **/proc/meminfo**.<br>Unit: GB | ≥ 0 GB | BMS | 1 minute |
| mem_buffers | (Agent) Buffer | Memory that is being used for buffers<br>Obtain its value by checking the file **/proc/meminfo**.<br>Run the **top** command to check the **KiB Mem:buffers** value.<br>Unit: GB | ≥ 0 GB | BMS | 1 minute |
| mem_cached | (Agent) Cache | Memory that is being used for file caches<br>Obtain its value by checking the file **/proc/meminfo**.<br>Run the **top** command to check the **KiB Swap:cached Mem** value.<br>Unit: GB | ≥ 0 GB | BMS | 1 minute |

  

**Table 1-6** Disk metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mountPointPrefix_disk_free | (Agent) Available Disk Space | Available disk space of the monitored object<br>Run the **df -h** command to check the data in the **Avail** column.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: GB | ≥ 0 GB | BMS | 1 minute |
| mountPointPrefix_disk_total | (Agent) Disk Storage Capacity | Disk storage capacity of the monitored object<br>Run the **df -h** command to check the data in the **Size** column.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: GB | ≥ 0 GB | BMS | 1 minute |
| mountPointPrefix_disk_used | (Agent) Used Disk Space | Used disk space of the monitored object<br>Run the **df -h** command to check the data in the **Used** column.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: GB | ≥ 0 GB | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mountPointPrefix_disk_usedPercent | (Agent) Disk Usage | Disk usage of the monitored object. It is calculated as follows: **Disk Usage** = **Used Disk Space**/**Disk Storage Capacity**.<br>**Disk Usage** = **Used Disk Space**/**Disk Storage Capacity**<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: percent | 0-100% | BMS | 1 minute |

**Table 1-7** Disk I/O metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mountPointPrefix_disk_agt_read_bytes_rate | (Agent) Disks Read Rate | Volume of data read from the monitored object per second<br>The disk read rate is calculated by checking data changes in the sixth column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: byte/s | ≥ 0 bytes/s | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mountPointPrefix_disk_agt_read_requests_rate | (Agent) Disks Read Requests | Number of read requests sent to the monitored object per second<br>The disk read requests are calculated by checking data changes in the fourth column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: request/s | ≥ 0 | BMS | 1 minute |
| mountPointPrefix_disk_agt_write_bytes_rate | (Agent) Disks Write Rate | Volume of data written to the monitored object per second<br>The disk write rate is calculated by checking data changes in the tenth column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: byte/s | ≥ 0 bytes/s | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mountPointPrefix_disk_agt_write_requests_rate | (Agent) Disks Write Requests | Number of write requests sent to the monitored object per second<br>The disk write requests are calculated by checking data changes in the eighth column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: request/s | ≥ 0 | BMS | 1 minute |
| disk_readTime | (Agent) Average Read Request Time | Average amount of time that read requests have waited on the disks<br>The average read request time is calculated by checking data changes in the seventh column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: ms/count | ≥ 0 ms/Count | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| disk_writeTime | (Agent) Average Write Request Time | Average amount of time that write requests have waited on the disks<br>The average write request time is calculated by checking data changes in the eleventh column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: ms/count | ≥ 0 ms/Count | BMS | 1 minute |
| disk_ioUtils | (Agent) Disk I/O Usage | Disk I/O usage of the monitored object<br>Check the data changes in the thirteenth column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: percent | 0-100% | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| disk_queue_length | (Agent) Disk Queue Length | Average number of read or write requests to be processed for the monitored disk in the monitoring period<br>The average disk queue length is calculated by checking data changes in the fourteenth column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: count | ≥ 0 | BMS | 1 minute |
| disk_write_bytes_per_operation | (Agent) Average Disk Write Size | Average number of bytes in an I/O write for the monitored disk in the monitoring period<br>The average disk write size is calculated by dividing the data changes in the tenth column of the corresponding device by that of the eighth column in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: KB/op | ≥ 0 KB/op | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| disk_read_bytes_per_operation | (Agent) Average Disk Read Size | Average number of bytes in an I/O read for the monitored disk in the monitoring period<br>The average disk read size is calculated by dividing the data changes in the sixth column of the corresponding device by that of the fourth column in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: KB/op | ≥ 0 KB/op | BMS | 1 minute |
| disk_io_svctm | (Agent) Disk I/O Service Time | Average time in an I/O read or write for the monitored disk in the monitoring period<br>The average disk I/O service time is calculated by dividing the data changes in the thirteenth column of the corresponding device by the sum of data changes in the fourth and eighth columns in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: ms/op | ≥ 0 ms/op | BMS | 1 minute |

  

**Table 1-8** File system metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| disk_fs_rwstate | (Agent) File System Read/Write Status | Read and write status of the mounted file system of the monitored object Possible values are **0** (read and write) and **1** (read only).<br>Check file system information in the fourth column in the **/proc/mounts** file. | 0 and 1 | BMS | 1 minute |
| disk_inodesTotal | (Agent) Disk inode Total | Total number of index nodes on the disk Run the **df -i** command to check information in the **Inodes** column.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). | ≥ 0 | BMS | 1 minute |
| disk_inodesUsed | (Agent) Total inode Used | Number of used index nodes on the disk<br>Run the **df -i** command to check data in the **IUsed** column.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). | ≥ 0 | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| disk_inodesUsedPercent | (Agent) Percentage of Total inode Used | Percentage of used index nodes on the disk<br>Run the **df -i** command to check data in the **IUse%** column.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: percent | 0-100% | BMS | 1 minute |

**Table 1-9** NIC metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| net_bitRecv | (Agent) Inbound Bandwidth | Number of bits received by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: bit/s | ≥ 0 bits/s | BMS | 1 minute |
| net_bitSent | (Agent) Outbound Bandwidth | Number of bits sent by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: bit/s | ≥ 0 bits/s | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| net_packetRecv | (Agent) NIC Packet Receive Rate | Number of packets received by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: count/s | ≥ 0 counts/s | BMS | 1 minute |
| net_packetSent | (Agent) NIC Packet Send Rate | Number of packets sent by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: count/s | ≥ 0 counts/s | BMS | 1 minute |
| net_errin | (Agent) Receive Error Rate | Percentage of receive errors detected by this NIC per second<br>Unit: percent | 0-100 % | BMS | 1 minute |
| net_errout | (Agent) Transmit Error Rate | Percentage of transmit errors detected by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: percent | 0-100 % | BMS | 1 minute |
| net_dropin | (Agent) Received Packet Drop Rate | Percentage of packets discarded by this NIC to the total number of packets received by the NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: percent | 0-100 % | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| net_dropout | (Agent) Transmitted Packet Drop Rate | Percentage of packets transmitted by this NIC which were dropped per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: percent | 0-100 % | BMS | 1 minute |

**Table 1-10** Software RAID metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| md1_status_device:1 | (Agent) Status | Software RAID status of the monitored object. Its value is **0** if the RAID is abnormal.<br>Run the plug-in script **/usr/local/telescope/plugins/raid-monitor.sh** in a collection period. Obtain its value by checking data changes in the **/proc/mdstat** file and run **mdadm -D/dev/md0** (**md0** indicates the RAID name). | 0 and 1 | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| md1_active_device:2 | (Agent) Active Disks | Number of active disks in software RAID of the monitored object. Its value is **-1** if the RAID is abnormal.<br>Run the plug-in script **/usr/local/telescope/plugins/raid-monitor.sh** in a collection period. Obtain its value by checking data changes in the **/proc/mdstat** file and run **mdadm -D/dev/md0** (**md0** indicates the RAID name). | ≥ 0, –1 | BMS | 1 minute |
| md1_working_device:2 | (Agent) Working Disks | Number of working disks in software RAID of the monitored object. Its value is **-1** if the RAID is abnormal.<br>Run the plug-in script **/usr/local/telescope/plugins/raid-monitor.sh** in a collection period. Obtain its value by checking data changes in the **/proc/mdstat** file and run **mdadm -D/dev/md0** (**md0** indicates the RAID name). | ≥ 0, –1 | BMS | 1 minute |
| md1_failed_device:0 | (Agent) Failed Disks | Number of failed disks in software RAID of the monitored object. Its value is **-1** if the RAID is abnormal.<br>Run the plug-in script **/usr/local/telescope/plugins/raid-monitor.sh** in a collection period. Obtain its value by checking data changes in the **/proc/mdstat** file and run **mdadm -D/dev/md0** (**md0** indicates the RAID name). | ≥ 0, –1 | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| md1_spare_device:0 | (Agent) Spare Disks | Number of spare disks in software RAID of the monitored object. Its value is **-1** if the RAID is abnormal.<br>Run the plug-in script **/usr/local/telescope/plugins/raid-monitor.sh** in a collection period. Obtain its value by checking data changes in the **/proc/mdstat** file and run **mdadm -D/dev/md0** (**md0** indicates the RAID name). | ≥ 0, –1 | BMS | 1 minute |

**Table 1-11** Process metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| proc_pHashId_cpu | CPU Usage | CPU consumed by a process. **pHashId** (process name and process ID) is the value of **md5**.<br>Check the metric value changes in the **/proc/pid/stat** file.<br>Unit: percent | 0-100 % | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| proc_pHashId_mem | Memory Usage | Memory consumed by a process. **pHashId** (process name and process ID) is the value of **md5**.<br>**Memory Usage** = **RSS** x **PAGESIZE**/**MemTotal**<br>• Obtain the **RSS** value by checking the second column of the file **/proc/pid/statm**.<br>• Obtain the **PAGESIZE** value by running the **getconf PAGESIZE** command.<br>• Obtain the **MemTotal** value by checking the file **/proc/meminfo**.<br>Unit: percent | 0-100 % | BMS | 1 minute |
| proc_pHashId_file | Opened Files | Number of files opened by a process. **pHashId** (process name and process ID) is the value of **md5**.<br>Run the **ls -l /proc/pid/fd** command to view the number of opened files. | ≥0 | BMS | 1 minute |
| proc_running_count | (Agent) Running Processes | Number of running processes<br>You can obtain the status of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. | ≥0 | BMS | 1 minute |
| proc_idle_count | (Agent) Idle Processes | Number of idle processes<br>You can obtain the status of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. | ≥0 | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| proc_zombie_count | (Agent) Zombie Processes | Number of zombie processes<br>You can obtain the status of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. | ≥0 | BMS | 1 minute |
| proc_blocked_count | (Agent) Blocked Processes | Number of blocked processes<br>You can obtain the status of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. | ≥0 | BMS | 1 minute |
| proc_sleeping_count | (Agent) Sleeping Processes | Number of sleeping processes<br>You can obtain the status of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. | ≥0 | BMS | 1 minute |
| proc_total_count | (Agent) Total Processes | Total number of processes on the monitored object<br>You can obtain the status of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. | ≥0 | BMS | 1 minute |

  

# 1.5 Monitored Metrics

## Description

NOTE

After installing the Agent on a BMS, you can view its OS monitoring metrics. Monitoring data is collected at an interval of 1 minute.

## Namespace

SERVICE.BMS

## Metrics

Table 1-12 lists the metrics supported by BMS.

**Table 1-12** Metrics

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| cpu_usage | (Agent) CPU Usage | CPU usage of the monitored object<br>Obtain its value by checking metric value changes in the **/proc/stat** file in a collection period.<br>Run the **top** command to check the **%Cpu(s)** value.<br>Unit: percent | 0-100 % | BMS | 1 minute |
| load_average5 | (Agent) 5-Minute Load Average | CPU load averaged from the last 5 minutes<br>Obtain its value by dividing the **load5/** value in **/proc/loadavg** by the number of logical CPUs.<br>Run the **top** command to check the **load5** value in the **/proc/loadavg** file. | ≥ 0 | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mem_usedPercent | (Agent) Memory Usage | Memory usage of the monitored object<br>Obtain its value by checking the file **/proc/meminfo**. **Memory Usage** = (**MemTotal** - **MemAvailable**)/**MemTotal**<br>Unit: percent | 0-100 % | BMS | 1 minute |
| mountPointPrefix_disk_free | (Agent) Available Disk Space | Available disk space of the monitored object<br>Run the **df -h** command to check the data in the **Avail** column.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: GB | ≥ 0 GB | BMS | 1 minute |
| mountPointPrefix_disk_usedPercent | (Agent) Disk Usage | Disk usage of the monitored object. It is calculated as follows: **Disk Usage** = **Used Disk Space**/**Disk Storage Capacity**.<br>**Disk Usage** = **Used Disk Space**/**Disk Storage Capacity**<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: percent | 0-100 % | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| mountPointPrefix_disk_ioUtilsandvolumePrefix_disk_ioUtils | (Agent) Disk I/O Usage | Disk I/O usage of the monitored object<br>Obtain its value by checking data changes in the thirteenth column of the corresponding device in the **/proc/diskstats** file in a collection period.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: percent | 0-100 % | BMS | 1 minute |
| mountPointPrefix_disk_inodesUsedPercent | (Agent) Percentage of Total inode Used | Percentage of used index nodes on the disk<br>Run the **df -i** command to check data in the **IUse%** column.<br>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).<br>Unit: percent | 0-100 % | BMS | 1 minute |
| net_bitRecv | (Agent) Inbound Bandwidth | Number of bits received by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: bit/s | ≥ 0 bit/s | BMS | 1 minute |
| net_bitSent | (Agent) Outbound Bandwidth | Number of bits sent by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: bit/s | ≥ 0 bit/s | BMS | 1 minute |

  

| Metric ID | Metric | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|---|---|---|---|---|---|
| net_packetRecv | (Agent) NIC Packet Receive Rate | Number of packets received by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: count/s | ≥ 0 counts/s | BMS | 1 minute |
| net_packetSent | (Agent) NIC Packet Send Rate | Number of packets sent by this NIC per second<br>Check metric value changes in the **/proc/net/dev** file in a collection period.<br>Unit: count/s | ≥ 0 counts/s | BMS | 1 minute |
| net_tcp_total | (Agent) TCP TOTAL | Total number of TCP connections of this NIC | ≥0 | BMS | 1 minute |
| net_tcp_established | (Agent) TCP ESTABLISHED | Number of ESTABLISHED TCP connections of this NIC | ≥0 | BMS | 1 minute |

# 1.6 FAQs

## 1.6.1 Why Doesn't the Cloud Eye Console Display Monitoring Data or Why Is There a Delay in Data Display After Agent Has Been Installed and Configured?

1. After the Agent is installed successfully, server monitoring data is displayed on the Cloud Eye console after two minutes. If **BMS** is not displayed on the **Monitoring Overview** page after five minutes, check whether the time of the BMS is the same as that of the client where you are using the management console.

   The time when the Agent reports data depends on the local time of the BMS. The time when the console delivers requests is related to the browser time of the client. If the two are inconsistent, no monitoring data is displayed on the Cloud Eye console.

2. Log in to the BMS and run the **service telescoped status** command to check the status of Agent. If the following information is displayed, Agent is running properly:

   ```
   Telescope process is running well.
   ```

  

If monitoring data is still not displayed, check the configuration as instructed in **Manually Configuring the Agent for Linux**.

  

# 2 Backup

## 2.1 Overview

### Solution Introduction

To prevent data loss of a system deployed on a BMS caused by external virus intrusions, misoperation, or software bugs, you can use the Cloud Backup and Recovery (CBR) service to back up your BMS periodically and automatically (the minimum backup period is one hour). You can use a backup to quickly restore your BMS in any of the cases mentioned above.

CBR backs up the configurations and all EVS disks of your BMS to highly reliable Object Storage Service (OBS) to ensure the security of your data.

**Figure 2-1** Backup and restoration

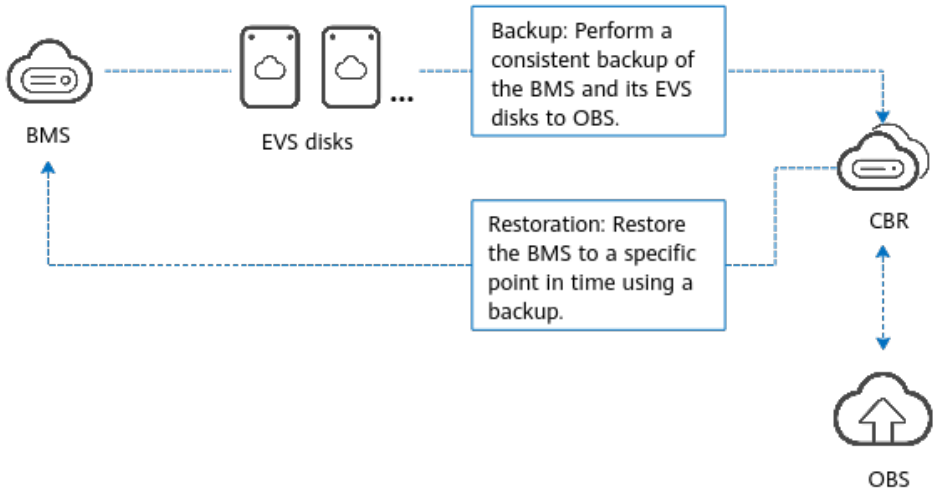


### Scenarios

Enterprise core databases and finance services require high data security. In such cases, you are advised to back up BMSs to improve data reliability.

  

### Advantages

- Easy: You can create a policy easily for online backup.
- Efficient: Incremental backups shorten the time required for backup by 95%. With instant restoration, the RPO is as low as 1 hour and RTO is only several minutes.
- Reliable: A consistent backup of multiple disks ensures data security and reliability.

### Constraints

- A BMS backup cannot be used to create an image.
- When a BMS is restored using a backup, it will be stopped automatically, which will interrupt user services. After it is stopped, it will be locked for a period of time during which users cannot perform any operations on it.

### Prerequisites

1. Create a key pair.

   For system security, you are advised to use a key pair to authenticate users who attempt to log in to a BMS. You can use an existing key pair or create a new one for remote login authentication.

   For details, see **Using an SSH Key Pair**. If you already have a key pair, skip this step.
2. Create a VPC.

   BMS uses networks provided by a Virtual Private Cloud (VPC), including subnets and security groups.

   For details, see **Creating a VPC with a Subnet**.

### Procedure

To back up a BMS, do as follows:

1. **Create a backup policy.** Set the backup time, backup period, and retention rules to back up a BMS automatically.
2. **Buy a server backup vault.** Buy a vault to store backups.
3. **Associate your server with a vault.** Associate your BMS with the vault to back up and replicate data.
4. **Create a server backup.** Back up your BMS manually.
5. **Check backups and restore data.** Check the backups created automatically or manually on the management console. Restore your BMS to a specific point in time as you need.

## 2.2 Creating a Backup Policy

1. Log in to the **CBR console**.
2. Choose **Policies** in the navigation pane. In the upper right corner, click **Create Policy**. **Figure 2-2** shows an example.

  

**Figure 2-2** Creating a backup policy

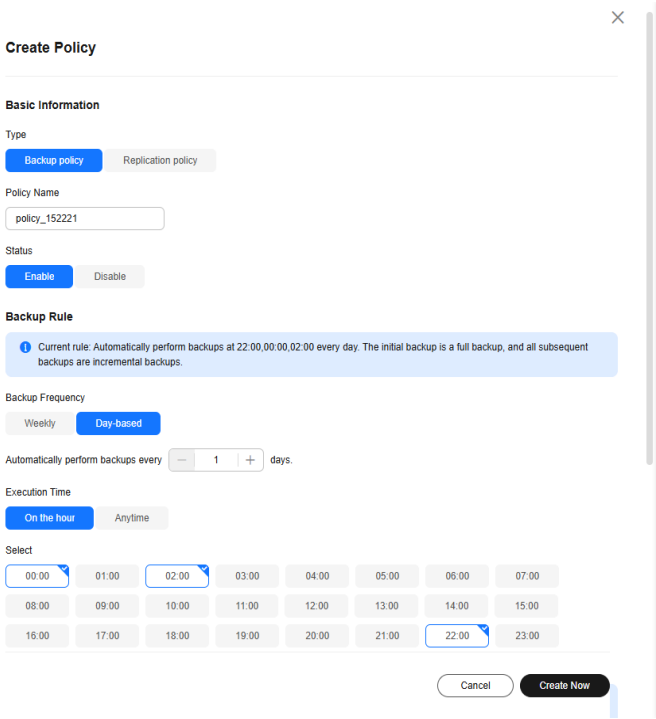


3. Set the backup policy parameters.

**Table 2-1** Parameter description

| Parameter | Description | Example Value |
| --- | --- | --- |
| Type | Policy type. | Backup |
| Policy Name | Backup policy name. You can enter a custom name or use the default name **policy**_*xxxx*.<br>A name contains a maximum of 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed. | backup_policy |
| Status | Whether to enable the backup policy. By default, a backup policy is enabled.<br>CBR backs up servers and EVS disks to a vault and deletes expired backups only after a backup policy is applied to the vault. | Enable |

  

| Parameter | Description | Example Value |
| --- | --- | --- |
| Backup Frequency | Frequency for executing a backup task. By default, a backup task is executed automatically every Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday.<br>• **Weekly**<br>Specify on which days of each week a backup task will be executed. You can select multiple days.<br>• **Day-based**<br>Specify the interval (1 to 30 days) for executing a backup task.<br>If you select **Day-based**, the first backup time is supposed to be on the day when the backup policy was created. If the specified backup time has passed when you create a backup policy, the initial backup will be performed in the next backup cycle.<br>You are advised to execute backup tasks during off-peak hours or when there are no services running. | Every 1 day |

  

| Parameter | Description | Example Value |
| --- | --- | --- |
| Execution Time | Execution time of a backup task in a day. The default backup time is 22:00.<br>You can select multiple times, either on or off the hour, for backup. If you select a time that is not on the hour, you can customize the time (0 to 59 minutes). For example, you can select a time that is not on the hour, such as 00:10 and 01:10.<br>You are advised to execute backup tasks during off-peak hours or when there are no services running.<br>The backup service experiences peak usage between 22:00 and 08:00, during which delays may occur. To ensure optimal performance, it is recommended that you evaluate your service types and stagger backups across discrete time periods.<br>**NOTICE**<br>• There may be a time difference between the scheduled backup time and the actual backup time.<br>• If you have a large amount of data to back up, you are advised to set a less frequent backup schedule. If a backup task takes longer than the backup interval, the system will skip the next backup execution time.<br>For example, a disk is scheduled to be backed up at 00:00, 01:00, and 02:00. A backup task starts at 00:00. Because a large amount of incremental data needs to be backed up or a heap of backup tasks are executed at the same time, this backup task takes 90 minutes and completes at 01:30. The scheduled backup at 01:00 will be skipped. There will be only two backups generated, one at 00:00, and the other at 02:00.<br>• The execution times refer to the local times of clients, not the time zone or times of the region. | 00:00, 02:00 |

  

| Parameter | Description | Example Value |
| --- | --- | --- |
| Full Backup | Whether to perform periodic full backups. By default, full backup is disabled.<br>• **Enable**<br>Enabling full backup improves your data reliability, but full backups will use more storage space.<br>Configure a full backup frequency. The value ranges from **0** to **100**. **0** means that a full backup will be performed in every backup task.<br>• **Disable**<br>Incremental backups will be performed based on the backup policy.<br>**NOTICE**<br>• A full backup usually takes a long period of time. If a full backup of a resource is in progress, other policy or manual backups of this resource will not be performed. You are advised to back up data during off-peak hours.<br>• When backups are kept by quantity, full backups can be performed only when the full backup frequency configured is less than the number of retained backups.<br>• If full backup is not enabled, to ensure data security, a full backup is performed after 365 incremental backups by default. | 7 |

  

| Parameter | Description | Example Value |
| --- | --- | --- |
| Retention Rule | Rule that specifies how backups will be retained. By default, backups are retained for one month.<br>• **Time period**<br>You can choose to retain replicas for one month, three months, six months, or one year, or for any desired number (**2** to **99999**) of days.<br>• **Backup quantity**<br>You can set the maximum number of backups to retain for each cloud server. The value ranges from **2** to **99999**.<br>• **Advanced Options**<br>You can also set long-term retention rules with advanced options. Long-term retention rules and quantity-based retention rules will be both applied.<br>– **Day-Based**: The most recent backup of each day is retained. Range: 0–100<br>– **Weekly**: The most recent backup of each week is retained. Range: 0–100<br>– **Monthly**: The most recent backup of each month is retained. Range: 0–100<br>– **Yearly**: The most recent backup of each year is retained. Range: 0–100<br>For example, if you select **Day-Based**, the system retains the most recent backup each day. If a disk is backed up for multiple times in a day, only the most recent backup of that day is retained. If you set the value for **Day-Based** to **5**, the system retains the most recent backup from each of the last five days that have backups generated. If there are more than five backups, the earliest backup will be deleted automatically. If **Day-Based**, **Weekly**, **Monthly**, and **Yearly** are all configured, the union | Select **Backup quantity** and set to keep 3 backups. In addition, select **Advanced Options** and set to keep the most recent backup from each of the last two weeks. If today is the 30th of a month, the execution of this policy is shown in **Figure 2-4**.<br>• Dates with a time indicate the days that have backups generated.<br>• Dates with the time in gray indicate that the backups have been deleted.<br>• Dates with the time in green indicate that the backups are retained.<br>If the weekly retention rule in advanced options is not configured, only the backups generated on the 25th, 26th, and 29th will be retained. |

  

| Parameter | Description | Example Value |
| --- | --- | --- |
| | backups will be selected for retention. For example, if **Day-Based** is set to **5** and **Weekly** to **1**, five backups will be retained. The long-term retention rule and the quantity-based retention rule will both apply.<br>• **Permanent** | |

  

| Parameter | Description | Example Value |
| --- | --- | --- |
| | **NOTE**<br>• The system automatically deletes the earliest and expired backups every other day to avoid exceeding the maximum number of backups to retain or retaining any backup longer than the maximum retention period.<br>• Retained backups may not be deleted immediately after the expiration time specified in the backup policy. There may be a delay. Generally, they are deleted in batches **from 8:00 to 20:00 after the expiration time**. For example, if a backup expired at 20:00 on November 23, 2024, it was deleted from 08:00 to 20:00 on November 24, 2024. In this way, backup data can be deleted during off-peak hours.<br>• A retention rule only applies to backups generated by automatically executing a backup policy. It does not take effect for those generated by manually executing a backup policy. You can manually delete them from the backup list.<br>• If a backup is used to create an image, the backup will not be deleted by a retention rule. Instead, it will be forcibly retained. After the image created from the backup is deleted, the retention rule will apply to the backup. That is, if the backup expires or is not within the most recent backups, it will be deleted automatically.<br>• A maximum of 10 backups can be retained for failed periodic backup tasks. They are retained for one month and can be deleted manually.<br>• A backup cannot be deleted before the subsequent backup task is complete. | |

  

**Figure 2-3** Retention rule with advanced options configured

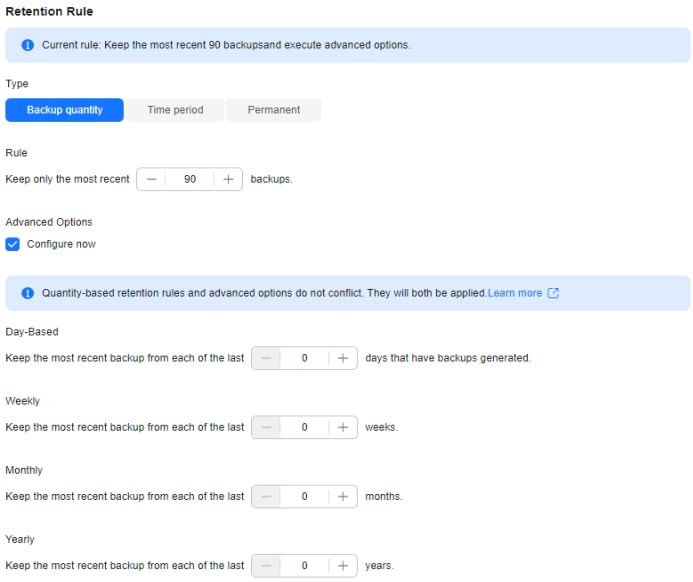


**Figure 2-4** Retention rule execution example

| Sun | Mon | Tues | Wed | Thur | Fri | Sat |
|---|---|---|---|---|---|---|
| | | | | **1** 23:00 | 2 | 3 |
| **4** 23:00 | **5** 23:00 | 6 | 7 | **8** 23:00 | 9 | 10 |
| **11** 23:00 | **12** 23:00 | 13 | 14 | **15** 23:00 | 16 | 17 |
| **18** 23:00 | **19** 23:00 | 20 | 21 | **22** 23:00 | 23 | 24 |
| **25** 23:00 | **26** 23:00 | 27 | 28 | **29** 23:00 | 30 | 31 |

NOTE

Frequent backups, more retained backups, or a longer retention can improve data protection, but those backups will also take up more space. Configure an appropriate backup policy based on the data importance and service volume.

4. Click **Create Now**. After the backup policy is created, you can see it in the backup policy list.

# 2.3 Buying a Server Backup Vault

1. Go to the **Buy Server Backup Vault**.
2. Set parameters as instructed in the tables below. For parameters not included in the tables, retain the default values.

   Complete **Basic Settings** and select a **Protection Type**.

  

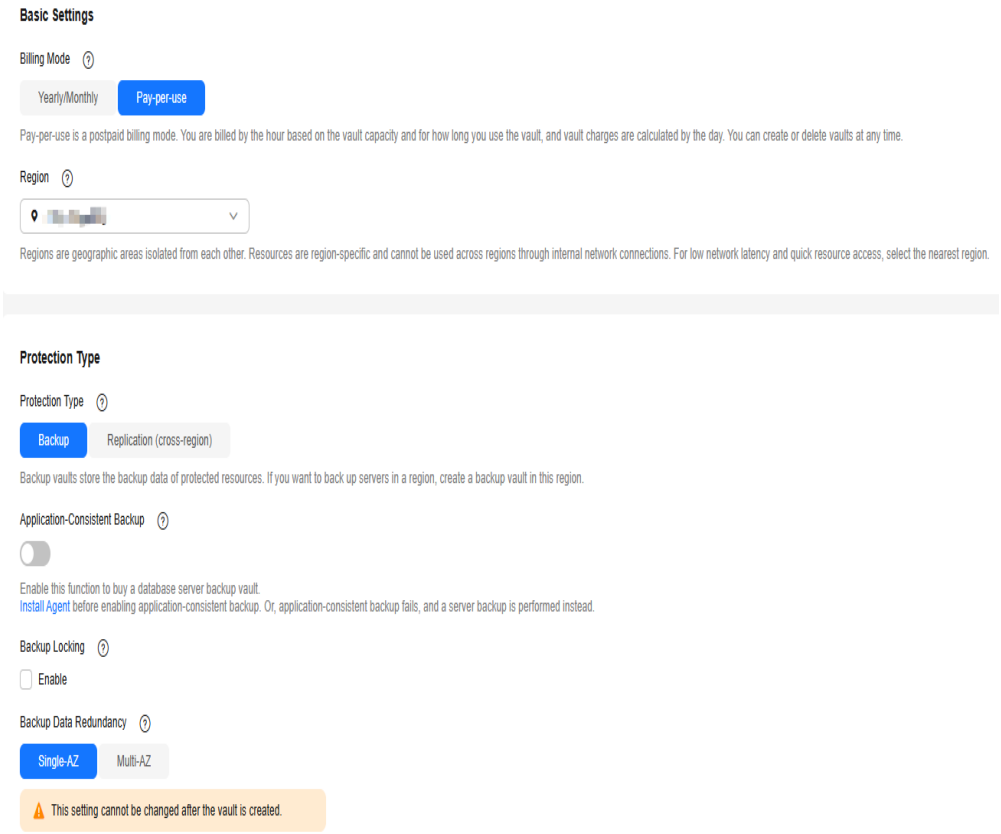


| Parameter | Description | Example Value |
| --- | --- | --- |
| Billing Mode | • Pay-per-use is a postpaid billing mode. You are billed based on your resource usage. You can buy or delete vaults anytime. Expenditures are deducted from the account balance.<br>• Yearly/Monthly is a prepaid billing mode. You are billed based on the subscription duration you specify. This mode provides lower prices and is ideal when the resource use duration is predictable. | **Pay-per-use** |
| Region | Resources in different regions cannot communicate with each other over internal networks. For lower latency and faster access, select the region nearest to you.<br>Once a vault is created, its region cannot be changed. | CN-Hong Kong |

  

| Parameter | Description | Example Value |
| --- | --- | --- |
| Protection Type | • **Backup**: The vault will be used to store server backups.<br>• **Replication (cross-region)**: The vault will be used to store replicas of server backups. If you select **Replication (cross-region)**, you do not need to select a server.<br>For example, if you want to back up a server, select **Backup** for the vault protection type. If you want to replicate backups of a server from one region to another, select **Replication (cross-region)** for the vault in this other region. | Backup |
| Backup Data Redundancy | • **Single-AZ**: Backup data is stored in a single AZ, with lower costs.<br>• **Multi-AZ**: Backup data is stored in multiple AZs for higher reliability. If an AZ is unavailable, backup data can still be accessed from other AZs.<br>Once a vault is created, its backup data redundancy policy cannot be changed. Plan and select a policy that best suits your service needs. | Single-AZ |

Configure **Resources to Back Up**, **Backup Capacity**, and **Backup Policy**.

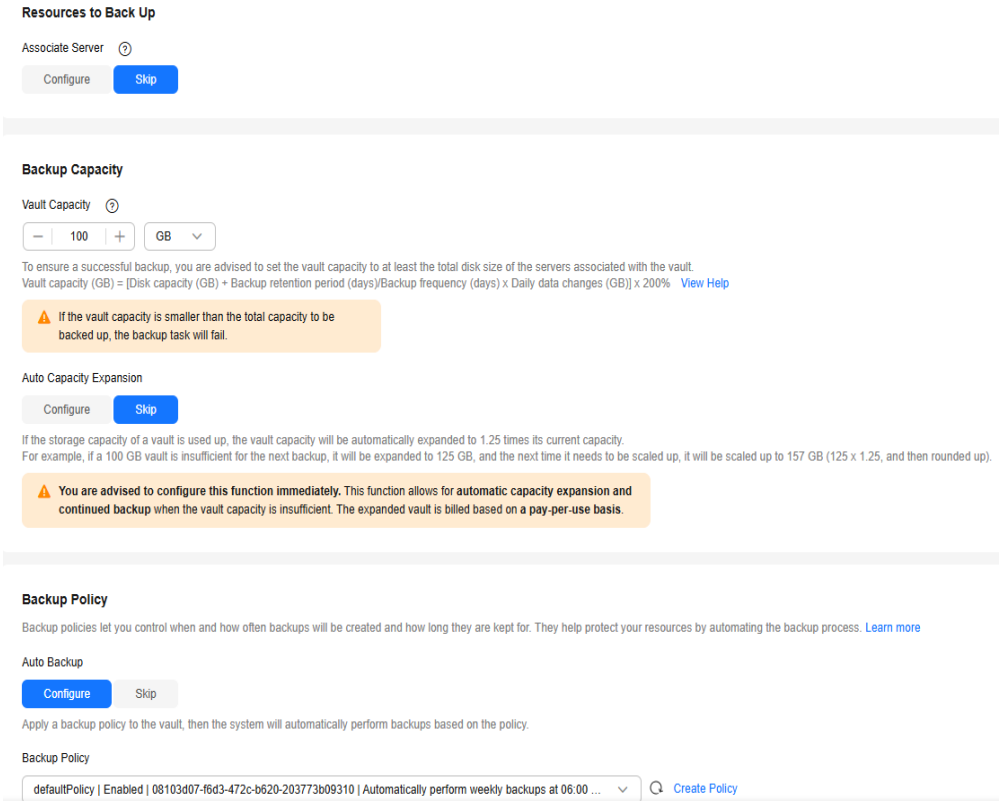


  

| Parameter | Description | Example Value |
| --- | --- | --- |
| Resources to Back Up | • **Configure**: Select the servers to associate.<br>• **Skip**: If no server is available, you can purchase a vault first and associate servers later. | Skip |
| Backup Capacity | Set the vault capacity. The value ranges from 10 GB to 10,485,760 GB.<br>You need to **properly plan the vault capacity** and ensure that it is not smaller than the size of the servers you want to back up. If automatic association is enabled and a backup policy is applied to the vault, plan more capacity as needed.<br>As the vault's used space grows, you can expand the vault capacity if it becomes insufficient. | 100 GB |
| Auto Capacity Expansion | If you select **Configure** for **Auto Capacity Expansion**, the vault capacity will be expanded automatically to 1.25 times the size of the original capacity when the capacity is about to be used up.<br>For example, if the vault capacity is 100 GB but more than 100 GB is required by a backup, the vault capacity will be expanded to 125 GB automatically. When the 125 GB of capacity is about to be used up, the vault capacity will be expanded to 157 GB. | Skip |
| Auto Backup | • **Configure**: A backup policy can be applied to this vault. Servers associated with this vault will be backed up automatically based on the policy. You can select an existing backup policy or create a new one. For example, select the backup policy created in **Creating a Backup Policy**.<br>• **Skip**: Servers associated with this vault will not be backed up automatically. If needed, you can create a policy and apply it to this vault later to back up servers periodically and automatically. For details, see **Policy Management**. | Configure |
| Automatic Association | • **Configure**: In the next backup cycle, CBR will automatically scan for unprotected servers, associate them with the vault, and perform backups.<br>• **Skip**: To enable automatic association later, see **Associating Resources with a Vault**. | Skip |

Configure **Enterprise Project** and **Advanced Settings**.

  

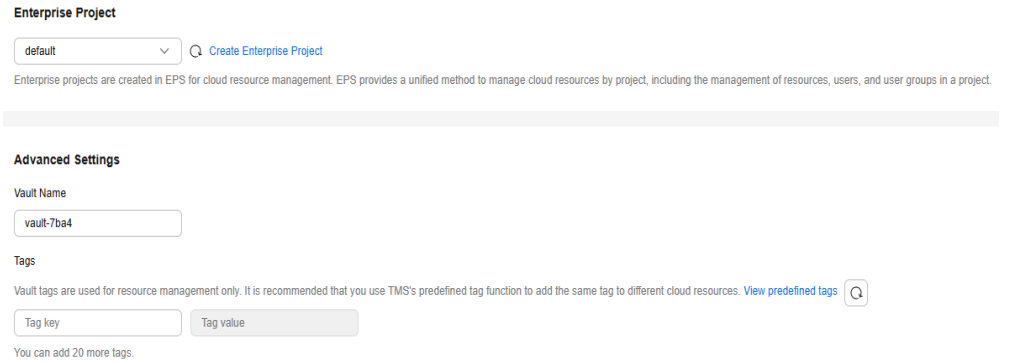


| Parameter | Description | Example Value |
| --- | --- | --- |
| Enterprise Project | Add the vault to an existing enterprise project.<br>This parameter is only available for enterprise users who have enabled enterprise project management.<br>**NOTE**<br>If the **CBR FullAccess** permissions have been assigned to IAM users, enterprise projects will not be displayed for you to choose from when you create a vault. Go to the Enterprise Project Management console and assign the **CBR FullAccess** permissions to the target user group. | default |
| Vault Name | Enter a name for the vault.<br>A name contains a maximum of 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed. You can also use the default name **vault**_*xxxx*. | vault-317a |

3. Click **Next**.
4. Click **Submit**.

   Go back to the **Cloud Server Backups** page, you can view the created vault in the vault list. The vault status is **Available**.

# 2.4 Associating Servers with a Vault

After purchasing a vault, associate servers with it.

The selected servers must have not been associated with any vault and are in the **Running** or **Stopped** state.

1. On the **Cloud Server Backups** page, locate the vault and click **Associate Server** in the **Operation** column.
2. Select the BMSs you want to associate with the vault. You can choose to back up entire servers or select some disks for backup.

  

**NOTICE**

- If a new disk is attached to an associated server, the system automatically identifies the new disk and includes it in subsequent backup tasks.
- To avoid data inconsistency after restoration, you are advised to back up entire servers.

  If you want to back up only some of the disks to reduce costs, ensure that data on these disks does not depend on other disks. Or, data inconsistency may occur.

  For example, data of an Oracle application is scattered across different disks. If only some of the disks are backed up, restoration restores only the data of the disks that have been backed up, with the rest data unchanged. As a result, data inconsistency occurs and even the application cannot be started.
- You are not advised to associate a shared disk with multiple servers for cloud server backup. If multiple servers need to be backed up, associate the vault with only certain disks and do not associate it with shared disks. For details, see **Associating Resources with a Vault**.

3. Click **OK**. Then on the **Associated Servers** tab, you can view the servers that have been associated with.

# 2.5 Creating a BMS Backup

1. On the **Cloud Server Backups** page, click the **Vaults** tab and locate the vault that is associated with the server.
2. Click **Perform Backup** in the **Operation** column. In the server list, select the server you want to back up. After a server is selected, it is added to the list of selected servers.
3. Set **Name** and **Description** for the backup.

| Parameter | Description | Example Value |
| --- | --- | --- |
| Name | Name of the backup to be created.<br>A name contains a maximum of 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed.<br>**NOTE**<br>You can also use the default name **manualbk**_*xxxx*.<br>If multiple servers are backed up, the system automatically adds suffixes to their backup names, for example, **backup-0001** and **backup-0002**. | manualbk_d819 |
| Description | Description of the backup to be created.<br>It cannot exceed 255 characters. | - |

4. Determine whether to enable **Full Backup**. If full backup is enabled, the system will perform a full backup for every associated server. This requires a larger capacity compared with an incremental backup.

  

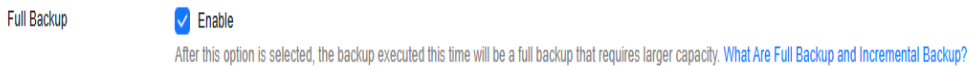


**NOTE**

By default, the first backup is a full backup, and subsequent backups are incremental backups.

For data security, after 100 incremental backups are performed, a full backup will be performed by default.

For details, see **What Are Full Backup and Incremental Backup?**

5. Click **OK**. The system automatically creates a backup for the server.

   On the **Backups** tab, when the status of the backup changes to **Available**, the backup task is successful.

**NOTE**

When manually backing up a cloud server, you can restart the server after the backup progress exceeds 10%. However, to ensure data integrity, you are not advised to restart it until the backup is complete.

# 2.6 Checking Backups and Restoring Data

## Checking Backups Created Automatically or Manually

You can check all backups on the **Backups** tab, as shown in **Figure 2-5**.

Names of backups created automatically start with **autobk_** and those created manually start with **manualbk_**.

**Figure 2-5** Backup list

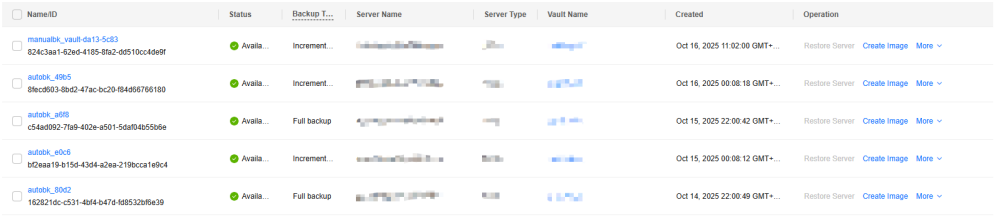


You can click a server backup name to view the backup details, including disk backup details.

**Figure 2-6** Backup details

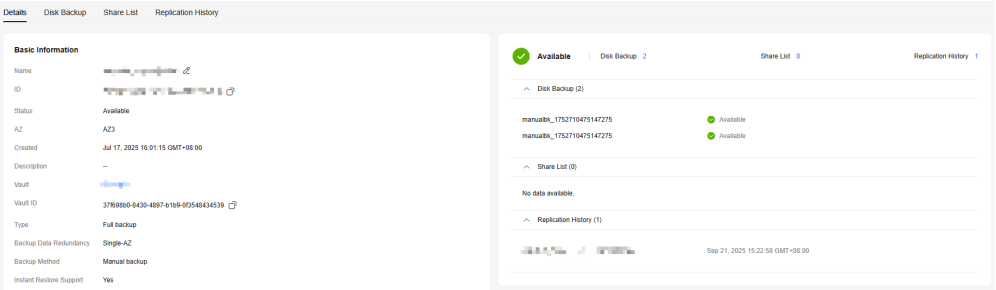


  

For your convenience, backups are also displayed on the **Disks** tab of your BMS details page.

## Restoring Data

When disks on a server are faulty or their data is lost, you can use a backup to restore the server to its state when the backup was created.

1. Click the **Backups** tab and locate the desired backup.
2. Click **Restore Server** in the **Operation** column, as shown in **Figure 2-7**.

**WARNING**

- The current server data will be overwritten by the data captured at the time of backup. **The restoration cannot be undone.**
- Servers will be shut down during restoration. It is recommended that you perform restoration during off-peak hours.

**Figure 2-7** Restoring a server

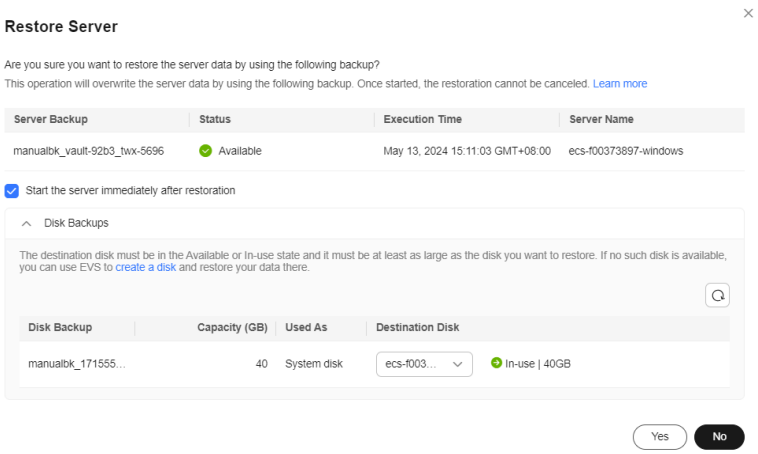


3. (Optional) Deselect **Start the server immediately after restoration**.

   In this case, you need to start the server manually after the restoration is complete.
4. In the **Destination Disk** drop-down list, select the disk that the backup will be restored to.

**WARNING**

If the number of disks to be restored is greater than the number of disks that were backed up, restoration may cause data inconsistency.

For example, if Oracle data is scattered across multiple disks and only some of the disks are restored, data may become inconsistent and the application may fail to start.

  

NOTE

- If a server has only one disk, the backup is restored to the only disk by default.
- If a server has multiple disks, the backup is restored to the original disks by default. You can also restore the backup to other disk by selecting that disk from the **Destination Disk** drop-down list. That disk must not be smaller than the original disk.
- Data on data disks cannot be restored to system disks.

5. Click **OK** and confirm the restoration is successful.

   In the backup list, check the restoration status. When the backup status changes to **Available** and there are no new failed restoration tasks in **Tasks**, the restoration is successful. Data is restored to the state when the backup was created.

  