

Bare Metal Server

FAQs

Issue 09
Date 2022-09-30



Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks 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 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 General FAQ.....	1
1.1 What Restrictions Are There on BMS Use?.....	1
1.2 How Are BMSs Different from ECSs?.....	2
1.3 What Are the Differences Between BMSs and Traditional Physical Servers?.....	2
1.4 What Are the Differences Between BMS and Dedicated Host (DeH)?.....	2
1.5 How Do BMSs Ensure Data Security?.....	3
1.6 Can I Use APIs to Access BMSs?.....	3
1.7 What OS Images Does BMS Support?.....	3
1.8 Will Services Be Affected If Hyper-Threading Is Configured for BMS?.....	4
1.9 How Do I View and Increase the BMS Quota?.....	4
1.10 How Do I Synchronize the Time of a BMS?.....	5
2 Instance FAQ.....	6
2.1 How Long Does It Take to Create a BMS?.....	6
2.2 What Do I Do If I Can't Find My BMS on the Management Console?.....	6
2.3 How Can I Obtain board_type of a BMS Flavor?.....	6
2.4 Why Is the BMS Creation Task Displayed as Failed But the BMS List Shows the BMS?.....	7
2.5 How Do I Create a BMS That Can be Quickly Provisioned?.....	8
2.6 What Are the Advanced Features of BMSs Using EVS Disks?.....	8
2.7 Can I Transfer a BMS to Another Account?.....	9
2.8 Is the BMS Host Name with Suffix novalocal Normal?.....	9
2.9 Why Does the System Display a Message Indicating That the BMS Cannot Be Deleted?.....	10
2.10 How Do I Monitor BMS Metrics?.....	10
2.11 How Do I Create an Agency for Server Monitoring of the BMS?.....	10
3 Billing FAQ.....	12
3.1 Where Can I Query the BMS Price?.....	12
3.2 What Are Pre-payment and Post-payment? How Do I Choose Between Them?.....	12
3.3 What Is the Impact on Billing if I Unsubscribe from a BMS?.....	12
3.4 How Do I Set Automatic Renewal for BMSs Billed on a Yearly/Monthly Basis?.....	13
3.5 When Will a BMS Be Released After It Expires?.....	14
3.6 How Do I Unfreeze a BMS?.....	14
4 Key Pair and Password FAQ.....	15
4.1 How Do I Change the OS Password of a BMS?.....	15

5 Login FAQ.....	18
5.1 How Do I Prepare to Log In to a BMS?.....	18
5.2 What Can I Do If I Selected the Key Pair Authentication When I Created a BMS But Want to Log In to the BMS Using a Password?.....	18
5.3 What Do I Do If I Can't Log In to a Linux BMS?.....	19
5.4 What Browser Versions Can Be Used to Remotely Log In to a BMS?.....	21
5.5 What Do I Do If the Login Page Does Not Respond?.....	22
5.6 What Do I Do If the BMS Console Is Displayed Improperly After I Remotely Log In to a BMS?.....	24
5.7 What Do I Do If the Numeric Keypad Does Not Work During Remote Login?.....	25
5.8 What Do I Do If the SSH Login or Data Transmission Is Slow?.....	25
6 Network and Security FAQ.....	27
6.1 Can BMSs of Different Accounts Communicate with Each Other over an Internal Network?.....	27
6.2 How Do Two BMSs in the Same Region But in Different AZs Communicate with Each Other?.....	27
6.3 Are My BMSs in the Same Subnet?.....	27
6.4 Can I Associate a BMS with Multiple Security Groups?.....	27
6.5 Can BMSs Communicate with ECSs in the Same VPC?.....	28
6.6 What Are the Differences Between the Primary and Extension NICs of BMSs?.....	28
6.7 Can I Bind Multiple EIPs to a BMS?.....	28
6.8 Can I Configure an EIP?.....	28
6.9 Will I Obtain an EIP That Has Been Released?.....	28
6.10 What Are the Differences Between EIPs, Private IP Addresses, and Virtual IP Addresses?.....	29
6.11 How Can I Modify the Network Configuration or Restart the Network If I Can Log In to a BMS Using Only SSH?.....	29
6.12 How Do I Handle the Failure to Ping a CentOS 7 Extension NIC?.....	29
6.13 What Do I Do If the Communication Between the Primary NIC and Extension NIC of the BMS is Abnormal?.....	30
6.14 How Can I Configure a Static IP Address for a BMS?.....	31
6.15 What Do I Do If the Network Performance Is Poor When 128 Ethernet Network Devices Are Bound to a BMS?.....	32
7 Disk FAQ.....	34
7.1 Can EVS Disks Be Attached to BMSs? How Many Data Disks Can Be Attached to a BMS?.....	34
7.2 What Are the Restrictions for Attaching a Disk to a BMS?.....	34
7.3 How Do I Know Whether EVS Disks Are Available in a Flavor?.....	35
7.4 How Do I Change the Disk Identifier in the fstab file to UUID?.....	35
7.5 How Do I Obtain the Drive Letter of an EVS Disk?.....	36
7.6 Are the EVS Disk Device Names on the Console and the Device Names in BMS OSs Consistent?.....	36
7.7 Why Is the EVS Disk Size Not Updated in the BMS OS After the EVS Disk Capacity Has Been Expanded?.....	40
7.8 How Can I Restore System Disk Data Using the Snapshot?.....	40
7.9 What Do I Do to Prevent Risks of Attaching or Detaching the System Disk?.....	41
7.10 How Do I Select Storage?.....	41
7.11 Why Is the Disk Capacity Displayed in the BMS OS Less Than That Displayed on the Official Website?.....	41

8 OS FAQ.....	42
8.1 Can I Install or Upgrade BMS OSs by Myself?.....	42
8.2 Can I Change the OS of a BMS?.....	42
8.3 Is a GUI Provided for BMS OSs?.....	42
8.4 Is an Upload Tool Delivered with BMS OSs?.....	42
8.5 Does the Public Image Used to Create a BMS Have a Swap Partition?.....	42
8.6 How Do I Configure the Static Host Name of a BMS?.....	43
8.7 How Do I Set the Password Validity Period?.....	45
8.8 How Do I Set SSH Configuration Items?.....	46
8.9 How Can I Handle the Eight-Hour Difference Between the Windows BMS and Local Time.....	47
8.10 How Do I Change the SID of a Windows Server 2012 BMS?.....	48
8.11 How Do I Change the Kernel Version of CentOS 7 BMSs?.....	50
8.12 How Do I Reserve Log Space If the Root Partition Automatically Expands Disks?.....	51
8.13 How Do I Roll Back the Kernel Version If I Mistakenly Upgrade the Kernel?.....	55
8.14 How Do I Increase the Swap Partition Size?.....	56
8.15 How Do I Increase the Size of the Root Partition of a BMS Which Is Quickly Provisioned?.....	57
8.16 Common Linux Commands.....	59
8.17 How Do I Update the Disk Metadata After the LVM Volume Is Remounted?.....	60
8.18 How Do I Handle a Network Failure After Services Are Switched from a Windows BMS Booted from an EVS Disk to an HA BMS?.....	60

1 General FAQ

1.1 What Restrictions Are There on BMS Use?

- External hardware devices (such as USB devices, bank USB keys, external hard disks, and dongles) cannot be loaded.
- Out-of-band management is not supported. Your BMSs are managed and maintained by Huawei Cloud.
- Live migration is not supported. If a BMS is faulty, your services running on it may be affected. It is good practice to deploy your services in a cluster or in primary/standby mode to ensure high availability.
- You cannot create a server without an OS, that is, a BMS must have an OS.
- After a BMS is created, you cannot change its VPC.
- Virtualization software cannot be installed on a BMS.
- When you create a BMS, you can only select a flavor with specified CPU, memory, and local disks but cannot configure them separately. After a BMS is created, you can expand the capacity of attached EVS disks but cannot modify the BMS CPU, memory, or local disks.
- You can only attach EVS disks whose device type is **SCSI** to a BMS.
- You cannot attach EVS disks to BMSs of certain flavors or BMSs created from certain images because these BMSs do not have SDI iNICs or lack compatibility.
- Do not delete or modify built-in plug-ins of an image, such as Cloud-Init and bms-network-config. Otherwise, basic BMS functions may be unavailable.
- If you choose to assign an IP address automatically when you create a BMS, do not change the private IP address of the BMS after the BMS is provisioned. Otherwise, the IP address may conflict with that of another BMS.
- BMSs do not support bridge NICs because they will cause network interruptions.
- Do not upgrade the OS kernel. Otherwise, the hardware driver may become incompatible with the BMS and adversely affect the BMS reliability.
- You can stop a BMS only on the management console and cannot run **shutdown** to stop it.

- Quotas are enforced for service resources on the platform to prevent unforeseen spikes in resource usage. The BMS service has no independent quota. It shares the number of instances, CPU cores, and memory with the ECS service. You can view BMS quota in the **Elastic Cloud Server** row. For details see [How Do I View My Quotas?](#)

1.2 How Are BMSs Different from ECSs?

BMSs allow users to exclusively access physical resources but ECS users can only share these resources. BMSs are recommended for deploying mission-critical applications and services that require high performance and a secure and reliable running environment, such as big data clusters and enterprise middleware systems.

For details about the differences between BMSs, physical servers, and ECSs, see [Comparison Between BMSs, Physical Servers, and ECSs](#).

1.3 What Are the Differences Between BMSs and Traditional Physical Servers?

Compared with traditional physical servers, BMSs support automatic provisioning, automatic O&M, communication through the VPC, and interconnection with shared storage. BMSs have all the features and advantages of physical servers. Your applications can access the physical CPU and memory without any virtualization overhead.

For details about the differences between BMSs, physical servers, and ECSs, see [Comparison Between BMSs, Physical Servers, and ECSs](#).

1.4 What Are the Differences Between BMS and Dedicated Host (DeH)?

Both BMS and DeH allow you to exclusively use physical servers but they have the following differences:

- BMS uses the bare metal architecture and does not provide a virtualization platform by default. You need to install virtualization software on a BMS before you provision an ECS on it.
- DeH provides a virtualization platform by default. After purchasing a DeH, you can directly provision ECSs on it.

For more information, see [Table 1-1](#).

Table 1-1 Comparison between BMS and DeH

Item	BMS	DeH
Whether virtualization is provided	No	Yes

Item	BMS	DeH
How to use	Use each BMS as a physical server.	Provision ECSs on a DeH.
Specifications	BMS specifications	DeH specifications and ECS specifications
Image	BMS images	ECS images

1.5 How Do BMSs Ensure Data Security?

- BMSs offer physical server-level performance and isolation. They provide dedicated computing resources without any loss due to virtualization. The high-performance and high-reliability BMSs can ensure the security of data stored on them.
- If a BMS uses local disks, RAID can be used to improve fault tolerance and ensure data security.
- If a BMS uses EVS disks, the EVS disks can be backed up based on the consistency snapshot technology. You can use the backups to restore BMS data, ensuring data security and reliability.

1.6 Can I Use APIs to Access BMSs?

Yes.

The BMS service provides various APIs, including BMS APIs and native OpenStack APIs. If you are familiar with network service protocols and programming languages, you can use BMS APIs to manage your cloud resources.

For details, see [Bare Metal Server API Reference](#).

1.7 What OS Images Does BMS Support?

BMS supports mainstream Linux and Windows image versions.

BMS supports the following Linux images:

- Red Hat Linux Enterprise 6.5/6.7/6.8/6.9/7.2/7.3/7.4/7.5 64-bit
- CentOS 6.8/6.9/7.2/7.3/7.4 64-bit
- SUSE Linux Enterprise Server 11.4/12.1/12.2/12.3 64-bit
- Oracle Linux Server Release 6.8/6.9/7.2/7.3/7.4 64-bit
- Ubuntu 14.04/16.04 64-bit
- Debian 8.6 64-bit
- EulerOS 2.2/2.3 64-bit

BMS supports the following Windows images:

- 2016 Standard 64

- 2012 R2 Standard 64-bit

Note that the OSs supported for BMSs of different flavors vary. For details, see [OSs Supported by Different Types of BMSs](#).

1.8 Will Services Be Affected If Hyper-Threading Is Configured for BMS?

Yes. Hyper-threading is configured in the BMS BIOS, during which the BMS needs to be restarted. As a result, the services are affected.


1.9 How Do I View and Increase the BMS Quota?

What Is Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of ECSs or EVS disks that can be created.

If the existing resource quota cannot meet your service requirements, you can apply for a higher quota.

How Do I View My Quotas?

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, choose **Resources > My Quotas**. The **Service Quota** page is displayed.
4. View the used and total quota of each type of resources on the displayed page.
If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources > My Quotas**. The **Service Quota** page is displayed.
3. Click **Increase Quota** in the upper right corner of the page.
4. On the **Create Service Ticket** page, configure parameters as required.
In the **Problem Description** area, fill in the content and reason for adjustment.
5. After all necessary parameters are configured, select **I have read and agree to the Ticket Service Protocol and Privacy Statement** and click **Submit**.

1.10 How Do I Synchronize the Time of a BMS?

Time consistency is important to BMSs and may affect task execution results. Use the Huawei Cloud NTP server for time synchronization. For details, see [Does Huawei Cloud Provide the NTP Server and How Can I Install It?](#)

2 Instance FAQ

2.1 How Long Does It Take to Create a BMS?

Generally, a Linux BMS is created within 30 minutes and a Windows BMS is created within one to two hours. If a BMS supports quick provisioning, it can be created within about five minutes.

2.2 What Do I Do If I Can't Find My BMS on the Management Console?

Symptom

I have purchased a BMS but cannot find it on the management console.

Causes

Your BMS is not in the selected region or project.

Solution

1. Log in to the BMS management console.
2. In the upper left corner of the management console, select the region where your BMS is. Then, your BMS is displayed.

If your BMS is purchased under a project in a region, switch to that project to view the BMS.

2.3 How Can I Obtain board_type of a BMS Flavor?

Symptom

The public images supported by BMSs vary depending on BMS flavors. You can view the public images supported by each flavor on the management console or using the [Querying Images](#) API. When you use the API to query for the public

images supported by a BMS flavor, you need to enter **board_type** of the flavor. This section describes how to obtain the value of **board_type**.

Solution

The format of a BMS flavor is **physical.AB.C**, for example, **physical.s1.large**.

In the flavor format:

- *A* specifies the BMS type. For example, **s** indicates a general-purpose BMS, **c** a computing BMS, and **m** a memory-optimized BMS.
- *B* specifies the BMS serial number. For example, **1** in **s1** indicates the first generation of general-purpose BMSs.
- *C* specifies the flavor size, such as medium, large, or xlarge.

Use *AB+First one or more letters of C* as the value of **board_type**. For example, if the flavor is **physical.s1.large**, **board_type** is **s1l**. For **board_type** of more flavors, see [Table 2-1](#).

Table 2-1 board_type of a BMS flavor

BMS Flavor	board_type
physical.m2.medium	m2m
physical.h2.large	h2l
physical.hs2.large	hs2l
physical.io2.xlarge	io2xl
physical.kl1.3xlarge	kl13xl

2.4 Why Is the BMS Creation Task Displayed as Failed But the BMS List Shows the BMS?

Symptom

After you applied for a BMS configured with an EIP on the management console, the BMS application request was successfully processed but the EIP could not be bound to the BMS due to insufficient EIPs. In this case, **Failed** will be displayed for the task in the **Task Status** area. However, the BMS that you applied for will be displayed in the BMS list.

Root Cause

- The BMS list shows all the BMSs whose application requests have been processed.
- The **Task Status** area shows the processing status of the BMS creation task, including statuses of sub-tasks, such as preparing BMS resources and binding an EIP. Only when all subtasks have succeeded, the task status changes to **Succeeded**. Otherwise, the task status is displayed as **Failed**.

The BMS is only temporarily displayed in the BMS list. After the system rolls back the failed task, the BMS will be removed from the list.

2.5 How Do I Create a BMS That Can be Quickly Provisioned?

When you create a common BMS (that is, a BMS booted from a local disk), its OS needs to be downloaded from the cloud and it also takes some time to install the OS. When you create a BMS that uses an EVS as its system disk, the OS has been installed on the disk and does not need to be downloaded or installed. In this way, the BMS can be provisioned within a short time when you apply for it.

On the page for creating a BMS, select a flavor that supports quick BMS provisioning, set the system disk type and capacity, and configure other required parameters.

Figure 2-1 Creating a BMS

The screenshot displays the BMS creation configuration page. It is divided into three main sections: Flavor, Image, and Disk.

Flavor Section: A table lists various flavors. The 'physicalS4.xlarge' flavor is selected and highlighted with a red box. Below the table, the current specifications are shown as 'physicalS4.xlarge | 28 cores | 12*16 GB'. A descriptive paragraph explains that this flavor provides a 10GE dual-port NIC and 2*10GE extension NICs for high-speed interconnections.

Flavor	Flavor name	CPU	Memory	Local Disk	Extended Configuration
<input type="checkbox"/>	physical.ks1.2xlarge	128 cores Hisilicon Kunp...	16*16 GB DDR4	960G SAS SSD	2 x 2*25GE
<input type="checkbox"/>	physicalS4.3xlarge	44 cores Intel(R) Xeon(...	12*32 GB DDR4	NA	2 x 2*10GE
<input type="checkbox"/>	physicalS4.large	20 cores 2*10Core 4114...	12*16 GB DDR4	NA	2x2*10GE
<input type="checkbox"/>	physicalS4.medium	20 cores Intel(R) Xeon(...	8*16 GB DDR4	NA	2 x 2*10GE
<input checked="" type="checkbox"/>	physicalS4.xlarge	28 cores Intel(R) Xeon(...	12*16 GB DDR4	NA	2 x 2*10GE

Current specifications: **physicalS4.xlarge | 28 cores | 12*16 GB**

This flavor provides a 10GE dual-port NIC for connecting to the VPC and 2*10GE extension NICs for high-speed interconnections between BMSs. You can configure VLANs on the extension NICs as needed. [Learn how to configure user-defined VLANs.](#)
This flavor supports boot from EVS disks, and provides EVS disk HA and quick service restoration. [Learn more about advanced features.](#)

Image Section: Three tabs are visible: 'Public Image', 'Private Image', and 'Shared Image'. The 'Public Image' tab is active. Two dropdown menus are shown, both highlighted with a red box. The first dropdown is set to 'CentOS' and the second to 'CentOS 7.6 x86 sdi2 for BareMetal'.

Disk Section: The 'EVS' tab is selected. A red box highlights the 'System Disk' configuration area. It shows a dropdown menu set to 'General Purpose SSD', a size of '150 GB', and IOPS limits: 'IOPS limit: 3,600, IOPS burst limit: 8,000'. Below this, there is a link to 'Add Data Disk' with the text 'You can attach 39 more disks.'

2.6 What Are the Advanced Features of BMSs Using EVS Disks?

BMSs that use EVS disks as their system disks can be provisioned within a shorter time, which facilitates quick service recovery.

Such BMSs have the following advanced features:

- BMSs booted from EVS disks can be provisioned within about 5 minutes.
- CSBS backups ensure data security.
- BMS rebuilding upon faults is supported, enabling quick service recovery.
- An image of a BMS can be exported to apply configurations of the BMS to other BMSs, eliminating the need to repeatedly configure BMSs.

Helpful Links

[Creating a BMS Supporting Quick Provisioning](#)

2.7 Can I Transfer a BMS to Another Account?

A BMS cannot be directly transferred to another account. To transfer it, perform the following operations:

1. Use the BMS to create a private image.
For details, see [Creating a Private Image from a BMS](#).
2. Share the image with the target account.
For details, see [Sharing Specified Images](#).
3. Use the target account to create a BMS from the shared private image.
4. Unsubscribe from the original BMS.

2.8 Is the BMS Host Name with Suffix `novalocal` Normal?

Symptom

Host names of some BMSs have suffix `.novalocal`.

For example, you set the host name to `abc` during BMS creation. [Table 2-2](#) lists the host names (obtained by running the `hostname` command) of BMSs created using different images and those displayed after the BMSs are restarted.

Table 2-2 Hostnames of BMSs created from different images

Image	Host Name Before BMS Restart	Host Name After BMS Restart
CentOS 6.8	abc	abc.novalocal
CentOS 7.3	abc.novalocal	abc.novalocal
Ubuntu 16	abc	abc

Host names of BMSs created from some types of images have suffix `.novalocal`, whereas others do not.

Troubleshooting

This is a normal phenomenon. You can ignore it.

The static host name of a Linux BMS is user-defined and injected using Cloud-Init during the BMS creation. According to the test results, Cloud-Init adapts to OSs differently. As a result, hostnames of some BMSs have suffix **.novalocal**, whereas others do not.

If you really do not want any host names with the suffix **.novalocal**, you can change the hostname. For details, see [How Do I Configure the Static Host Name of a BMS?](#)

2.9 Why Does the System Display a Message Indicating That the BMS Cannot Be Deleted?

Currently, BMSs only support the yearly/monthly billing mode and cannot be deleted. Therefore, when you click **Delete**, the system displays a message indicating that the BMS cannot be deleted.

If you no longer need a BMS, you can unsubscribe from it using either of the following methods:

- Locate the row that contains the BMS, click **More** in the **Operation** column, and select **Unsubscribe** from the drop-down list. On the **Unsubscribe** page, select a reason for the unsubscription and click **Confirm**.
- Choose **Billing Center > Unsubscriptions and Changes > Unsubscriptions**. Locate the row that contains the BMS, click **More** in the **Operation** column, and select **Unsubscribe** from the drop-down list.

2.10 How Do I Monitor BMS Metrics?

Cloud Eye can be used to monitor BMS metrics only after Agent is installed on the BMS. For details about the installation guide and supported monitoring metrics, see [Server Monitoring](#).

2.11 How Do I Create an Agency for Server Monitoring of the BMS?

1. On the management console, choose **Service List > Identity and Access Management**.
2. In the navigation pane on the left, choose **Agencies** and then click **Create Agency** in the upper right corner.
 - **Agency Name**: Enter **bms_monitor_agency**.
 - **Agency Type**: Select **Cloud service**.
 - **Cloud Service**: Select **Elastic Cloud Server (ECS) and Bare Metal Server (BMS)** from the drop-down list.
 - **Validity Period**: Select **Unlimited**.

- **Description:** Enter **Support BMS server monitoring**.
- 3. Click **Next**. On the **Select Policy/Role** page, search for and select **CES Administrator**.
- 4. Click **Next**. On the **Select Scope** page, select **All resources** or **Region-specific projects**.

 **NOTE**

If the BMS belongs to a sub-project, ensure that the sub-project has the CES Administrator permission.

- 5. Click **OK**.

3 Billing FAQ

3.1 Where Can I Query the BMS Price?

The prices of BMSs vary depending on the region and flavor. You can query the prices from [Product Pricing Details](#).

3.2 What Are Pre-payment and Post-payment? How Do I Choose Between Them?

Pre-payment refers to the yearly/monthly billing mode. The system calculates the price based on the BMS type and quantity you select and deducts the sum from your account.

Post-payment refers to the pay-per-use billing mode, which means that you pay for what you use. BMSs do not support this mode.

If your services are stable and you need BMSs for a long term, pre-payment is recommended. If your service demands fluctuate greatly and are hard to predict accurately, or temporary or abrupt demand spikes may occur, post-payment may better suit your needs.

3.3 What Is the Impact on Billing if I Unsubscribe from a BMS?

It depends on the BMS status and how long you have used it. [Table 3-1](#) provides the details.

Table 3-1 Unsubscription rules

Unsubscription Type	Number of Unsubscriptions	Handling Fee Charged	Amount Due Charged	Cash Coupons Refunded	Discount Coupons Returned
Unsubscribing from an in-use BMS	Unlimited	Yes	Yes	No	No
Unsubscribing from an inactive BMS	Unlimited	No	No	Yes	No
Unsubscribing from a BMS that failed to be created	Unlimited	No	No	Yes	No
Unsubscribing from an inactive renewal period	Unlimited	No	No	Yes	No


For more information, see [Unsubscription Rules](#).

3.4 How Do I Set Automatic Renewal for BMSs Billed on a Yearly/Monthly Basis?

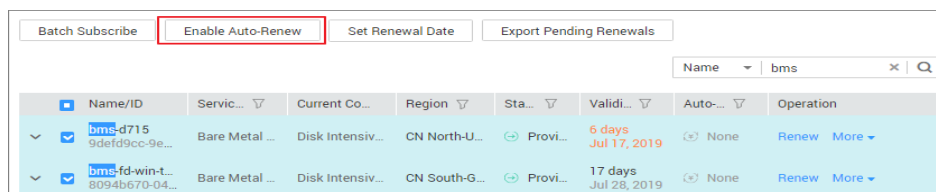
To prevent BMSs from being deleted when they expire, you can enable automatic renewal for BMSs that you want to use for a long term.

Procedure

1. Log in to the management console and choose **Billing > Renewal**.
2. Query the target BMSs by name, order No., or ID.
3. To enable automatic renewal for a BMS or multiple BMSs, do as follows:
 - To enable automatic renewal for a BMS, locate the row that contains the BMS and click **Enable Auto-Renew**. On the displayed page, set **New Auto-Renew Period** and click **OK**.



- To enable automatic renewal for multiple BMSs, select the BMSs and click **Enable Auto-Renew** on top of the **Renewals** list. On the displayed page, set **New Auto-Renew Period** and click **OK**.



The screenshot shows a management interface for Bare Metal Servers. At the top, there are four buttons: 'Batch Subscribe', 'Enable Auto-Renew' (highlighted with a red border), 'Set Renewal Date', and 'Export Pending Renewals'. Below the buttons is a search bar with the text 'Name' and 'bms'. The main part of the interface is a table with the following columns: Name/ID, Service, Current Co., Region, Sta., Validity, Auto-Renew, and Operation. Two rows of server data are visible:

Name/ID	Service	Current Co.	Region	Sta.	Validity	Auto-Renew	Operation
bms-d715 9defd9cc-9e...	Bare Metal ...	Disk Intensiv...	CN North-U...	Provi...	6 days Jul 17, 2019	None	Renew More
bms-fd-win-t... 8094b670-04...	Bare Metal ...	Disk Intensiv...	CN South-G...	Provi...	17 days Jul 28, 2019	None	Renew More

3.5 When Will a BMS Be Released After It Expires?

If you do not renew a BMS before it expires, it will be retained for a period of time.

The retention period is 15 days. For details, see [What Is a Retention Period?](#)

3.6 How Do I Unfreeze a BMS?

If a BMS has expired and is not renewed, it will be frozen and automatically enter a retention period, during which you cannot access or use it. For example, you cannot download data from the BMS. To unfreeze the BMS, renew it.

4 Key Pair and Password FAQ

4.1 How Do I Change the OS Password of a BMS?

Change the Password for Logging In to a Linux BMS

CentOS 7.5 is used as an example.


1. [Remotely log in to the BMS](#).
2. Run the `passwd root` command, enter the new password, and retype the new password.

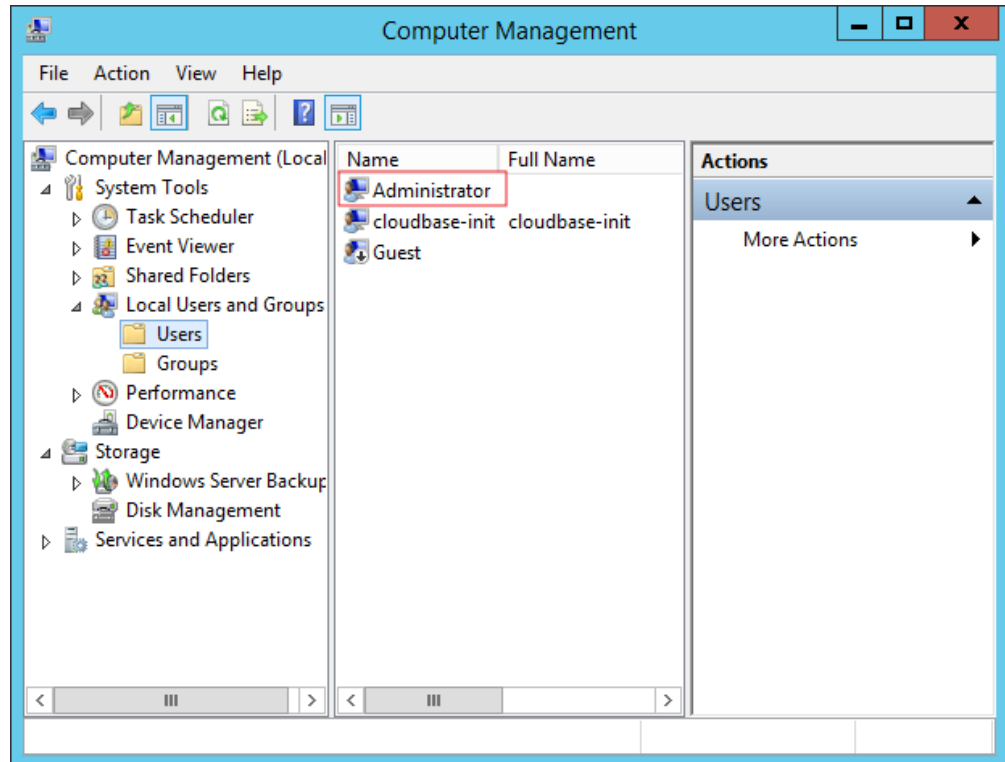
```
[root@ ~]-77a9 ~]# passwd root
Changing password for user root.
New password: _____
Retype new password: _____
passwd: all authentication tokens updated successfully.
[root@ ~]-77a9 ~]#
```

3. Enter `exit` to log out and use the new password to log in to the BMS to check whether you have changed the password successfully.


Change the Password for Logging In to a Windows BMS

Windows Server 2012 R2 is used as an example.


1. [Log in to the BMS remotely using MSTSC](#).
2. Click  in the lower left corner, choose **Windows PowerShell** and enter `compmgmt.msc` to open **Computer Management**.
3. Choose **System Tools > Local Users and Groups > Users**, right-click the target username, for example, **Administrator** in the following figure.




- 4. Select **Set Password**.



- 5. Click **Continue**.



6. Enter the new password, confirm the password, and click **OK**.



7. Press **Ctrl+Alt+Delete** to lock the BMS screen. Then unlock the screen and use the new password to log in.

5 Login FAQ

5.1 How Do I Prepare to Log In to a BMS?

Windows

- Obtain the login password.
When you create a Windows BMS, you can only select the key pair login. So, before you log in to a Windows BMS, you need to parse the key file into a password. For details, see [Obtaining the Password of a Windows BMS](#).
- Ensure that an EIP is bound to the BMS.
For details, see [Binding an EIP](#).

Linux

- Obtain the login password.
 - If the authentication mode is SSH key pair, use the key pair you selected when you created the BMS. If you have misplaced your private key file, enable [Data Encryption Workshop](#) to reset the key pair.
 - If the authentication mode is password, use the password you set when you created the BMS.
- Ensure that an EIP is bound to the BMS (not required for remote login).
If you want to log in to the BMS by using an SSH key pair, or a username and a password, you need to ensure that you have bound an EIP to the BMS.
For details, see [Binding an EIP](#).

5.2 What Can I Do If I Selected the Key Pair Authentication When I Created a BMS But Want to Log In to the BMS Using a Password?

When you create a Windows BMS, you can select only key pair authentication. To log in to the BMS using a password, you can:

- Use the private key file obtained when you created the BMS to obtain the password. For details, see [Obtaining the Password of a Windows BMS](#).

When you create a Linux BMS, you can select key pair or password authentication. If key pair authentication is selected but you want to use a password, you can:

- Log in to the BMS as instructed in [Logging In to a BMS Using an SSH Key Pair](#) and run the `passwd` command to set a password.

5.3 What Do I Do If I Can't Log In to a Linux BMS?

If you cannot log in to a BMS using SSH, it is recommended that you log in to the BMS through the console.

Check Whether You Can Remotely Log In to the BMS Through the Console

If the SSH login fails, check whether you can remotely log in to the BMS through the management console.

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.
3. Locate the row that contains the target BMS and click **Remote Login** in the **Operation** column.

After about one minute, the login page is displayed. Press **Enter** and enter username **root** and the password.

NOTE

Learn about the [preparations](#) for logging in to a BMS.

If you are still unable to log in to the BMS, record the related BMS information and the time when the login failure occurred. Then, [submit a service ticket](#). Contact Huawei Cloud technical support.

Troubleshoot

If you can log in to the BMS remotely but cannot log in to it using SSH, do as follows:

1. [Check Network Connectivity](#)
2. [Check Whether the Security Group Is Correctly Configured](#)
3. [Check Whether Non-System Disk Information Is Commented Out in the /etc/fstab File](#)
4. [Check the Port for Remotely Accessing the BMS](#)
5. [Check the CPU Load](#)

Check Network Connectivity

Check whether the EIP bound to the BMS can be pinged. If the EIP cannot be pinged, check whether the following rule has been added to the security group. If it is not, add it.

Protocol	Direction	Port Range	Source
ICMP	Inbound	All	0.0.0.0/0

Then, try to remotely log in to the BMS again.

Check Whether the Security Group Is Correctly Configured

Check whether port 22 of the security group is allowed. If it is not, add a security group rule to allow it.

- [Default Security Groups and Security Group Rules](#)
- [Adding Security Group Rules](#)

Then, try to remotely log in to the BMS again.

Check Whether Non-System Disk Information Is Commented Out in the /etc/fstab File

1. Log in to the BMS and run the following command to edit the `/etc/fstab` file:
`vi /etc/fstab`
2. Comment out the data disk configuration in the `/etc/fstab` file.

The `/etc/fstab` file contains information about the file systems and storage devices automatically attached to the BMS when the BMS starts. The configuration information about data disks automatically attached to the BMS needs to be commented out. For example, the last row shown in [Figure 5-1](#) is the data disk configuration to be commented out in the `/etc/fstab` file.

Figure 5-1 Data disk configuration in the fstab file

```
[root@ ~]# cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Wed Feb 27 06:58:16 2019
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=4c2c090d-4228-49fc-9cbe-3920b3bf287c / ext4 defaults 1 1
UUID=9c29104b-31b8-4421-a207-102f86ec7ae5 /mnt/test ext4 defaults 1 1
```

After performing the preceding operations, restart the BMS and try to remotely connect to it.

Check the Port for Remotely Accessing the BMS

Check the BMS settings.

1. Check whether the BMS sshd process is running.
2. Check whether the BMS rejects access by the local PC.
 - a. Log in to the BMS and run the following command:
`vi /etc/hosts.deny`

- b. If the IP address of the local PC is contained in the file, the IP address is rejected. In such a case, delete the IP address from the file.
3. Open the `/etc/ssh/ssh_config` file on the local PC and check the default login port. In addition, check whether the value of the `port` field (SSH-enabled port) in the `/etc/ssh/sshd_config` file of the BMS has been changed (the default port number is 22).

```
# semanage port -a -t ssh_port_t -p tcp #PORTNUMBER
#
#Port 22
#AddressFamily any
```

Then, try to remotely log in to the BMS again.

Check the CPU Load

Check whether the login failure is caused by high CPU usage. If yes, perform the following operations to reduce the CPU usage:

- Stop some processes that are not used temporarily and try again.
- You can also **restart the BMS**.
The restart operation will stop the server and interrupt services. Exercise caution when performing this operation.
- **Reinstall the OS**.
Reinstalling the OS is a high-risk operation. Back up data before the reinstallation.

Then, try to remotely log in to the BMS again.

5.4 What Browser Versions Can Be Used to Remotely Log In to a BMS?

When you use a browser to remotely log in to a BMS, ensure that the browser version meets the requirements listed in [Table 5-1](#).

Table 5-1 Browser version requirements

Browser	Version
Google Chrome	31.0-75.0
Mozilla FireFox	27.0-62.0
Internet Explorer	10.0-11.0

5.5 What Do I Do If the Login Page Does Not Respond?

Symptom

On the page for remotely logging in to a BMS, after you press **Enter**, the page does not respond.

Possible Causes

The BMS OS configuration does not allow remote login to the BMS.

Solution

Use a key pair to log in to the BMS and configure the OS as required. The configuration varies depending on the OS. The following part provides configurations of some OSs as examples. For details, see [Configuring Remote Login to a BMS](#) in *Bare Metal Server Private Image Creation Guide*.

1. Modify the configuration file.
 - For SUSE Linux Enterprise Server 12 SP2, SUSE Linux Enterprise Server 12 SP1, Ubuntu 16.04 Server, CentOS Linux 7.3, and EulerOS 2.2, use the vi editor to open the `/etc/default/grub` file and add `console=tty0 console=ttyS0` after `GRUB_CMDLINE_LINUX`.

Figure 5-2 Example

```
# If you change this file, run 'grub2-mkconfig -o /boot/grub2/grub.cfg' afterwards to update
# /boot/grub2/grub.cfg.
GRUB_DISTRIBUTOR=""
GRUB_DEFAULT=saved
GRUB_HIDDEN_TIMEOUT=0
GRUB_HIDDEN_TIMEOUT_QUIET=true
GRUB_TIMEOUT=8
GRUB_CMDLINE_LINUX_DEFAULT="resume=/dev/sda1 splash=silent quiet showopts crashkernel=99M,high crashkernel=72M,low"
# kernel command line options for failsafe mode
GRUB_CMDLINE_LINUX_RECOVERY=single
GRUB_CMDLINE_LINUX="console=tty0 console=ttyS0"
# Uncomment to enable BadRAM filtering, modify to suit your needs
# This works with Linux (no patch required) and with any kernel that obtains
# the memory map information from GRUB (GNU Mach, kernel of FreeBSD ...)
#GRUB_BADRAM=0x01234567,0xfefefefe,0x89abcdef,0xefefefef
# Uncomment to disable graphical terminal (grub-pc only)
GRUB_TERMINAL=gfxterm
# The resolution used on graphical terminal
# note that you can use only modes which your graphic card supports via VBE
# you can see them in real GRUB with the command 'vbeinfo'
GRUB_GFXMODE=auto
"grub" 40L, 2090C 15,46 Top
```

- For Oracle Linux 7.3 and Red Hat Enterprise Linux 7.3, use the vi editor to open the `/etc/sysconfig/grub` file and add `console=tty0 console=ttyS0` after `GRUB_CMDLINE_LINUX`.

Figure 5-3 Example

```
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="crashkernel=auto vconsole.font=latarcyrheb-sun16 rd.lvm.lv=ol/swap rd.lvm.lv=ol/root vconsole.keymap=us rhgb quiet "console=tty0 console=ttyS0"
GRUB_DISABLE_RECOVERY="true"
-
-
-
```

2. Update the configuration.
 - For SUSE Linux Enterprise Server 12 SP2, Oracle Linux 7.3, Red Hat Enterprise Linux 7.3, CentOS Linux 7.3, and EulerOS 2.2, run the following commands to update the configuration:
stty -F /dev/ttyS0 speed 115200
grub2-mkconfig -o /boot/grub2/grub.cfg
systemctl enable serial-getty@ttyS0
 - For Ubuntu 16.04 Server, run the following commands to update the configuration:
stty -F /dev/ttyS0 speed 115200
grub-mkconfig -o /boot/grub/grub.cfg
systemctl enable serial-getty@ttyS0
3. (Optional) Modify the security configuration file.
If you log in to the BMS through the serial port as user **root**, you need to modify the security configuration file. Add the following information to the end of **/etc/securetty**:

Figure 5-4 Example

```
vc/2
vc/3
vc/4
vc/5
vc/6
vc/7
vc/8
vc/9
vc/10
vc/11
tty1
tty2
tty3
tty4
tty5
tty6
tty7
tty8
tty9
tty10
tty11
ttyS0
"securetty" 39L, 221C
```

4. Run the **reboot** command to restart the OS.

After configuring the BMS OS, check whether you can log in to the BMS remotely.

5.6 What Do I Do If the BMS Console Is Displayed Improperly After I Remotely Log In to a BMS?

Symptom

The following symptoms occur:

- After you exit the vim editor, only half space of the screen is editable.
- When you enter more than 80 characters, the current row is covered.
- If you adjust the size of the browser window when using a text editor such as vim, rows are broken on the screen.

Possible Causes

Remote login to a BMS is subject to the communication on the serial port. The BMS console cannot automatically adapt to the screen. The default number of rows is 24, and that of columns is 80.

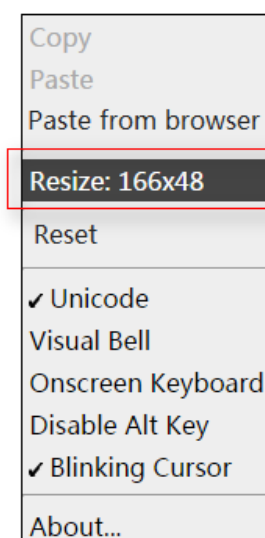
Solution

After you log in to the BMS remotely, right-click the blank area and select **Resize: xxx**. A command will be pasted on the command line, such as **stty cols 166 rows 48**. Then press **Enter** and adjust the console size.

Figure 5-5 Selecting Resize: xxx

```
Discovered PICMG Extension 2.2  
Discovered IPMB-0 address 0x20  
[SQL Session operational. Use ~? for help]
```

```
linux-8nad:~ # █
```



CAUTION

When you are using a text editor such as vim, do not adjust the window size. If you do need to adjust the window size, exit the editor first, adjust the window size, and adjust the console size based on the solution provided in this section.

5.7 What Do I Do If the Numeric Keypad Does Not Work During Remote Login?

Symptom

When I enter numbers using the numeric keypad for remote login, the numbers are not displayed properly.

```
Authorized users only. All activities may be monitored and reported.
Activate the web console with: systemctl enable --now cockpit.socket

bms-██████████-0001 login: ^[[4~^[[B^[[6~^[[D^[[G^[[C^C
```

Solution

Run the Linux **setleds** command to turn on the numeric keypad.

1. On the remote login page, run the following command to query the status of the numeric keypad:

```
setleds -F
```

```
[root@arm-autoinstaller ~]# setleds -F
Current flags:      NumLock off  CapsLock off  ScrollLock off
```

NumLock is **off**, indicating that the numeric keypad is turned off.

2. Run the following command to turn on the numeric keypad:

```
setleds +num
```

3. Run the **setleds -F** command again. If **NumLock** changes to **on**, the issue is fixed.

5.8 What Do I Do If the SSH Login or Data Transmission Is Slow?

Symptom

The login to Linux BMSs or data transmission between Linux BMSs in SSH mode is slow because UseDNS is enabled for SSH.

UseDNS is an enhanced security feature of SSH and is enabled by default. In such a case, the server obtains the host name of a client by locating the PTR record of the client IP address through a reverse DNS query. Then, the server performs a DNS query based on the obtained client host name and checks whether the

obtained IP address is the same as the original IP address, preventing client spoofing. Generally, clients use a dynamic IP address and do not have a corresponding PTR record. Therefore, this feature is invalid for information comparison. However, this feature increases the delay and thereby slows down the client connection.

Solution

1. Log in to the BMS remotely.
2. Run the following command to open the `/etc/ssh/sshd_config` file:
vi /etc/ssh/sshd_config
3. Find the following field:
`#UseDNS yes`
Add the following information in a new line after the field:
`UseDNS no`
4. Save the configuration and restart SSH.
service sshd restart

6 Network and Security FAQ

6.1 Can BMSs of Different Accounts Communicate with Each Other over an Internal Network?

Generally, BMSs of different accounts cannot communicate with each other for security concerns.

However, if you do need to allow BMSs of different accounts to communicate with each other through an internal network, you can create a VPC peering connection between VPCs in different accounts. For details, see [Creating a VPC Peering Connection with a VPC in Another Account](#).

6.2 How Do Two BMSs in the Same Region But in Different AZs Communicate with Each Other?

If they are in the same VPC, they communicate with each other through an internal network. If they are on the same subnet of a VPC, they communicate with each other through the layer-2 network. If they are on different subnets of a VPC, they communicate with each other through the layer-3 network. An EIP must be bound to the primary NIC of each BMS so that they can communicate with each other.

6.3 Are My BMSs in the Same Subnet?

You can customize your networks. Therefore, no matter whether your BMSs use the common network or high-speed network, you can control whether they are in the same subnet.

6.4 Can I Associate a BMS with Multiple Security Groups?

Yes. For details, see [Changing a Security Group](#).

6.5 Can BMSs Communicate with ECSs in the Same VPC?

Yes, BMSs can communicate with ECSs in the same VPC.

Your VPC may consist of multiple network segments. If the BMSs and ECSs are in the same segment, they communicate with each other through the Layer 2 network. If they are in different segments, they communicate with each other through the Layer 3 network.

In addition, you must configure security group rules for the BMSs to communicate with the ECSs. In addition, to enable an ECS to access a Windows BMS, disable the firewall of Windows.

6.6 What Are the Differences Between the Primary and Extension NICs of BMSs?

They are different in the following ways:

- Generally, the OS default routes preferentially use the primary NICs. If the OS default routes use the extension NICs, network communication will be interrupted. Then, you can check the route configuration to rectify the network communication error.
- By default, only the primary NICs can communicate with the public service zone (zone where PaaS and DNS services are deployed). The extension NICs cannot communicate this zone.

6.7 Can I Bind Multiple EIPs to a BMS?

Only one EIP can be bound to a NIC. If you want to bind multiple EIPs to a BMS, you can bind them to extension NICs and then perform required operations on the BMS, such as adding policy-based routes or namespaces, to ensure network connectivity. For details about how to configure policy-based routes, see [How Do I Access the Internet Using an EIP Bound to an Extension NIC?](#)

6.8 Can I Configure an EIP?

No. EIPs are automatically allocated from the DHCP address pool.

6.9 Will I Obtain an EIP That Has Been Released?

You may not. After an EIP is released, it is randomly assigned among users. If you just stop using an EIP temporarily and want to use it in the future, please do not release the EIP.

6.10 What Are the Differences Between EIPs, Private IP Addresses, and Virtual IP Addresses?

An EIP can be used to access the Internet and can only be bound to one BMS.

A private IP address is used for communication within an internal network and cannot be used to access the Internet.

Virtual IP addresses, also called floating IP addresses, are used for active and standby switchover of servers to achieve high availability. If the active server is faulty and cannot provide services, the virtual IP address is dynamically switched to the standby server to provide services.

6.11 How Can I Modify the Network Configuration or Restart the Network If I Can Log In to a BMS Using Only SSH?

The network automatically allocated by the BMS cannot be modified. If you modify the network configuration, you may fail to log in to the BMS. If the BMS has a NIC of the user-defined VLAN, you can configure or modify the network to which the NIC connects.

6.12 How Do I Handle the Failure to Ping a CentOS 7 Extension NIC?

Cause

A known kernel issue of the OS

NOTE

CentOS 7.4 and earlier CentOS 7 have this issue.

Solution

This issue has been rectified in CentOS 7.5. To use extension NICs, you are advised to use CentOS 7.5 or upgrade the OS kernel version to 3.10.0-862. (In quick provisioning scenarios, you can only resolve this issue by changing the OS to CentOS 7.5.).

1. Upload the CentOS 7.5 kernel file downloaded from the official website to the BMS and run the following command to update the kernel:

```
yum install kernel-3.10.0-862.el7.x86_64.rpm
```

NOTE

If you have configured automatic EVS disk attaching to the BMS in `/etc/fstab`, comment out the corresponding configuration item in `/etc/fstab`. Otherwise, you may fail to enter the BMS OS when restarting the BMS.

- Restart the OS. After entering the OS, reinstall the SDI iNIC driver, RAID controller card driver, and IB driver for CentOS 7.5 by following the instructions in the *Bare Metal Server Private Image Creation Guide*.

6.13 What Do I Do If the Communication Between the Primary NIC and Extension NIC of the BMS is Abnormal?

Cause

If two NICs on the same network segment are added to a BMS, communication between the primary NIC and extension NIC is abnormal because the BMS gateway strictly verifies the source MAC addresses. For example, in [Figure 6-1](#), the primary NIC and extension NIC are both on the 172.22.9.X network segment. A policy-based route needs to be configured to enable communication between the NICs.

Figure 6-1 Network segment of the NICs

```
10: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group default qlen 1000
    link/ether fa:16:3e:e5:b9:9d brd ff:ff:ff:ff:ff:ff
    inet 172.22.9.7/24 brd 172.22.9.255 scope global bond0
        valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fee5:b99d/64 scope link
        valid_lft forever preferred_lft forever
11: bond0.3935@bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group default qlen 1000
    link/ether fa:16:3e:54:2d:3b brd ff:ff:ff:ff:ff:ff
    inet 172.22.9.206/24 brd 172.22.9.255 scope global bond0.3935
        valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fe54:2d3b/64 scope link
        valid_lft forever preferred_lft forever
```

Solution

- Run the following command to add two routing table names (**net1** and **net2**) and priorities (**252** and **251**) to the **/etc/iproute2/rt_tables** file:

```
vi /etc/iproute2/rt_tables
```

```
252 net1
251 net2
```

- Run the following command to add the NIC routing information to the **/etc/rc.local** file:

```
vi /etc/rc.local
```

For example, the IP address of the primary NIC is 172.22.9.7, that of the extension NIC is 172.22.9.206, and that of the BMS gateway is 172.22.9.1, add the following routes:

```
ip route add 172.22.9.0/24 dev bond0 src 172.22.9.7 table net1
ip route add default via 172.22.9.1 dev bond0 table net1
ip route add 172.22.9.0/24 dev bond0.3935 src 172.22.9.206 table net2
ip route add default via 172.22.9.1 dev bond0.3935 table net2
ip rule add from 172.22.9.7/32 table net1
ip rule add from 172.22.9.206/32 table net2
```

6.14 How Can I Configure a Static IP Address for a BMS?

Scenarios

To customize the DNS server information of a BMS, you need to configure a static IP address for the BMS. If you change the IP address assignment mode from DHCP to the static mode, the IP address and gateway must be consistent with those when the BMS is provisioned. Otherwise, network disconnections may occur. This section takes CentOS 7 as an example to describe how to configure a static IP address for a BMS.

Procedure

1. Query the IP address and gateway of the BMS.

Run the following command to query the IP address of the BMS:

ifconfig bond0

```
[root@bms-2178 ~]# ifconfig bond0
bond0: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST> mtu 8888
    inet 192.168.20.238 netmask 255.255.255.0 broadcast 192.168.20.255
    inet6 fe80::f816:3eff:fe4b:c31c prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:4b:c3:1c txqueuelen 1000 (Ethernet)
    RX packets 7153 bytes 644462 (629.3 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 9435 bytes 1703746 (1.6 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

Run the following command to query the gateway address of the BMS:

ip ro

```
[root@bms-2178 ~]# ip ro
default via 192.168.20.1 dev bond0
169.254.0.0/16 dev bond0 scope link metric 1008
169.254.169.254 via 192.168.20.1 dev bond0 proto static
192.168.20.0/24 dev bond0 proto kernel scope link src 192.168.20.238
```

2. Modify the network configuration file.

Run the **vi /etc/sysconfig/network-scripts/ifcfg-bond0** command to open the **/etc/sysconfig/network-scripts/ifcfg-bond0** file, change the network information from DHCP to static, or delete **PERSISTENT_DHCLIENT=1** and add configuration items **IPADDR**, **NETMASK**, and **GATEWAY** (indicating the IP address, subnet mask, and gateway).

Figure 6-2 Modifying the network configuration file

```
USERCTL=no  
#PERSISTENT DHCPCLIENT=1  
BONDING_MASTER=yes  
ONBOOT=yes  
NM_CONTROLLED=no  
BOOTPROTO=static  
IPADDR=192.168.20.238  
NETMASK=255.255.255.0  
GATEWAY=192.168.20.1  
BONDING_OPTS="mode=4 xmit_hash_policy=layer3+4 miimon=100"  
DEVICE=bond0  
TYPE=Bond
```

NOTE

The IP address, subnet mask, and gateway must be consistent with those when the BMS is provisioned. Otherwise, network disconnections may occur.

3. Run the **systemctl disable bms-network-config.service** command to disable the bms-network-config network script.
4. Restart the BMS to make the network configuration take effect, or run the **kill dhclient** command to restart the network service to make the configuration take effect.

6.15 What Do I Do If the Network Performance Is Poor When 128 Ethernet Network Devices Are Bound to a BMS?

Symptom

The network performance is poor when multiple Ethernet network devices (for example, 128 single-queue NICs) are bound to an EulerOS 2.3/Euler OS 2.5 BMS. This is because the interrupts of all the 128 NIC queues are processed by the same CPU, causing a critical bottleneck of the CPU and affecting the network performance. The root cause is that **--hintpolicy=subset** is set for irqbalance of the Euler OS 2.3/Euler OS 2.5 BMS. To solve the problem, change it to **--hintpolicy=ignore** as instructed in [Procedure](#).

NOTE

hintpolicy is a policy used by irqbalance to balance interrupts across CPUs based on **affinity_hint** of each interrupt. **affinity_hint** indicates the CPU affinity of an interrupt. The value of **hintpolicy** can be:

- **exact**: irqbalance never violates **affinity_hint**.
- **subset**: irqbalance distributes interrupts to a subset of **affinity_hint**.
- **ignore**: irqbalance completely ignores **affinity_hint**.

Procedure

The following operations use EulerOS 2.3 as an example:

1. Log in to the BMS as user **root**.
2. Run the following command to open the `/etc/sysconfig/irqbalance` file:
vi /etc/sysconfig/irqbalance
3. The original setting is `--hintpolicy=subset`.

```
#  
# IRQBALANCE_ARGS  
#   append any args here to the irqbalance daemon as documented in the man page  
#  
IRQBALANCE_ARGS="--policyscript=/etc/sysconfig/irqbalance.rules --hintpolicy=subset"  
~
```

Change the setting to `--hintpolicy=ignore`.

```
#  
# IRQBALANCE_ARGS  
#   append any args here to the irqbalance daemon as documented in the man page  
#  
IRQBALANCE_ARGS="--policyscript=/etc/sysconfig/irqbalance.rules --hintpolicy=ignore"  
~
```

7 Disk FAQ

7.1 Can EVS Disks Be Attached to BMSs? How Many Data Disks Can Be Attached to a BMS?

Yes. Ultra-high I/O, general purpose SSD, high I/O, and common I/O (previous generation) EVS disks can be attached to BMSs. However, some BMSs support only local disks and EVS disks cannot be attached to them. In this case, you are not allowed to add an EVS disk when creating a BMS.

A maximum of 60 data disks can be attached to a BMS.

NOTE

If you need to attach an EVS disk to an existing BMS, **Device Type** of the EVS disk must be **SCSI**. If you need to create an EVS disk and attach it to the BMS, you must select **SCSI** in **Advanced Settings** when you create the EVS disk.

7.2 What Are the Restrictions for Attaching a Disk to a BMS?

- The disk and the target BMS must be located in the same AZ.
- The BMS must be in **Running** or **Stopped** state.
- **Device Type** of the EVS disk must be **SCSI**.
- A non-shared EVS disk must be in **Available** state.
A shared EVS disk must be in **In-use** or **Available** state.
- BMSs using some flavors or images cannot have EVS disks attached because the servers do not have SDI iNICs or for other reasons.

7.3 How Do I Know Whether EVS Disks Are Available in a Flavor?

Some BMSs do not support EVS disks because the BMSs do not have SDI iNICs or for other reasons. Therefore, certain flavors with EVS disks unavailable are provided for creating such BMSs.

You can call the [Querying Details About Flavors and Extended Flavor Information](#) API to check whether a BMS flavor supports EVS disks. If the value of `baremetal:_support_evs` in the response is `true`, EVS disks are available in the flavor. If the parameter value is `false` or the parameter is not displayed in the response, EVS disks are unavailable in the flavor.

7.4 How Do I Change the Disk Identifier in the `fstab` file to UUID?

Scenarios

After attaching disks to a Linux BMS, you must change the disk identifier in the `fstab` file to UUID. Otherwise, you cannot enter the BMS OS or the BMS becomes unavailable due to a mount point disorder after you stop and start the BMS, or restart the BMS.

NOTE

Universally Unique Identifier (UUID) is a 128-bit number used to identify information in computer systems.

Procedure

This section takes CentOS 7 as an example to describe how to change the disk identifier in the `fstab` file to UUID.

1. Log in to the BMS as user `root`. Run the `blkid` command to query all types of file systems that have been mounted to the BMS and UUIDs of the corresponding devices.

```
/dev/sda2: UUID="4eb40294-4c6f-4384-bbb6-b8795bbb1130" TYPE="xfs"  
/dev/sda1: UUID="2de37c6b-2648-43b4-a4f5-40162154e135" TYPE="swap"
```

2. Run the `cat /etc/fstab` command to open the `fstab` file.

```
/dev/sda2 / xfs defaults 0 0  
/dev/sda1 swap swap defaults 0 0
```

3. Check the disk identifier in the `fstab` file.

- If the disk identifier is UUID, no further action is required.
- If the disk identifier is the device name, go to [4](#).

4. Run the `vi /etc/fstab` command to open the `fstab` file, press `i` to enter editing mode, and change the disk identifier to UUID.

```
UUID=4eb40294-4c6f-4384-bbb6-b8795bbb1130 / xfs defaults 0 0  
UUID=2de37c6b-2648-43b4-a4f5-40162154e135 swap swap defaults 0 0
```

Press `Esc` and enter `:wq` to save and exit the file.

7.5 How Do I Obtain the Drive Letter of an EVS Disk?

After a BMS is restarted, the drive letter of an EVS disk attached to the BMS may change. This section describes how to find the mapping between an EVS disk and its drive letter.

1. Record **Device Identifier** of the EVS disk on the BMS details page.
2. Log in to the BMS OS, switch to the `/dev/disk/by-id` directory, and run the `ll` command to check the mapping between the WWN and drive letter. In Linux, WWN is in the format `wwn-0x + Device identifier`, for example, `wwn-0x50000397c80b685d -> ../../sdc`.

Figure 7-1 Checking the mapping between the WWN and drive letter

```
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c8088c61 -> ../../sdb
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c80b2539 -> ../../sde
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c80b685d -> ../../sdc
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c80ba3e9 -> ../../sdg
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c80bb905 -> ../../sdf
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c810e531 -> ../../sdd
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x600508e0000000002ab14603b88fa90b -> ../../sda
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e0000000002ab14603b88fa90b-part1 -> ../../sda1
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e0000000002ab14603b88fa90b-part2 -> ../../sda2
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e0000000002ab14603b88fa90b-part3 -> ../../sda3
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e0000000002ab14603b88fa90b-part4 -> ../../sda4
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e0000000002ab14603b88fa90b-part5 -> ../../sda5
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x68886030000369fafaf17a17502223655 -> ../../sdh
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x68886030000369fafaf17a17502223655-part1 -> ../../sdh1
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x68886030000369fafaf17a17502223655-part2 -> ../../sdh2
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x68886030000369fafaf17a17502223655-part3 -> ../../sdh3
rwXgwxgwx. 1 root root 9 Mar 21 14:16 wwn-0x6888603000036b61fa17a17502223655 -> ../../sdo
```

NOTE

You are advised to use the WWN to perform operations on disks. For example, run the `mount wwn-0x50000397c80b685d Folder name` command to attach a disk. You are not advised to use the drive letter directly because drive letter drift may cause the failure to find the disk.

Obtaining the drive letter of a disk by using the WWN is only supported by Linux.


7.6 Are the EVS Disk Device Names on the Console and the Device Names in BMS OSs Consistent?

Local System Disk

The EVS disk device names displayed on the BMS details page on the VPC console are inconsistent with the device names displayed in the BMS OS. To prevent impact of device name changes on services, you are advised to use EVS disks by UUID.

If EVS disks are specified during BMS allocation, the EVS disk device names displayed on the BMS details page start from `/dev/sdb` and the device names displayed in the BMS OS start after the BMS local disk names, as shown in [Figure 7-2](#).

Figure 7-2 Device names in the BMS OS



If EVS disks are attached to an allocated BMS, the device names displayed on the BMS details page are those specified by the tenant during disk attaching. After the EVS disks are detached from the BMS, the disks will not be displayed on the BMS details page, and the device names will be released.

If EVS disks are detached from an allocated BMS, the device names displayed in the BMS OS vary depending on whether the BMS OS restarts.

After EVS disks are attached to a BMS, if the BMS OS does not restart, the device names displayed in the BMS OS start from the smallest device name that is not used by other devices. For example, if device names `/dev/sda` and `/dev/sdc` are in use, the device names will start from `dev/sdb`. After EVS disks are detached from the BMSs, if the BMS OS does not restart, the BMS OS will release the device names.

If the BMS OS restarts, the device names displayed in the BMS OS will change based on the number of disks the BMS has and the disk attaching sequence. [Figure 7-3](#) shows the device names displayed in the BMS OS after EVS disks are attached to the BMS (before and after BMS restart). [Figure 7-4](#) shows the device names displayed in the BMS OS after EVS disks are detached from the BMS (before and after BMS restart).

Figure 7-3 Attaching EVS disks to a BMS

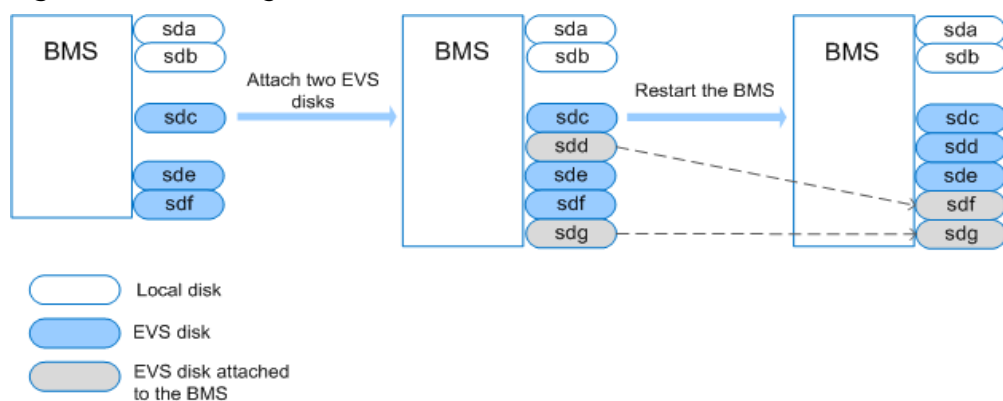



Figure 7-4 Detaching EVS disks from a BMS



EVS System Disk

The EVS disk device names displayed on the BMS details page on the VPC console may be inconsistent with the device names displayed in the BMS OS.

If EVS disks are specified during BMS allocation, the EVS disk device names displayed on the BMS details page start from **/dev/sda** and the device names in the BMS OS are displayed in a sequence determined by system scanning. There are two situations as shown in [Figure 7-5](#) and [Figure 7-6](#), and the EVS system disk always has the smallest drive letter of all the EVS disks.

Figure 7-5 Device names in the BMS OS (situation 1)


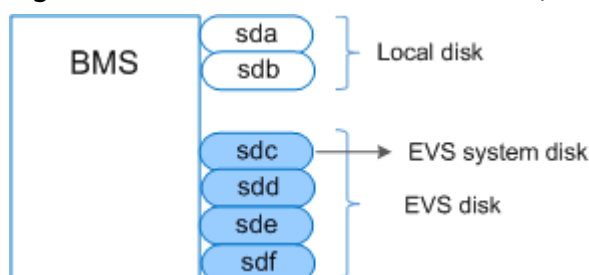


Figure 7-6 Device names in the BMS OS (situation 2)



If EVS disks are attached to an allocated BMS, the device names displayed on the BMS details page are those specified by the tenant during disk attaching. After the EVS disks are detached from the BMS, the disks will not be displayed on the BMS details page, and the device names will be released.

If EVS disks are detached from an allocated BMS, the device names displayed in the BMS OS vary depending on whether the BMS OS restarts.

After EVS disks are attached to a BMS, if the BMS OS does not restart, the device names displayed in the BMS OS start from the smallest device name that is not

used by other devices. For example, if device names `/dev/sda` and `/dev/sdc` are in use, the device names will start from `dev/sdb`. After EVS disks are detached from the BMSs, if the BMS OS does not restart, the BMS OS will release the device names.

If the BMS OS restarts, the device names displayed in the BMS OS will change based on the number of disks the BMS has and the disk attaching sequence. [Figure 7-7](#) and [Figure 7-8](#) show the device names displayed in the BMS OS after EVS disks are attached to the BMS (before and after BMS restart). [Figure 7-9](#) and [Figure 7-10](#) show the device names displayed in the BMS OS after EVS disks are detached from the BMS (before and after BMS restart).

Figure 7-7 Attaching an EVS disk (before the BMS restart)




Figure 7-8 Attaching an EVS disk (after the BMS restart)




Figure 7-9 Detaching an EVS disk (before the BMS restart)





Figure 7-10 Detaching an EVS disk (after the BMS restart)**NOTE**

To identify EVS disks, log in to the BMS and run the `ls SCSI` command. Disks whose type is **Huawei** are EVS disks.

7.7 Why Is the EVS Disk Size Not Updated in the BMS OS After the EVS Disk Capacity Has Been Expanded?

If this occurs, scan block devices in the BMS OS. Take the sdh disk of Red Hat as an example, run the `echo 1 > /sys/block/sdh/device/rescan` command.

7.8 How Can I Restore System Disk Data Using the Snapshot?

You can create snapshots of the BMS system disk on the EVS console periodically. To restore the system disk data, mount the target system disk to the `sda` mount point.

1. Power off the BMS.
 - a. Log in to the management console.
 - b. Under **Computing**, click **Bare Metal Server**.
The BMS console is displayed.
 - c. Locate the target BMS and click **Stop**.
2. Detach the system disk.
 - a. Click the BMS after it is powered off.
The page showing details of the BMS is displayed.
 - b. Locate the target system disk and click **Detach**.
In the displayed dialog box, click **OK**.
3. Attach the system disk.
 - a. On the page showing the BMS details, click **Attach Disk**.
The **Attach Disk** page is displayed.
 - b. Select the system disk and mount point `/dev/sda`, and click **Attach Disk**.
In the displayed dialog box, click **OK**.

7.9 What Do I Do to Prevent Risks of Attaching or Detaching the System Disk?

Attaching or detaching the system disk is a high-risk operation. You can attach or detach the system disk only when you need to restore the system disk data using the snapshot. In other cases, you are forbidden to attach or detach the system disk.

7.10 How Do I Select Storage?

When you create a BMS, you can select one from the following storage types:

- Elastic Volume Service (EVS): provides EVS disks of different QoS configurations to meet performance requirements in various scenarios.
- Dedicated Distributed Storage Service (DSS): provides exclusive storage resources. You can create disks of different specifications as needed and attach them to BMSs.

7.11 Why Is the Disk Capacity Displayed in the BMS OS Less Than That Displayed on the Official Website?

Possible causes of this issue are as follows:

1. Hardware vendors have a different method of calculating storage capacity from that of the OS. Hardware vendors use decimal notation to calculate disk capacity, in which 1 GB = 1000 x 1000 x 1000 bytes. In the OS, the capacity is calculated in binary mode, in which 1 GB = 1024 x 1024 x 1024 bytes.
2. The system contains hidden partitions, such as the boot partition, system backup, and restoration partition.
3. The file system consumes some disk capacity. Before using a hard disk, the OS partitions the disk and initializes the file system. The configuration also occupies a small amount of disk capacity.
4. The RAID array occupies some disk capacity. For example, if two 600 GB hard disks form RAID 1, only 600 GB capacity of one disk can be used.

8 OS FAQ

8.1 Can I Install or Upgrade BMS OSs by Myself?

You can reinstall a BMS OS. If the kernel is upgraded, confirm with the cloud service vendor whether certain services (for example, network service) and hardware drivers (such as RAID controller card drivers and NIC drivers) need to be upgraded accordingly. If the service packages or drivers are not upgraded to match the kernel version, the OS may fail to start or basic functions of the OS may be unavailable.

8.2 Can I Change the OS of a BMS?

Yes. You can change the OS of a BMS.

8.3 Is a GUI Provided for BMS OSs?

The Linux OSs provided for BMSs are managed using the command line interface (CLI). If you want to manage OSs using GUI, configure the GUI.

8.4 Is an Upload Tool Delivered with BMS OSs?

No. You must install and configure the upload tool, for example, the FTP tool, by yourself.

8.5 Does the Public Image Used to Create a BMS Have a Swap Partition?

Yes.

To increase the swap partition size, see [How Do I Increase the Swap Partition Size?](#)

8.6 How Do I Configure the Static Host Name of a BMS?



Symptom

The static host name of a Linux BMS is user-defined and injected on the console during the BMS creation. You can use the console or run the **hostname** command to change the host name of a BMS. However, if you restart the BMS, its host name will be automatically changed to the user-defined one injected on the console.

Automatically Updating the Host Name (Recommended)

Change the host name of the BMS on the console and enable automatic host name synchronization in the BMS OS. In this way, after the BMS is restarted, it automatically synchronizes the host name from the console.

This method has the following restrictions:

- The host name contains a maximum of 63 characters.
 - Special characters except hyphens (-), underscores (_), and periods (.) are not supported.
 - Uppercase letters are not supported.
 - This method does not apply to Windows BMSs.
1. Log in to the management console, click **Bare Metal Server** under **Computing**.
 2. Click the name of the BMS whose name is to be changed.
 3. On the displayed page, click  next to **Name**, enter a new name that meets the preceding requirements, and click  to save the change.
 4. Log in to the BMS OS and run the following command to enable automatic hostname synchronization:

```
sed -i 's/auto_synchronize_hostname.*/auto_synchronize_hostname = True/g' `find / -name bms-network-config.conf`
```

Check that automatic synchronization is enabled.

```
cat `find / -name bms-network-config.conf`
```

```
[NETWORK_CONFIG]
enable_bms_network = True
enable_bms_udev_rules = False
bsdtar_path=C:\Program Files\Cloudbase Solutions\Cloudbase-Init\bin\bsdtar.exe
mtu_use_dhcp_config = True
is_distributed_bms = False

[METADATA]
enable_preserve_hostname = False
auto_synchronize_hostname = True

[IB]
enable_ib = True

[ROCE]
enable_roce = True
```


5. Log in to the management console again. Locate the row that contains the BMS, click **More** in the **Operation** column, and select **Restart**.
After about 10 minutes, verify that the BMS is restarted and its hostname is automatically updated.

 **NOTE**

If you set the value of **auto_synchronize_hostname** in step 4 to **False**, the host name configured during BMS creation will be retained.

Manually Updating the Host Name

To make the changed host name take effect even after the BMS is stopped or restarted, save the changed name into configuration files.

For example, if the changed host name is *new_hostname*, perform the following steps:

1. Modify the **/etc/hostname** configuration file.
 - a. Run the following command to edit the **/etc/hostname** configuration file:
sudo vim /etc/hostname
 - b. Change the host name to *new_hostname*.
 - c. Run the following command to save and exit the configuration file:
:wq
2. (Optional) For Red Hat Enterprise Linux, CentOS, and Fedora 6, modify the configuration file **/etc/sysconfig/network**.
 - a. Run the following command to edit the **/etc/sysconfig/network** configuration file:
sudo vim /etc/sysconfig/network
 - b. Change the **HOSTNAME** value to *new_hostname*.
HOSTNAME=new_hostname
 - c. Run the following command to save and exit the configuration file:
:wq
3. Modify the **/etc/cloud/cloud.cfg** configuration file.
 - a. Run the following command to edit the **/etc/cloud/cloud.cfg** configuration file:
sudo vim /etc/cloud/cloud.cfg
 - b. Use either of the following methods to modify the configuration file:
 - Method 1: Change the **preserve_hostname** parameter value or add the **preserve_hostname** parameter to the configuration file.
If **preserve_hostname: false** is already available in the **/etc/cloud/cloud.cfg** configuration file, change it to **preserve_hostname: true**.
If **preserve_hostname: false** is unavailable in the **/etc/cloud/cloud.cfg** configuration file, add **preserve_hostname: true** before **cloud_init_modules**.
 - Method 2: Delete or comment out the following content:
update_hostname

- c. Run the following command to save and exit the configuration file:
:wq
4. Change the BMS network configuration script **bms-network-config.conf**.
The value of parameter **enable_preserve_hostname** in the **bms-network-config.conf** file is **False** by default, indicating that the host name is updated each time the board resets. To disable this function, change its value to **True**.
 - a. Change the value of **enable_preserve_hostname** in the **bms-network-config.conf** file to **True**:
sed -i 's/enable_preserve_hostname.*/enable_preserve_hostname = True/g' `find / -name bms-network-config.conf`
5. (Optional) For SUSE, modify the configuration file **/etc/sysconfig/network/dhcp**.
 - a. Run the following command to edit the **/etc/sysconfig/network/dhcp** configuration file:
sudo vim /etc/sysconfig/network/dhcp
 - b. Set the value of **DHCLIENT_SET_HOSTNAME** to **no** to ensure that DHCP does not automatically allocate host names.
DHCLIENT_SET_HOSTNAME="no"
 - c. Run the following command to save and exit the configuration file:
:wq
6. Run the following command to restart the BMS:
sudo reboot
7. Run the following command to check whether the static host name is changed:
sudo hostname
If the changed host name *new_hostname* is displayed in the command output, the host name is changed and the new name permanently takes effect.

8.7 How Do I Set the Password Validity Period?

If you cannot log in to a BMS due to password expiry, contact the operation administrator.

If you can log in to the BMS, perform the following operations to set the password validity period:

1. Log in to the BMS OS and run the following command to query the password validity period:
vi /etc/login.defs
The value of parameter **PASS_MAX_DAYS** indicates the password validity period.
2. Run the following command to change the value of parameter **PASS_MAX_DAYS** in **1**:
chage -M 99999 user_name
99999 is the validity period of the password, and *user_name* is a system user.

You are advised to set the password validity period as needed and change it on a regular basis.

3. Run **vi /etc/login.defs** to verify that the configuration has taken effect.

Figure 8-1 Configuration verification

```
# Password aging controls:
#
#     PASS_MAX_DAYS   Maximum number of days a password may be used.
#     PASS_MIN_DAYS   Minimum number of days allowed between password changes.
#     PASS_MIN_LEN    Minimum acceptable password length.
#     PASS_WARN_AGE   Number of days warning given before a password expires.
#
PASS_MAX_DAYS 99999
PASS_MIN_DAYS 0
PASS_MIN_LEN 5
PASS_WARN_AGE 7
```

8.8 How Do I Set SSH Configuration Items?

You can select the BMS login mode or account type. If you require special configuration, perform the following operations:

1. To improve security of the BMS, disable remote login using the password and retain only the certificate login mode. Configure the following parameters:
 - Check whether the **/etc/cloud/cloud.cfg** file contains parameter **ssh_pwauth** and its value is **false**. If not, add the parameter or change its value to **false**. This ensures that password cannot be used to log in to the BMS using Xshell.
 - Check whether the value of parameter **ChallengeResponseAuthentication** in the **/etc/ssh/sshd_config** file is **no**. If not, change it to **no**. This ensures that password cannot be entered using the keyboard inactive method to log in to the BMS using Xshell.
2. To enable remote login as user **root** and enable SSH permissions of user **root**, perform the following operations:

 **CAUTION**

This operation may cause risks. Exercise caution before performing this operation.

- a. Modify the Cloud-Init configuration file.

Take CentOS 6.7 as an example. Modify the following parameters:

```
users:
- name: root
  lock_passwd: false

disable_root: 0
ssh_pwauth: 1
```

In the preceding information:

- If the value of **lock_passwd** is set to **false**, user password is not locked.

- **disable_root** specifies whether to disable remote SSH login as user **root**. Set the value to **0**, indicating that the remote SSH login as user **root** is enabled (In some OSs, value **true** indicates disabled and **false** indicates enabled).
 - **ssh_pwauth** specifies whether to support SSH password login. Set this parameter to **1**, indicating that SSH password login is supported.
- b. Run the following command to open the **/etc/ssh/sshd_config** file using the vi editor:

```
vi /etc/ssh/sshd_config
```

Change the value of **PasswordAuthentication** in the **sshd_config** file to **yes**.

 NOTE

- For SUSE and openSUSE, set **PasswordAuthentication** and **ChallengeResponseAuthentication** in the **sshd_config** file to **yes**.
 - For Ubuntu, set **PermitRootLogin** to **yes**.
- c. Lock the initial password of user **root** in the image template by modifying the **shadow** file to prevent risks.

- i. Run the following command to open the **/etc/shadow** configuration file using the vim editor:

```
vim /etc/shadow
```

Add **!!** to the password hash value of the root account. The modified configuration file is as follows:

```
# cat /etc/shadow | grep root
root:!!$6$SphQRPXu$Nvg6izXbhDPrcY3j1vRiHaQFVRpNiV3HD/
bjDgnZrACOWPXwJahx78iaut1liglUrwavVGSYQ1JOlw.rDlVh7.:17376:0:99999:7::
```

- ii. After the configuration file is modified, press **Esc** and enter **:wq** to save and exit the file.

 NOTE

For Ubuntu, delete the user created during the OS installation. For example, run the **userdel -rf ubuntu** command to delete user **ubuntu** created during OS installation.

8.9 How Can I Handle the Eight-Hour Difference Between the Windows BMS and Local Time

Cause

Linux uses the time of the motherboard CMOS chip as the Coordinated Universal Time (UTC) and determines the system time based on the configured time zone. However, Windows uses the CMOS time as the system time directly without converting it based on the time zone.

Solution

1. Log in to the Windows BMS.



2. Click  in the lower left corner, choose **Windows PowerShell**, and enter **regedit.exe** to open the registry.
3. In the displayed **Registry Editor** window, choose **HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Control > TimeZoneInformation**.
4. In the right pane, right-click a blank area and choose **New > DWORD (32-bit) Value** to add a REG_DWORD code. Set its name to **RealTimeIsUniversal** and value to **1**.

Figure 8-2 Adding a code



5. After the modification, restart the BMS.
After the BMS restarts, its system time is consistent with the local time.

8.10 How Do I Change the SID of a Windows Server 2012 BMS?

Scenarios

A Security Identifier (SID) is a unique value that identifies a user, group, or computer account (administrator account). When an account is created for the first time, a unique SID is assigned to each account on the network. A SID is determined by the computer name, current time, and CPU use time of the current user-mode thread.

A complete SID contains:

- User and group security description
- 48-bit ID authority
- Revision level
- Variable sub-authority values

An example SID is S-1-5-21-287469276-4015456986-3235239863-500.

S	1	5	21-287469276-4015456986-3235239863	500
---	---	---	------------------------------------	-----

The string is a SID.	SID version	SID authority, which is NT in this example	SID sub-authorities	Accounts and groups in the domain
----------------------	-------------	--	---------------------	-----------------------------------

Currently, all the Windows Server 2012 BMSs have the same SID. In the cluster deployment scenario, you need to change the SID of the BMSs to ensure that each BMS uses a unique SID.

Procedure


1. Log in to the BMS OS.
2. Click  in the lower left corner, choose **Windows PowerShell**, and run the **whoami /user** command to query the SID.

Figure 8-3 Querying the original SID

```

User Name                               SID
-----
-00\administrator S-1-5-21-287469276-4015456986-3235239863-500
PS C:\Users\Administrator>
    
```

3. Modify the Cloudbase-Init configuration files.
 - a. Open the **cloudbase-init.conf** and **cloudbase-init-unattend.conf** files.
File directory: **C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf**
 - b. Add **first_logon_behaviour=no** to both files.


```

[DEFAULT]
username=Administrator
groups=Administrators
first_logon_behaviour=no
netbios_host_name_compatibility=false
metadata_services=cloudbaseinit.metadata.services.httpserver
inject_user_password=true
                    
```
 - c. Delete **cloudbaseinit.plugins.common.sethostname.SetHostNamePlugin** from the **cloudbase-init-unattend.conf** configuration file.

Figure 8-4 Modifying the configuration file

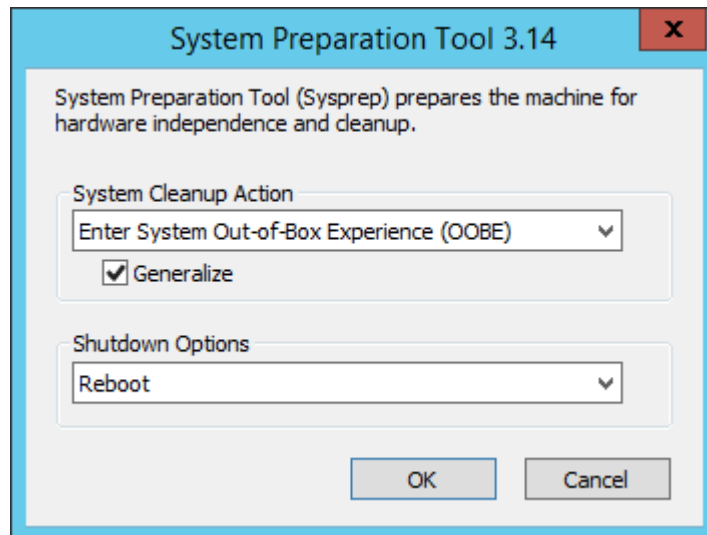
```

se Solutions\Cloudbase-Init\log\
.log
0, suds=INFO, iso8601=WARN, requests=WARN
M1, 115200, N, 8

iles\Cloudbase Solutions\Cloudbase-Init\LocalScripts\
.metadata.services.configdrive.ConfigDriveService, cloudbaseinit.metadata.services.httpserver
ommon.mtu.NTUPugin, cloudbaseinit.plugins.common.sethostname.SetHostNamePlugin, cloudbaseini
    
```

4. Open the CLI and run the following command to open the Sysprep window:
`C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf>
C:\Windows\System32\Sysprep\sysprep.exe /unattend:Unattend.xml`
5. In the **System Preparation Tool 3.14** dialog box, configure parameters and click **OK**.

Figure 8-5 System Preparation Tool settings



6. After the configuration is complete, the BMS automatically restarts. You need to encapsulate and decompress the package again. After the BMS restarts, you need to reset the password for the Windows OS. Contact the customer service.
7. Log in to the BMS OS and check the SID using the method in [2](#).

Figure 8-6 Querying the new SID

```

User Name                               SID
=====
win-ck7r022vrh0\administrator 5-1-5-21-3812874840-1741028955-636704118-500
PS C:\Users\Administrator>
    
```

As shown in the preceding figure, the SID has been changed successfully.

8.11 How Do I Change the Kernel Version of CentOS 7 BMSs?

Scenarios

Some special software runs properly only on specified Linux kernel versions. This section describes how to change the kernel version.

Solution

1. Log in to the BMS OS.

2. Run the following command to check the current kernel version:

uname -r

```
[root@xxxxxx~]# uname -r  
3.10.0-327.22.2.el7.x86_64
```

3. Run the following command to check the number of kernels in the OS:

cat /boot/grub2/grub.cfg | grep menuentry

```
[root@xxxxxx~]# cat /boot/grub2/grub.cfg | grep menuentry  
if [ x"${feature_menuentry_id}" = xy ]; then  
  menuentry_id_option="--id"  
  menuentry_id_option=""  
  export menuentry_id_option  
  menuentry 'CentOS Linux (3.10.0-327.22.2.el7.x86_64) 7 (Core)' --class centos --class gnu-linux --  
  class gnu --class os --unrestricted $menuentry_id_option 'gnulinux-3.10.0-327.el7.x86_64-  
  advanced-80b9b662-0a1d-4e84-b07b-c1bf19e72d97' {  
  menuentry 'CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core)' --class centos --class gnu-linux --class  
  gnu --class os --unrestricted $menuentry_id_option 'gnulinux-3.10.0-327.el7.x86_64-  
  advanced-80b9b662-0a1d-4e84-b07b-c1bf19e72d97' {  
  menuentry 'CentOS Linux (0-rescue-7d26c16f128042a684ea474c9e2c240f) 7 (Core)' --class centos --  
  class gnu-linux --class gnu --class os --unrestricted $menuentry_id_option 'gnulinux-0-  
  rescue-7d26c16f128042a684ea474c9e2c240f-advanced-80b9b662-0a1d-4e84-b07b-c1bf19e72d97' {
```

4. Run the following command to set the kernel that is started by default, for example, CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core):

grub2-set-default "CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core)"

5. Run the following command to check the kernel that is started by default:

grub2-editenv list

```
[root@xxxxxx~]# grub2-editenv list  
saved_entry=CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core)
```

6. Restart the BMS, enter its OS, and run the command in step 2 to check whether the kernel has been changed.

```
[root@xxxxxx~]# uname -r  
3.10.0-327.el7.x86_64
```

8.12 How Do I Reserve Log Space If the Root Partition Automatically Expands Disks?

Scenarios

In the scenario where the root partition automatically expands disks, the initial root partition may occupy all space of the system disk. This section describes how to reserve log space.

Procedure

1. Run the **lsblk** command. The following command output indicates that the initial root partition has occupied all space of the system disk.


```
Last login: Fri Mar  2 01:26:34 2018
root@bms-ubuntu-0001:~# lsblk
NAME        MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda          8:0    0 837.3G  0 disk
├─sda1       8:1    0   953M  0 part /boot
├─sda2       8:2    0    4.7G  0 part [SWAP]
├─sda3       8:3    0 831.6G  0 part /
└─sda4       8:4    0    64M   0 part
root@bms-ubuntu-0001:~#
```

2. Run the following command to create a directory for storing logs:

mkdir log

```
root@bms-ubuntu-0001:~# mkdir log
root@bms-ubuntu-0001:~# ll
total 44
drwx----- 6 root root 4096 May 31 08:48 ./
drwxr-xr-x 24 root root 4096 May 31 08:47 ../
-rw----- 1 root root    1 Mar  2 01:35 .bash_history
-rw-r--r-- 1 root root 3106 Feb 19  2014 .bashrc
drwx----- 2 root root 4096 Dec 22 23:49 .cache/
drwxr-xr-x  2 root root 4096 May 31 08:48 log/
drwxr-xr-x  2 root root 4096 Feb 28 01:41 .oracle_jre_usage/
-rw-r--r-- 1 root root  140 Feb 19  2014 .profile
drwx----- 2 root root 4096 Dec 11 22:21 .ssh/
-rw----- 1 root root 4835 Mar  2 01:35 .viminfo
```

3. Run the following command to create a 200 GB image file for storing logs.

dd if=/dev/zero of=disk.img bs=1M count=200000

```
root@bms-ubuntu-0001:~# dd if=/dev/zero of=disk.img bs=1M count=200000
200000+0 records in
200000+0 records out
209715200000 bytes (210 GB) copied, 807.411 s, 260 MB/s
root@bms-ubuntu-0001:~# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda3       819G  198G  588G  26% /
none            4.0K   0  4.0K  0% /sys/fs/cgroup
udev            158G  12K  158G  1% /dev
tmpfs           32G  1.1M   32G  1% /run
none            5.0M   0  5.0M  0% /run/lock
none            158G   0  158G  0% /run/shm
none            100M   0  100M  0% /run/user
/dev/sda1       922M   54M  806M  7% /boot
root@bms-ubuntu-0001:~#
```

4. Run the following commands to virtualize the generated file into a block device and format it:

losetup /dev/loop0 disk.img

mkfs.ext4 /dev/loop0

```
root@bms-ubuntu-0001:~# losetup /dev/loop0 disk.img
root@bms-ubuntu-0001:~# mkfs.ext4 /dev/loop0
mke2fs 1.42.9 (4-Feb-2014)
Discarding device blocks: done
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
12804096 inodes, 51200000 blocks
2560000 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
1563 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

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

5. Run the following command to mount the image file to the log directory:

mount disk.img log

```
root@bms-ubuntu-0001:~# mount disk.img log/
root@bms-ubuntu-0001:~# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda3	819G	1.7G	784G	1%	/
none	4.0K	0	4.0K	0%	/sys/fs/cgroup
udev	158G	12K	158G	1%	/dev
tmpfs	32G	1.1M	32G	1%	/run
none	5.0M	0	5.0M	0%	/run/lock
none	158G	0	158G	0%	/run/shm
none	100M	0	100M	0%	/run/user
/dev/sda1	922M	54M	806M	7%	/boot
/dev/loop1	193G	60M	183G	1%	/root/log

6. Create a file in the log directory.

```
root@bms-ubuntu-0001:~# cd log/
root@bms-ubuntu-0001:~/log# ll
total 24
drwxr-xr-x 3 root root 4096 May 31 09:09 ./
drwx----- 6 root root 4096 May 31 08:50 ../
drwx----- 2 root root 16384 May 31 09:09 lost+found/
root@bms-ubuntu-0001:~/log# vim test
root@bms-ubuntu-0001:~/log# cat test
helloworld!
```

7. Run the following command to add the mount command to `/etc/rc.local`:
mount /root/disk.img /root/log

```
#
# By default this script does nothing.
mount /root/disk.img /root/log
exit 0
~
```

8. Run the following command to restart the OS:
reboot

```
The system is going down for reboot NOW!
Connection closing...Socket close.

Connection closed by foreign host.

Disconnected from remote host(10.185.78.41:22) at 21:20:32.
```

9. Run the **lsblk** command. The command output indicates that the image file has been mounted.

```
Last login: Thu May 31 08:51:44 2018 from 10.190.179.88
root@bms-ubuntu-0001:~# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda          8:0    0 837.3G  0 disk
├─sda1       8:1    0   953M  0 part /boot
├─sda2       8:2    0   4.7G  0 part [SWAP]
├─sda3       8:3    0 831.6G  0 part /
└─sda4       8:4    0    64M  0 part
loop0       7:0    0 195.3G  0 loop /root/log
root@bms-ubuntu-0001:~# cat /root/log/test
helloworld!
root@bms-ubuntu-0001:~#
```

8.13 How Do I Roll Back the Kernel Version If I Mistakenly Upgrade the Kernel?

Scenarios

SDI, RAID, and IB hardware drivers of the BMS are related to the kernel. You are not advised to upgrade the kernel version.

If you have upgraded the kernel, perform the operations in this section. This section uses CentOS 7.2 as an example to describe how to set the BMS OS to start from the default kernel if you have upgraded the kernel.

Upgrade Scenario

1. Run the **uname -a** command to query the current kernel version.

```
[root@bms-centos ~]# uname -a  
Linux bms-centos 3.10.0-327.el7.x86_64 #1 SMP Thu Nov 29 14:49:43 UTC 2018 x86_64 x86_64  
x86_64 GNU/Linux
```
2. Run the **yum update kernel** command to upgrade the kernel.
3. Run the **cat /boot/grub2/grub.cfg |grep menuentry** command to check the kernel information of the OS after the upgrade.

As shown in the following figure, **3.10.0-327.el7.x86_64** is the default kernel version and **3.10.0-862.3.2.el7.x86_64** is the upgraded kernel version.

```
if [ x"${feature_menuentry_id}" = xy ]; then  
  menuentry_id_option="--id"  
  menuentry_id_option=""  
export menuentry_id_option  
menuentry 'CentOS Linux (3.10.0-862.3.2.el7.x86_64) 7 (Core)' --class centos --class gnu-linux --class  
ctd $menuentry_id_option 'gnulinux-3.10.0-327.el7.x86_64-advanced-4c147502-c776-4ca9-8657-fb4c8e8c9794'  
menuentry 'CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core)' --class centos --class gnu-linux --class gnu  
$menuentry_id_option 'gnulinux-3.10.0-327.el7.x86_64-advanced-4c147502-c776-4ca9-8657-fb4c8e8c9794' {  
menuentry 'CentOS Linux (0-rescue-2b86009638bb45c9ad2f4e3d14ba820a) 7 (Core)' --class centos --class gnu  
ss os --unrestricted $menuentry_id_option 'gnulinux-0-rescue-2b86009638bb45c9ad2f4e3d14ba820a-advanced-  
b4c8e8c9794' {
```

Emergency Settings After Kernel Upgrade

1. Run the following commands to set the original kernel version as the default startup kernel and verify the modification result:

```
grub2-set-default "CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core)"
```

```
grub2-editenv list
```

```
[root@bms-centos ~]# grub2-editenv list  
saved_entry=CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core)
```

2. After the verification is complete, restart the OS from the default kernel.

```
CentOS Linux (3.10.0-862.3.2.el7.x86_64) 7 (Core)  
CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core)  
CentOS Linux (0-rescue-2b86009638bb45c9ad2f4e3d14ba820a) 7 (Core)
```

3. Run the **uname -a** command to check whether the kernel version is restored.

8.14 How Do I Increase the Swap Partition Size?

Scenarios

When you install the Oracle database for a Linux OS, the swap partition size will be checked. If the swap partition cannot meet requirements, you can perform the operations in this section to increase the swap partition size.

NOTE

The swap partition is similar to the virtual memory of the Windows OS. When the memory is insufficient, some hard disk space is virtualized into memory to improve the system running efficiency.

Procedure

1. Log in to the BMS OS.
2. Run the **lsblk** command to check the size of the swap partition.

```
[root@bms- ~]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda         8:0    0  1.1T  0 disk
├─sda1      8:1    0   500M  0 part /boot
├─sda2      8:2    0  29.5G  0 part
├─rhel-root 253:0   0  26.5G  0 lvm  /
├─rhel-swap 253:1   0    3G  0 lvm  [SWAP]
└─sda3      8:3    0    64M  0 part
```

The size of the swap partition is 3 GB.

3. Run the following command to increase the swap partition size by 5 GB (example):

```
dd if=/dev/zero of=/swapfile bs=1M count=5000
```

```
chmod 600 /swapfile
```

```
mkswap /swapfile swapon /swapfile echo "/swapfile swap swap defaults 0 0" >>/etc/fstab
```

4. Run the **lsblk** command to check the size of the expanded swap partition.

```
[root@bms- ~]# free
              total        used         free   shared  buff/cache   available
Mem:           263564592    87360740    18486896    805268    157716956    174200612
Swap:           8265716      2362592     5903124
```

The size of the swap partition is 8 GB.

8.15 How Do I Increase the Size of the Root Partition of a BMS Which Is Quickly Provisioned?

Scenarios

If the root partition size of a BMS which is quickly provisioned cannot meet service requirements, you can increase the root partition size by performing the operations in this section.

Procedure

Assume that the drive letter of the BMS system disk is `/dev/sdf` and its initial size is 40 GB. To increase the partition size to 140 GB, perform the following operations:

```
sdf      8:80    0    40G    0 disk
├─sdf1   8:81    0    500M   0 part /boot
├─sdf2   8:82    0     5G    0 part [SWAP]
├─sdf3   8:83    0   34.5G  0 part /
└─sdf4   8:84    0    64M   0 part
```

1. On the EVS console, add 100 GB capacity to the disk.

Current Configuration			
Disk Name	bms-6acd	Disk ID	64307021-1178-4e68-9dc6-b599f403176b
Current Capacity (GB)	40	AZ	eu-de-01
Region	ComputeService		
Disk Type	Common I/O		

Add Capacity (GB)
New Capacity (GB) **140**

After the capacity expansion, check that the size of system disk `/dev/sdf` has been increased from 40 GB to 140 GB. The 64 MB `sdf4` partition is the configdriver partition that stores the BMS configuration information.

```
sdf      8:80    0   140G    0 disk
├─sdf1   8:81    0    500M   0 part /boot
├─sdf2   8:82    0     5G    0 part [SWAP]
├─sdf3   8:83    0   34.5G  0 part /
└─sdf4   8:84    0    64M   0 part
```

2. Run the following command to back up the content in the configdriver partition:

```
dd if=/dev/sdf4 of=/root/configdriver.img
```

```
[root@bms-6acd ~]# dd if=/dev/sdf4 of=/root/configdriver.img
131072+0 records in
131072+0 records out
67108864 bytes (67 MB) copied, 0.291739 s, 230 MB/s
```


- Run the **fdisk /dev/sdf** command and perform the operations shown in the following figure to delete the configdriver partition.

```

[root@bms-6acd ~]# fdisk /dev/sdf
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): d
Partition number (1-4, default 4): 4
Partition 4 is deleted

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource
busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
    
```

Run the **partprobe** command to refresh the partitions. The system disk contains the following partitions:

```

sdf      8:80    0    140G    0 disk
├─sdf1   8:81    0    500M    0 part  /boot
├─sdf2   8:82    0     5G     0 part  [SWAP]
├─sdf3   8:83    0   34.5G   0 part  /
└─sdf4   8:84    0    97.7M   0 part
    
```

- Create a 100 MB configdriver partition in the system disk. Assume that the available sector range is 83755008-293601279, the initial value of the new partition is equal to the maximum value of the available sector minus 200000, that is, 293401279. The maximum value of the new partition is 293601279.

```

Command (m for help): n
Partition type:
  p   primary (3 primary, 0 extended, 1 free)
  e   extended
Select (default e): p
Selected partition 4
First sector (83755008-293601279, default 83755008): 293401279
Last sector, +sectors or +size{K,M,G} (293401279-293601279, default 293601279):
Using default value 293601279
Partition 4 of type Linux and of size 97.7 MiB is set

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource
busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
    
```

Run the **partprobe** command to refresh the partition details.

```

sdf      8:80    0    140G    0 disk
├─sdf1   8:81    0    500M    0 part  /boot
├─sdf2   8:82    0     5G     0 part  [SWAP]
├─sdf3   8:83    0   34.5G   0 part  /
└─sdf4   8:84    0    97.7M   0 part
    
```

- Run the following command to expand the root partition:
growpart /dev/sdf 3

```
root@bms-6acd ~]# growpart /dev/sdf 3
CHANGED: partition=3 start=11511808 old: size=72243200 end=83755008 new: size=281889471,end=293401279
```

The size of the extended root partition is as follows:

```
sdf      8:80  0  140G  0 disk
├─sdf1   8:81  0   500M  0 part /boot
├─sdf2   8:82  0    5G   0 part [SWAP]
├─sdf3   8:83  0 134.4G  0 part /
└─sdf4   8:84  0  97.7M  0 part
```

Run the **resize2fs /dev/sdf3** command to expand the file system of the root partition.

```
root@bms-6acd ~]# resize2fs /dev/sdf3
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/sdf3 is mounted on /; on-line resizing required
old_desc_blocks = 5, new_desc_blocks = 17
The filesystem on /dev/sdf3 is now 35236183 blocks long.

You have new mail in /var/spool/mail/root
root@bms-6acd ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sdf3       133G  1.9G 125G   2% /
devtmpfs        63G   0    63G   0% /dev
tmpfs           63G   0    63G   0% /dev/shm
tmpfs           63G  9.0M  63G   1% /run
tmpfs           63G   0    63G   0% /sys/fs/cgroup
/dev/sdf1       477M 104M  344M  24% /boot
tmpfs          13G   0    13G   0% /run/user/0
```

6. Run the following command to restore the configdriver partition:

```
dd if=/root/configdriver.img of=/dev/sdf4
```

```
root@bms-6acd ~]# dd if=/root/configdriver.img of=/dev/sdf4
131072+0 records in
131072+0 records out
67108864 bytes (67 MB) copied, 0.372614 s, 180 MB/s
root@bms-6acd ~]#

/dev/sdf1: UUID="b9c472f9-6737-4200-910a-efa3af16629a" TYPE="ext4"
/dev/sdf2: UUID="b07ff4d0-8b0b-4c43-a40a-0b27290ea215" TYPE="swap"
/dev/sdf3: UUID="1e57f71e-6adc-4e98-9407-0f7d678d4525" TYPE="ext4"
/dev/sdf4: UUID="2018-09-27-19-13-01-00" LABEL="config-2" TYPE="iso9660"
root@bms-6acd ~]#
```

The capacity expansion of the BMS root partition is complete.

8.16 Common Linux Commands

lsblk

The **lsblk** command is used to list all available block devices and the dependencies between them, except the RAM disks. Block devices include hard disks, flash memory, and CD-ROM.

By default, the **lsblk** command lists all block devices in a tree structure. Start the terminal and run the following command:

```
lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 202:0 0 40G 0 disk
├─sda1 202:1 0 4G 0 part [SWAP]
```



```
└─sda2 202:2 0 36G 0 part /  
sdb 202:16 0 10G 0 disk
```

The parameters are as follows:

- **NAME:** block device name
- **MAJ:MIN:** primary and secondary device numbers
- **RM:** whether the device is removable. **0** indicates no and **1** indicates yes.
- **SIZE:** device capacity
- **RO:** whether the device is read-only. **0** indicates no and **1** indicates yes.
- **TYPE:** block device type (disk or a disk partition)
- **MOUNTPOINT:** mount point of a device

8.17 How Do I Update the Disk Metadata After the LVM Volume Is Remounted?

Scenarios

If the LVM volume is remounted when a BMS OS is reinstalled, you need to update the disk metadata in a timely manner. Otherwise, the OS will be unavailable after it is restarted.

Procedure

If a BMS uses LVM partitioning and the LVM volume is remounted when the BMS OS is reinstalled, update the disk metadata in a timely manner after the remount is complete. In this way, the disk metadata will be consistent with the disk mounting information after the OS is restarted. To update disk metadata, run the following commands (*sysvg* is the volume group (VG) name of the LVM volume):

```
lvm vgcfgrestore sysvg
```

```
lvm pvscan
```

```
lvm vgscan
```

```
lvm vgchange -ay
```

8.18 How Do I Handle a Network Failure After Services Are Switched from a Windows BMS Booted from an EVS Disk to an HA BMS?

Scenarios

After services are switched from a Windows BMS booted from an EVS disk to an HA BMS, the HA BMS cannot connect to the network.

Procedure


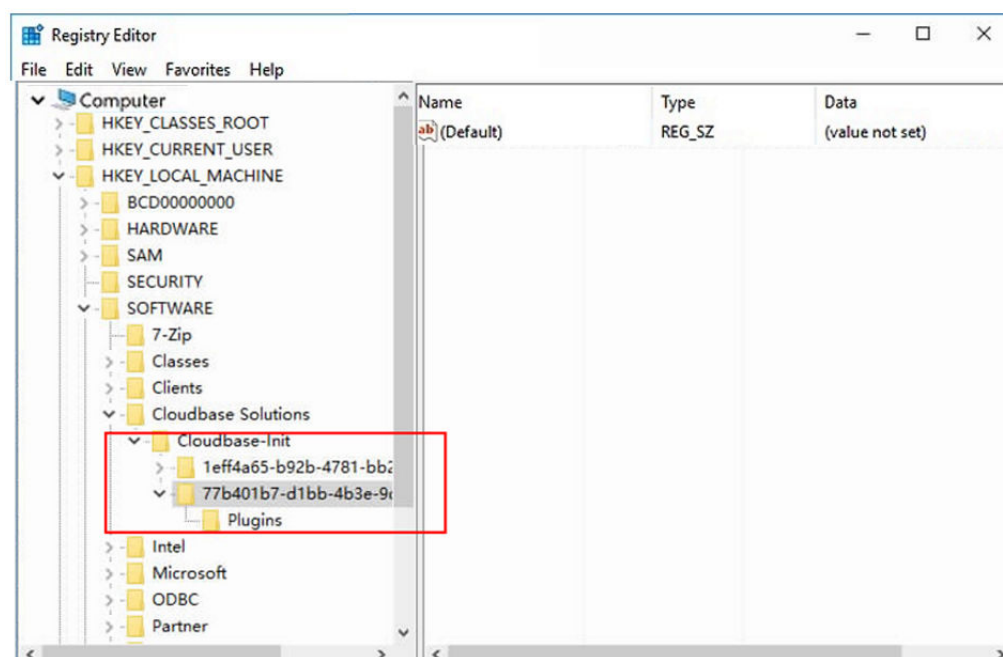

1. Log in to the Windows BMS and create a reserved administrator account.
Choose **Control Panel > User Accounts > Change account type > Add** to add an administrator account.
2. Click  in the lower left corner, choose **Windows PowerShell**, and enter **regedit.exe** to open the registry.
In the **Registry Editor** window, choose **HKEY_LOCAL_MACHINE > SOFTWARE > Cloudbase Solutions > Cloudbase-Init** and delete all folders from **Cloudbase-Init**.

Figure 8-7 Deleting Cloudbase-Init folders



3. Manually delete bond information.
 - a. Delete Team1.
Run the **Get-NetAdapter** command to check whether Team1 is there.
If there is a Team1, open **Windows PowerShell** as an administrator and run the following command to delete it:
Remove-NetLbfoTeam -Name Team1
 - b. Delete hidden network adapters.
Choose **Computer Management > Device Manager > View > Show hidden devices > Network adapters**. Right-click the adapter (it should be displayed in gray) and choose **Uninstall device**. Uninstall other grayed out adapters, one by one, in the same way.

Figure 8-8 Network adapters



4. Restart the Windows BMS.
5. Use the reserved administrator account to log in to the BMS. (The original administrator account becomes invalid after the BMS is restarted.)

After network access is restored, change the password of the original administrator account. Then, delete the reserved account.