Bare Metal Server

User Guide

Issue   08
Date   2020-09-10

Contents

1 Common Operations

When using BMSs, you may encounter various problems, such as remotely logging in to a BMS, expanding disk capacity, resetting the password, and reinstalling the OS. This section provides navigation to common operations to help you with these problems.

Create and Manage a BMS

General Operations

Perform the following steps to use a BMS:

1. Create a BMS by following the instructions in **Creating a BMS**.

   If the BMS quota is insufficient, you can apply to increase the quota by following the instructions in **Adjusting Resource Quotas**.

2. Log in to the BMS. The login mode varies depending on the BMS OS.

   – Linux BMS: **Remotely Logging In to a BMS**, **Logging In to a BMS Using an SSH Key Pair**, or **Logging In to a BMS Using a Password**

   – Windows BMS: **Logging In to a BMS Remotely Using MSTSC**

3. Stop the BMS by following the instructions in **Stopping a BMS**.

4. Delete the BMS by following the instructions in **Releasing a BMS**.

Billing Management

You can renew your yearly/monthly BMS in the following ways:

● **Manual Renewal**

● **Auto Renewal**

Configuration Change

If your BMS password is lost or has expired, you can reset the password by following the instructions in **Resetting the BMS Password**.

If the BMS OS cannot meet your requirements, you can reinstall or change the OS by following the instructions in **Reinstalling the OS**.

Refined BMS Control and Management

You can control and manage your BMS in a refined way using the following methods:

- **Injecting User Data into BMSs**
- **Metadata**

BMS Security and Reliability Configuration

You can improve the security and reliability of your BMS using the following methods:

- **Backing Up a BMS**

Create and Manage Private Images

By using private images, you can quickly deploy the service environment.

You can create a private BMS image in the following ways:

- **Creating a Private Image from a BMS**
- **Creating a Private Image from an External Image File**

You can perform the following operations on private images:

- **Share images** with other tenants.
- **Replicate images** across regions.
- **Export images** to your OBS bucket.

Create and Manage Disks

General Operations

To use a disk as a data disk, perform the following steps:

1. Create a disk in any of the following ways:
   - **Purchasing an EVS Disk**
   - **Creating a DSS Disk**

   For the differences between the two methods, see **Disk Types**.
2. **Attach the disk to a BMS**.
3. **Initialize the disk**.
4. **Detach the disk from the BMS**.
5. **Delete the disk**.

Configuration Changes

If the capacity of an existing system disk or data disk cannot meet requirements, you can expand the disk capacity. For details, see **Expanding the Capacity of an EVS Disk**. After the disk capacity has been expanded, the additional disk space needs to be allocated to an existing partition or a new partition.

Create and Use a Key Pair

To use a key pair, perform the following steps:

1. **Create** or **import** a key pair.
2. When creating a BMS, bind the key pair to the BMS.
3. **Log in to the BMS using the key pair**.
4. **Delete the key pair**.

Create and Manage a BMS Network

Security Group

To use a security group, perform the following steps:

1. **Create a security group**.
2. **Add a security group rule**.
3. When creating a BMS, add it to the security group.
4. **Delete the security group rule**.
5. **Delete the security group**.

EIP

To use an EIP, perform the following steps:

1. **Bind an EIP to a BMS**.
2. **Unbind the EIP from the BMS**.

VPC

You can bind an extra IP address (virtual or floating IP address) to a NIC to enable flexible network functions. You can also enable the source/destination check function of the NIC to prevent packet spoofing and improve security.

- **Managing Virtual IP Addresses**
- **Setting the Source/Destination Check for a NIC**

High-Speed Network

Operations related to the high-speed network include:

Managing High-Speed Networks

Enhanced High-Speed Network

The enhanced high-speed network is available only in Chinese mainland regions. Operations related to the enhanced high-speed network include:

- **Adding an Enhanced High-Speed NIC**
- **Deleting an Enhanced High-Speed NIC**
- After adding or deleting an enhanced high-speed NIC, configure the NIC in the OS. For details, see **Configuring an Enhanced High-Speed NIC (SUSE Linux Enterprise Server 12)** to **Configuring an Enhanced High-Speed NIC (Windows Server)**.

User-defined VLAN

Operations related to the user-defined VLAN include:

- **Overview**
- The method of configuring a user-defined VLAN varies for different OSs. For details, see sections **Configuring a User-defined VLAN (SUSE Linux Enterprise Server 12)** to **Configuring a User-defined VLAN (Windows Server)**.

IB Network

Operations related to the IB network include:

- **Overview**

User-defined Network and User-defined Network ACL

The user-defined network is a network type designed for virtualization on BMS. This network provides functions similar to those of the VPC.

- **Creating and Managing a User-defined Network**
- **Overview**
- For details about the application scenarios, see **XenServer on BMS**.

Virtualization

BMS supports virtualization. Cases of XenServer, VMware, and Hyper-V deployment on BMSs are provided as examples.

- **XenServer on BMS**
- **VMware on BMS**
- **Hyper-V on BMS**

Tags

You can use tags to identify various resources to improve the efficiency in classifying, querying, and managing resources. To use a tag, perform the following steps:

1. **Add a tag**.
2. **Query resources by tag**.
3. **Delete a tag**.

Monitor a BMS

To meet the basic monitoring and O&M requirements for servers, **Server Monitoring** monitors more than 40 metrics, such as CPU, memory, disk, and network metrics. You need to install Agent on the BMS to implement the monitoring. For details, see **Overview**.

For all the supported BMS metrics, see **Supported Monitoring Metrics (with Agent Installed)**.

2 Instance

2.1 Creating a BMS

2.1.1 Creating a BMS

Scenarios

This section describes how to create a BMS to deploy your services.

Prerequisites

- You have completed **Preparations**.
- To configure user data injection, you have prepared **user data scripts**.
- You have enabled Dedicated Cloud (DeC).

  For details, see **Creating a Dedicated BMS**.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Click **Buy BMS**.

   The page for you to purchase a BMS is displayed.
4. In the **Current Configuration** area on the right pane, confirm the billing mode. Currently, only the **Yearly/Monthly** billing mode is supported.

   ☐ NOTE

   - **Yearly/Monthly**: a prepaid billing mode, in which a BMS is billed based on the purchase period. This mode is more cost-effective than the pay-per-use mode and applies if the resource usage period can be estimated.
   - **Pay-per-use**: a postpaid billing mode. You can use BMSs without paying any fees in advance and you only need to pay for the time you use the BMSs.

5. Confirm **Region**.

   If the region is incorrect, click ⊙ in the upper left corner of the page to correct it.

6. Select an AZ.

   An AZ is a physical region where resources use independent power supply and networks. AZs are physically isolated but interconnected through an internal network.

   – It is recommended that you create BMSs in different AZs to ensure high availability of applications running on the BMSs.

   – To lower the network delay, create BMSs in the same AZ.

7. Select a flavor.

   Flavor contains the flavor name and the CPU, memory, local disks, and extended configuration of the BMS. After you select a flavor, the name and use scenarios of the flavor are displayed under the flavor list.

   Extended Configuration provides the NIC information of the selected flavor. For example, 2 x 2*10GE indicates that a 10GE NIC with two ports connects to the VPC network and a 10GE extended NIC with two ports supports high-speed interconnection between BMSs.

   📖 **NOTE**

   ● Configuration in the flavor, such as the CPU, memory, and local disks, cannot be changed.

   ● The bandwidth of different BMS flavors varies. Choose a flavor that meets your requirements.

   ● Some flavors support quick BMS provisioning. If you select a flavor of this type, parameter **System Disk** is displayed under **Disk**. The OS of this type of BMS is installed on an EVS disk.

8. Set **Image**.

   – Public Image

     A public image is a common standard OS image that is available to all users. It contains an OS and pre-installed public applications, such as the SDI iNIC driver, bms-network-config (a network configuration program), and Cloud-Init (an initialization tool). You can configure the applications and software in a public image as needed.

   – Private Image

     A private image is an image available only to the user who created it using an external image file or a BMS. It contains an OS, preinstalled public applications, and the user's private applications. Using a private image to create BMSs frees you from configuring BMSs repeatedly.

   – Shared Image

     A shared image is a private image shared by another public cloud user with you.

9. Set **License Type**.

   Set a license type for using an OS or software on the cloud platform. This parameter is available only if the image you selected is charged.

   – Use license from the system

Allows you to use the license provided by the cloud platform. Obtaining the authorization of such a license is charged.

– Bring your own licenses (BYOL)

Allows you to use your existing OS license. In such a case, you do not need to apply for a license again.

10. Set **Disk**.

Disks are classified as EVS disks and DSS disks based on whether the disks use dedicated storage resources. DSS disks provide dedicated storage resources.

– If you have applied for a storage pool on the DSS console and have obtained the pool, click the **DSS** tab and create disks in the storage pool.

– If you have not obtained a dedicated storage pool, click the **EVS** tab and create EVS disks that use public storage resources.

📖 NOTE

● When you use DSS resources to create a disk, the disk type must be the same as that of the requested storage pool. For example, both are of the high I/O type.

● For details about different disk types, see **Disk Types**.

A BMS has one system disk and one or more data disks. You can add multiple data disks for a BMS and customize the system disk size.

– System disk

If you select a flavor that supports quick provisioning, parameter **System Disk** is available. You can set the system disk type and size as needed.

– Data disk

You can add multiple data disks for a BMS and enable sharing for each data disk.

▪ Currently, BMSs only support SCSI disks.

▪ **Share**: indicates that the EVS disk can be shared. A shared disk can be attached to multiple BMSs simultaneously.

📖 NOTE

● After a system disk is detached from a BMS charged in yearly/monthly mode, the disk can only be used as a system disk and can only be attached to this BMS.

● If you detach a non-shared data disk purchased when you buy a BMS charged in yearly/monthly mode and want to attach it again, you can only attach it to the original BMS as a data disk.

● The non-shared data disk purchased when you buy a BMS charged in yearly/ monthly mode does not support separate renewal, unsubscription, automatic service renewal, conversion to pay-per-use payment, or release.

11. Configure automatic backup.

After automatic backup is enabled, the system automatically backs up the BMS based on the preset backup policy.

📖 NOTE

The automatic backup function applies only to BMSs that support quick provisioning. To enable this function, you must select a flavor that supports quick provisioning in step **7**.

a. Select **Enable auto backup**.

          b.   Configure **Backup Policy**.

In the drop-down list, select a backup policy. Alternatively, you can click **Manage Backup Policy** and set the backup policy on the Cloud Server Backup Service (CSBS) page. If you have not created any backup policy but select **Enable auto backup**, the system will use the default backup policy shown in **Figure 2-1**.

Figure 2-1 Default backup policy

For details about CSBS, see **Cloud Backup and Recovery Overview**.

12.   Set network parameters, including **VPC**, **NIC**, and **Security Group**.

When you use VPC for the first time, the system automatically creates a VPC for you, including the security group and NIC. The default subnet segment is 192.168.1.0/24 and the subnet gateway is 192.168.1.1. Dynamic Host Configuration Protocol (DHCP) is enabled for the subnet.

Table 2-1 Network parameters

Parameter	Description
VPC	You can select an existing VPC or create one.
NIC	Includes primary and extension NICs. You can add an extension NIC for a BMS and specify IP addresses for the primary and extension NICs. **CAUTION** - The primary NIC cannot be deleted because it is used to provide the default route. - If you choose to assign an IP address automatically, 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. - If a fixed IP address is assigned to a NIC, you cannot create BMSs in a batch.
High-Speed NIC	The high-speed NIC provides high-speed network ports for communication between BMSs. It provides high bandwidth. You cannot configure the same high-speed network for multiple high-speed NICs of a BMS. **NOTE** In some regions, high-speed networks have been upgraded to enhanced high-speed networks with higher performance. You need to configure enhanced high-speed NICs.

Parameter	Description
Enhanced High-Speed NIC	A BMS has a maximum of two enhanced high-speed NICs and depends on the total bandwidth of the extension NICs. For example, if the total bandwidth allowed for the extension NICs is 2 x 10GE and the bandwidth of the first enhanced high-speed NIC is 2 x 10GE, you cannot add another enhanced high-speed NIC. **NOTE** You can view the total bandwidth of extension NICs in **Extended Configuration**. ● If a flavor's **Extended Configuration** contains **2*10GE** (for example, the **Extended Configuration** of flavor physical.h2.large is **1*100G IB + 2*10GE**), BMSs of this flavor has only one NIC without extension NIC, and the total bandwidth of extension NICs is 0. ● If a flavor's **Extended Configuration** contains **2 x 2*10GE** (for example, the **Extended Configuration** of flavor physical.s3.large is **2 x 2*10GE**), BMSs of this flavor has two NICs, of which one is an extension NIC, and the total bandwidth of extension NICs is 2*10GE.

Parameter	Description
Security Group	A security group implements access control for BMSs within a security group and between different security groups. You can define different access control rules for a security group, and these rules take effect for all BMSs added to this security group. When creating a BMS, you can select only one security group. After a BMS is created, you can associate it with multiple security groups. For details, see **Changing a Security Group**. Security group rules determine BMS access and usage. For instructions about how to configure a security group rule, see **Adding Security Group Rules**. Enable the following common protocols and ports as needed: ● Port 80: used to view web pages by default through HTTP. ● Port 443: used to view web pages through HTTPS. ● ICMP: pings BMSs to check their communication statuses. ● Port 22: reserved for logging in to a Linux BMS using SSH. ● Port 3389: reserved for logging in to a Windows BMS using SSH. **NOTE** Before initializing a BMS, ensure that security group rules in the outbound direction meet the following requirements: ● Protocol: TCP ● Port Range: 80 ● Remote End: 169.254.0.0/16 If you use the default outbound security group rule, the preceding requirements are met, and the BMS can be initialized. The default outbound security group rule is as follows: ● Protocol: Any ● Port Range: Any ● Remote End: 0.0.0.0/16

Parameter	Description
EIP	An EIP is a static public IP address bound to a BMS in a VPC. Using the EIP, the BMS provides services externally. You can select one of the following three options for **EIP** as needed: ● **Automatically assign**: The system automatically assigns an EIP with a dedicated bandwidth to the BMS. The bandwidth is configurable. ● **Use existing**: An existing EIP is assigned to the BMS. ● **Not required**: The BMS cannot communicate with the Internet and can only be used to deploy services or clusters in a private network. **NOTE** If you select **Use existing**, you can create only one BMS at a time.
Specifications	This parameter is required when you select **Automatically assign** for **EIP**. ● **Dynamic BGP**: When changes occur on a network using dynamic BGP, network configurations can be promptly adjusted using the specified routing protocol, ensuring network stability and optimal user experience. ● **Static BGP**: When changes occur on a network using static BGP, carriers cannot adjust network configurations in real time to ensure optimal user experience.
Bandwidth Type	This parameter is mandatory when **EIP** is set to **Automatically assign**. ● **Dedicated**: The bandwidth can be used by only one EIP. ● **Shared**: The bandwidth can be used by multiple EIPs. **NOTE** ● A bandwidth can be shared between a limited number of EIPs. If the number of EIPs cannot meet service requirements, switch to a higher shared bandwidth or apply to expand the EIP quota of the existing bandwidth. ● EIPs that are charged yearly/monthly do not support shared bandwidths. ● When a shared bandwidth that is charged yearly/monthly expires, the system automatically deletes the bandwidth and creates an exclusive bandwidth charged by traffic for the EIPs sharing the deleted bandwidth.

Parameter	Description
Billed By	This parameter is required when you select **Automatically assign** for **EIP**. ● **Bandwidth**: You specify a maximum bandwidth and pay for the time you use the bandwidth. ● **Traffic**: You are charged based on the actual traffic you have used.
Bandwidth	This parameter is required when you select **Automatically assign** for **EIP**. Specifies the bandwidth size in Mbit/s.

13. Set the BMS login mode.

    Key pair is recommended because it features higher security than **Password**. If you select **Password**, ensure that the password meets complexity requirements described in **Table 2-2** to prevent malicious attacks.

    – Key pair

      A key pair is used for BMS login authentication. You can select an existing key pair, or click **Create Key Pair** to create one.

      ◫ NOTE

        If you use an existing key pair, ensure that you have saved the key file locally. Otherwise, you will fail to log in to the BMS.

    – Password

      In this mode, the initial password is used for authentication. You can log in to the BMS using the username and its initial password.

      If the BMS runs Linux, you can use username **root** and its initial password to log in to the BMS. If the BMS runs Windows, you can use username **Administrator** and its initial password to log in to the BMS. The passwords must meet the requirements described in **Table 2-2**.

Table 2-2 Password requirements

Parameter	Requirements	Example Value

Password	● Consists of 8 characters to 26 characters.
	● Must contain at least three of the following character types:
	  – Uppercase letters
	  – Lowercase letters
	  – Digits
	  – Special characters !@$%^-_=+[]{}:,./?
	● Cannot contain the username or the username spelled backwards.
	● Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows BMSs.)
		Test12$@

14. Configure **Enterprise Project**.

    This parameter is available only when you have enabled the enterprise project function, or your account is an enterprise account. To enable this function, contact your customer manager.

    An enterprise project is a cloud resource management mode, in which cloud resources and members are centrally managed by project. The default project is **default**.

    Select an enterprise project from the drop-down list. For details about enterprise projects, see **Enterprise Management User Guide**.

15. (Optional) Configure **Advanced Settings**.

    To use functions listed in **Advanced Settings**, click **Configure now**. Otherwise, click **Do not configure**.

    – **Agency**

      An agency provides BMSs with temporary security credentials for accessing other cloud services. The agency is created by the tenant administrator on the IAM console.

      If you have created an agency in IAM, you can select the agency from the drop-down list and obtain specified operation permissions. If you have no agency, click **Create Agency** to create one. Currently, agencies are mainly used for server monitoring. For more information, see **Overview**.

16. Set **BMS Name**.

    The name can be customized but can contain only letters, digits, underscores (_), hyphens (-), and periods (.).

    If you buy multiple BMSs at a time, suffixes will be added to the BMSs in sequence, such as **bms-0001**, **bms-0002**, … If you buy multiple BMSs again,

the values in the new BMS names increase from the existing maximum value. For example, the existing BMS with the maximum number in name is **bms-0010**. If you enter **bms**, the names of the new BMSs will be **bms-0011**, **bms-0012**, ... When the value reaches 9999, it will start from 0001 again.

17. Set **Required Duration** and **Quantity**.

    – **Required Duration**: Set the service duration if you select the **Yearly/Monthly** billing mode. The service duration ranges from one month to one year.

    📖 NOTE

    BMSs charged in yearly/monthly mode cannot be deleted. They support only resource unsubscription. 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 **Unsubscribes** page, click **Unsubscribe**.

    - 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.

    – **Quantity**: You can purchase BMSs of the remaining quota at a time.

    📖 NOTE

    If you manually set an IP address for the user-defined NICs or high-speed NIC, or specify an EIP, you can buy only one BMS at a time.

18. Click **Buy Now**. If you have any question about the price, click **Pricing details**.

    On the **Order Details** page, confirm the BMS information and click **Submit Order**.

19. Pay the fees as prompted and click **OK**.

    The BMS console is displayed.

20. Wait for the system to create your requested BMSs.

    The BMS status changes to **Running** after about 30 minutes. If you select a flavor that supports quick provisioning, you can obtain a BMS within about five minutes.

    📖 NOTE

    You can view the BMS creation status. For details, see **Viewing BMS Creation Statuses**.

Follow-up Operations

- After the BMS is created, you can view its details, such as name/ID, disks, and private IP address. For details, see **Viewing BMS Details**.

- After logging in to the BMS, you can install software or deploy services as needed. The login mode varies depending on the BMS OS. For details, see **Linux BMS Login Methods** or **Windows BMS Login Methods**.

- If you have created data disks when creating the BMS, you must format partitions of the data disks. For details, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

- Change the validity period of the password to prevent any inconvenience caused by password expiration. For detailed operations, see **How Do I Set the Password Validity Period?**

- BMSs created using public images have the one-click password reset plugin by default. If your BMS does not have the password reset plugin, or if you want to check whether the plugin is installed, see **(Optional) Installing the One-Click Password Reset Plug-in**.

- Some types of BMSs require drivers. For details about how to install drivers, see **Installing Drivers and Toolkits**.

- Currently, Windows Server 2012 BMSs have the same security identifier (SID), which is used to identify users, groups, and computer accounts. In cluster deployment scenarios, change the SIDs of BMSs by following the instructions in **How Do I Change the SID of a Windows Server 2012 BMS?** to ensure that each BMS has a unique SID.

2.1.2 Creating a BMS Supporting Quick Provisioning

Scenarios

When provisioning a common BMS, you need to download its OS from the cloud and install it. The download costs a long time. BMSs using EVS disks as system disks can be provisioned quickly.

BMSs supporting quick provisioning have the following advantages over other BMSs:

- BMSs booted from EVS disks can be provisioned within about 5 minutes.

- BMSs support CSBS backups, ensuring 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 without the need of configuring the BMSs again.

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 to obtain a BMS.

Procedure

You can create a BMS supporting quick provisioning by following the instructions in **Creating a BMS**.

When creating the BMS, pay attention to the following parameters:

- **Flavor**: Select **physical.s4.medium**, **physical.s4.large**, **physical.s4.xlarge**, **physical.s4.2xlarge**, or **physical.s4.3xlarge**. For more information about flavors, see **Instance Family**.

- **Image**: Select a public image that supports quick provisioning.

- **Disk**: Set the system disk type and size.

- **Auto Backup**: You are advised to select **Enable auto backup** and set **Backup Policy** to ensure data security.

2.1.3 Creating a Dedicated BMS

Scenarios

Resources in a DeC are physically isolated from those in public resource pools. If your services have high security compliance requirements, you can create BMSs in a DeC in either of the following ways:

- **Create a BMS on the DeC Console**
- **Create a BMS on the Cloud Server Console**

Before creating a BMS in a DeC, you must apply for a dedicated BMS resource pool.

Prerequisites

You have enabled DeC. For details, see **Enabling a DeC**.

Apply for a Dedicated BMS Pool

1. Log in to the management console.
2. Click the region name in the upper left corner and select the region where DeC resides from the drop-down list.

   Figure 2-2 Selecting the region where DeC resides

3. Choose **Service List** > **Dedicated Cloud** > **Dedicated Bare Metal Server**.

   The **Dedicated Bare Metal Server** page is displayed.
4. In the upper right corner, click **Apply for Resources**.
5. Select a flavor based on your service requirements and set the quantity and usage duration.
6. Click **Next**. After confirming that the configurations are correct, click **Submit**.

   Message **Request submitted successfully.** is displayed. The application will be reviewed by the O&M personnel. After the application is approved, you can choose **Fees** > **My Orders** and pay the order.
7. After paying the order, you can view information about the resource pool on the **Dedicated Bare Metal Server** page, such as **Resource Pool Type**, **CPU Allocation Rate**, and **Memory Allocation Rate**.

Method 1: Create a Dedicated BMS on the DeC Console

1. Log in to the management console.
2. Click the region name in the upper left corner and select the region where DeC resides from the drop-down list.

3. Choose **Service List** > **Dedicated Cloud** > **Dedicated Bare Metal Server**.

   The **Dedicated Bare Metal Server** page is displayed.

4. In the upper right corner of the page, click **Provision BMS in DeC**.

   The page for creating a BMS is displayed.

5. Set the parameters as prompted. These parameters are the same as those for creating a common BMS. For details, see **Creating a BMS**.

   After the BMS is created, the number of BMSs on the **Dedicated Bare Metal Server** page becomes **1**, and **CPU Allocation Rate** and **Memory Allocation Rate** increase.

Method 2: Create a Dedicated BMS on the Cloud Server Console

1. Log in to the management console.

2. Click the region name in the upper left corner and select the region where DeC resides from the drop-down list.

3. Choose **Service List** > **Computing** > **Bare Metal Server**.

   The Cloud Server Console is displayed.

4. On the BMS page, click **Provision BMS in DeC** in the upper right corner.

   The page for creating a BMS is displayed.

5. Set the parameters as prompted. These parameters are the same as those for creating a common BMS. For details, see **Creating a BMS**.

   After the BMS is displayed, click the **BMS Resource Pool** tab in the **Resource Usage Details** area on the **Dashboard** page. The number of BMSs is **1**, and **CPU Allocation Rate** and **Memory Allocation Rate** increase.

2.1.4 Creating a BMS Using a Private Image

Scenarios

If you want to create a BMS that has the same OS and applications as an existing BMS, you can create a private image using the existing BMS and then create a BMS using the private image. This frees you from repeatedly configuring BMSs and improves efficiency.

Background

You can create a private image using either of the following methods:

- **Creating a Private Image from a BMS**
- **Creating a Private Image from an External Image File**

Procedure

Create a BMS by following the instructions in **Creating a BMS**.

Note for setting the parameters:

- **Region**: Select the region where the private image is located.

- **Flavor**: OSs supported by different BMS flavors vary. For details, see **OSs Supported by Different Types of BMSs**. Select a flavor based on the private image OS.

- **Image**: Select **Private image** or **Shared image** and select the required image from the drop-down list.

- **Disk**: If the selected flavor supports quick provisioning, you are advised to increase **System Disk** by 2 GB or more.

2.2 Viewing BMS Information

2.2.1 Viewing BMS Creation Statuses

Scenarios

After clicking **Submit** to request a BMS, you can query the task status in the **Task Status** area. A task involves several sub-tasks, such as creating a BMS resource, binding an EIP, and attaching an EVS disk.

The task status may be either **Creating** or **Failed**:

- **Processing**: The system is processing the task.

- **Failed**: The system has failed to process the task. The system rolls back the failed task and displays an error code, for example, **(BMS.3033) Failed to create system disk**.

This section describes how to query BMS application processing status and the information displayed in the **Task Status** area.

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. **Task Status** is displayed on the right of common operations, such as **Start**, **Stop**, **Restart**, and **Delete**. After you purchase a BMS, the **Task Status** area will show the task processing status.

   Figure 2-3 BMS application status

	Name/ID	AZ	Status	Flavor
	bms-xxxx-zjq 6014f351-642f-...	eu-de-01	➡ Running	CPU:2*16Core Memory:24*16G DDR4 Local Disk:598GB Extended Configuration:2*10GE

   Start   Stop   Restart   Delete   Task Status: 🔴 4

4. Click the number displayed in the **Task Status** area to view details about the BMS application processing status. The tasks in **Processing** and **Failed** statuses are displayed.

📖 **NOTE**

If **Failed** is displayed for a task in the **Task Status** area, but the BMS list contains the BMS, handle this issue by following the instructions in **Why Is Failed Displayed for a BMS Application Task But the BMS List Contains the Obtained BMS?**

2.2.2 Viewing BMS Details
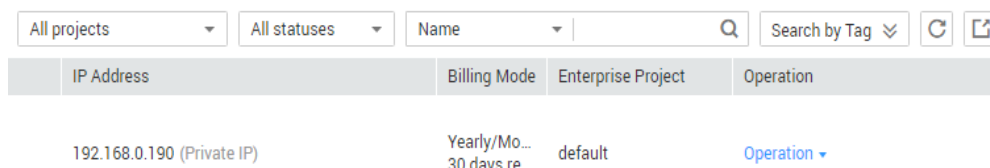
Scenarios

After you obtain a BMS, you can view and manage your BMS on the management console. This section describes how to query detailed information about a BMS, such as the BMS name/ID, disks, NICs, and EIP.

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   On the BMS list page, you can view your BMS and its flavor, image, and private IP address.

3. In the upper right corner of the BMS list, query BMSs by specifying the project, status, name, BMS ID, flavor, and private IP address. Alternatively, click **Search by Tag** above the upper right corner of the BMS list and search for a BMS by tag key and value.

   Figure 2-4 Searching BMSs

All projects ▾	All statuses ▾	Name ▾	🔍	Search by Tag ⌄	C	⬈
IP Address			Billing Mode	Enterprise Project	Operation
192.168.0.190 (Private IP)			Yearly/Mo… 30 days re…	default	Operation ▾

4. Click the name of the queried BMS.

   The page showing details of the BMS is displayed.

5. View the BMS details, such as name, status, flavor, and VPC. You can also click the **Disks**, **NICs**, **Security Groups**, **EIPs**, **Tag**, and **Monitoring** tabs to attach EVS disks to or detach EVS disks from the BMS, change the security group, bind an EIP to or unbind an EIP from the BMS, and create agencies.

   📖 **NOTE**

   The BMS monitoring data and charts are not displayed on the BMS details page. You need to view them on the Cloud Eye console. The prerequisite is that Agent has been installed on your BMS. For details, see **Overview**.

2.3 Logging In to a Linux BMS

2.3.1 Linux BMS Login Methods

Choose an appropriate method to log in to a Linux BMS based on the BMS network configuration and your on-premise OS.

Table 2-3 Linux BMS login methods

Access to the Internet	On-premise OS	Login Method
Yes/No	Windows or Linux	**Remotely Logging In to a BMS**
Yes	Windows	Use a remote login tool, such as PuTTY. ● For how to log in to a BMS using an SSH key pair, see **Logging In to a BMS Using an SSH Key Pair**. ● For how to log in to a BMS using a password, see **Logging In to a BMS Using a Password**.
Yes	Linux	Run commands. ● For how to log in to a BMS using an SSH key pair, see **Logging In to a BMS Using an SSH Key Pair**. ● For how to log in to a BMS using a password, see **Logging In to a BMS Using a Password**.

2.3.2 Remotely Logging In to a BMS

Scenarios

If common remote connection software (such as PuTTY) is unavailable, you can use the remote login function on the management console to log in to a BMS.

Constraints

● Only Linux BMSs support remote login.

● Only the user who creates a BMS or users with the Tenant Administrator or Server Administrator role can log in to the BMS remotely.

● The remote login page does not support Chinese input or desktop GUI-based operations.

● When you log in to a BMS remotely, shortcut keys such as Ctrl and Alt are not well supported. For example, if you enter **Alt +** *ASCII code*, multiple special characters are displayed.

● Before exiting the management console, log out of the OS.

Prerequisites

- The BMS must be in **Running** state.
- You have set a login password when creating the BMS. If you did not set a password or forget the password, you can reset the password by following the instructions in **Resetting the BMS Password with a Few Clicks**.
- If you selected the key pair login mode when creating the BMS, log in to the BMS by following the instructions in **SSH Key Pair** and set a password for the BMS. The detailed operations are as follows:

  Log in to the BMS using the key pair, switch to user **root**, and run the **passwd** command to set a password for user **root**.

  Figure 2-5 Setting a password for user **root**

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
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 password to log in.

   ☐ NOTE

   - If you do not log in within 10 minutes after obtaining the remote login link, it will become invalid.
   - If you do not perform any operation on the remote login page within 10 minutes, you need to obtain the link again.
   - If the login page does not respond after you press **Enter**, a possible cause is that remote login is not configured for the BMS image. You can resolve the issue by following the instructions in **What Do I Do If the Login Page Does Not Respond?**
   - If the BMS console is displayed improperly (such as broken lines and garbled characters) after you remotely log in to it, see **What Do I Do If the BMS Console Is Displayed Improperly After I Remotely Log In to a BMS?**
   - If numbers are not properly displayed after you enter them using the numeric keypad for remote login, see **What Do I Do If the Numeric Keypad Does Not Work During Remote Login?**

2.3.3 Logging In to a BMS Using an SSH Key Pair

Prerequisites

- The BMS must be in **Running** state.
- You have obtained the private key file used during BMS creation.

- You have bound an EIP to the BMS. For details, see **Binding an EIP to a BMS**.
- You have configured the inbound rules of the security group. For details, see section **Adding Security Group Rules**.
- The network connection between the login tool (such as PuTTY) and the target BMS is normal. For example, the default port 22 is not blocked by the firewall.

Logging In to the Linux BMS from a Windows Computer

You can use the following methods to log in to a Linux BMS from a local PC running Windows:

Method 1: Use PuTTY to log in to the BMS.

Before logging in to the BMS using PuTTY, ensure that the private key file has been converted to .ppk format.

1. Check whether the private key file has been converted to **.ppk** format.
   - If yes, go to step **7**.
   - If no, go to step **2**.
2. Visit the following website and download PuTTY and PuTTYgen:

   https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html

   📖 **NOTE**

   > PuTTYgen is a private key generator, which is used to create a key pair that consists of a public key and a private key for PuTTY.

3. Run PuTTYgen.
4. In the **Actions** area, click **Load** and import the private key file that you stored when creating the BMS.

   Ensure that the private key file is in the format of **All files (*.*)**.
5. Click **Save private key**.
6. Save the converted private key, for example, **kp-123.ppk**, to your local PC.
7. Double-click **PUTTY.EXE**. The **PuTTY Configuration** page is displayed.

Figure 2-6 PuTTY Configuration
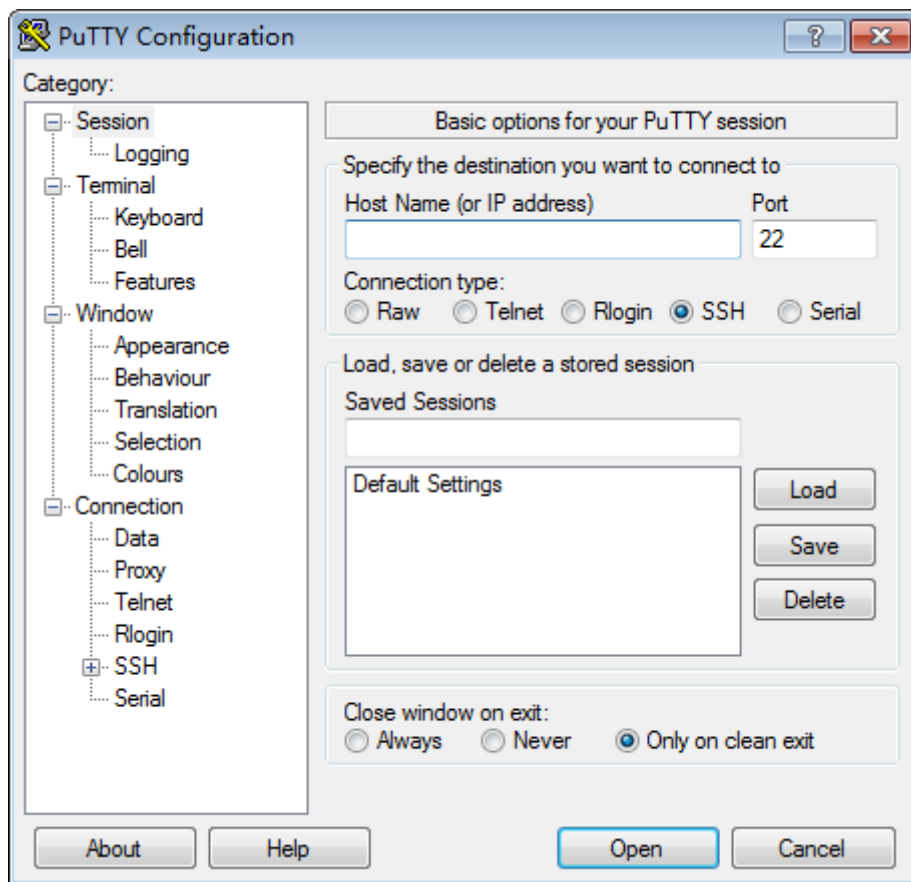
8. Choose **Connection** > **Data**. Enter the image username **root** in **Auto-login username**.

9. Choose **Connection** > **SSH** > **Auth**. In the last configuration item **Private key file for authentication**, click **Browse** and select the .ppk private key in step **6**.

10. Choose **Session** and enter the EIP of the BMS in the box under **Host Name (or IP address)**.

11. Click **Open**.

   Log in to a BMS.

Method 2: Use Xshell to log in to the BMS.

1. Start the Xshell tool.

2. Run the following command to remotely log in to the BMS through SSH:

   ssh *Username***@***EIP*

   Example:

   ssh root@192.168.0.1

3. (Optional) If the system displays the **SSH Security Warning** dialog box, click **Accept & Save**.

Figure 2-7 SSH Security Warning

4. Select **Public Key** and click **Browse** beside the user key text box.

5. In the user key dialog box, click **Import**.

6. Select the locally stored key file and click **Open**.

7. Click **OK** to log in to the BMS.

Logging In to the Linux BMS from a Linux Computer

Perform the following operations to log in to a Linux BMS from a local PC running Linux: The following procedure uses private key file **KeyPair-ee55.pem** as an example to describe how to log in to the BMS.

1. On the Linux CLI, run the following command to change operation permissions:

   chmod 400 /_path_/_KeyPair-ee55_

   📖 NOTE

   > In the preceding command, **_path_** refers to the path under which the key file is stored.

2. Run the following command to log in to the BMS:

   ssh -i /_path_/_KeyPair-ee55_ **root@**_EIP of the BMS_

   📖 NOTE

   - In the preceding command, _path_ refers to the path under which the key file is stored.
   - In the preceding command, **root** is the username of the BMS image.

2.3.4 Logging In to a BMS Using a Password

Prerequisites

- The BMS must be in **Running** state.
- You have bound an EIP to the BMS. For details, see **Binding an EIP to a BMS**.
- You have configured the inbound rules of the security group. For details, see **Adding Security Group Rules**.

- The network connection between the login tool (such as PuTTY) and the target BMS is normal. For example, the default port 22 is not blocked by the firewall.

📖 **NOTE**

If you want to use a password to log in a Linux BMS, log in to the BMS remotely by following the instructions in **Remotely Logging In to a BMS** and enable the SSH password login mode. For details, see **How Do I Set SSH Configuration Items?**

Log In to a BMS from a Windows PC

You can use the following methods to log in to a Linux BMS from a local PC running Windows (for example, use PuTTY):

📖 **NOTE**

Download PuTTY from **https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html**.

1. Run PuTTY.
2. In the navigation pane on the left, choose **Session**, enter the EIP of the BMS in the text box under **Host Name (or IP address)**, and select **SSH** for **Connection type**.
3. Choose **Windows** > **Translation** and select **UTF-8** from the **Received data assumed to be in which character set:** drop-down list box.
4. Click **Open**.
5. Enter username **root** and the password you set to log in to the BMS.

Log In to a BMS from a Linux PC

To log in to a Linux BMS from a Linux PC, run the following command:

ssh *EIP of the BMS*

2.4 Logging In to a Windows BMS

2.4.1 Windows BMS Login Methods

Currently, you can only log in to a Windows BMS remotely by running MSTSC on your local PC. An EIP must be bound to the BMS.

2.4.2 Logging In to a BMS Remotely Using MSTSC

Prerequisites

- The BMS must be in **Running** state.
- If a Windows BMS uses the key pair authentication mode, you have obtained the password for logging in to the BMS. For details, see **Obtaining the Password of a Windows BMS**.
- You have bound an EIP to the BMS. For details, see **Binding an EIP to a BMS**.

- You have configured the inbound rules of the security group. For details, see **Adding Security Group Rules**.

- The network connection between the login tool and the target BMS is normal. For example, the default port 3389 is not blocked by the firewall.

Procedure

The following procedure describes how to log in to a Windows BMS using **mstsc.exe**.

1. On the local PC, click **Start**.

2. In the **Search programs and files** box, enter **mstsc.exe** and press **Enter**.

3. Enter the EIP and username of the Windows BMS, click **Connect**, enter the password as prompted, and click **OK** to log in to the BMS.

2.5 Managing BMSs

2.5.1 Changing the Name of a BMS

Scenarios

To make it easy for you to identify and manage each BMS, HUAWEI CLOUD allows you to set BMS names and change the names at any time. The new name of a BMS takes effect after the BMS is restarted.

Constraints

The names of Windows BMSs cannot be changed.

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

    The BMS console is displayed.

3. Click the name of the BMS whose name is to be changed.

4. Click ✎ next to **Name**, enter a new name that meets requirements, and click

    ✔ to save the change.

    The BMS name can contain only letters, digits, hyphens (-), underscores (_), and periods (.).

5. Log in to the BMS OS and run the following command to enable automatic hostname synchronization:

    vi /opt/huawei/network_config/bms-network-config.conf

    Set the value of **auto_synchronize_hostname** to **True**.

    auto_synchronize_hostname = True

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

📖 NOTE

> If the value of **auto_synchronize_hostname** is **False**, after the BMS is restarted, the hostname will be automatically changed to that set during BMS creation.

6. 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.

2.5.2 Resetting the BMS Password

Scenarios

If you forget the password for logging in to a BMS or if you want to harden the password to improve security, you can reset the password on the console.

⚠️ CAUTION

If you change the password of a running BMS on the console, the BMS will be automatically restarted during password resetting. To prevent data loss, it is recommended that you reset the password during off-peak hours to minimize the impact on your services.

Prerequisites

- The password resetting function depends on the CloudResetPwdAgent plug-in, which is installed for public images by default. If your BMS is created from a private image, check whether the plug-in has been installed by following the instructions in **(Optional) Installing the One-Click Password Reset Plug-in**.

- Ensure that DHCP is enabled in the VPC to which the BMS belongs.

- The BMS network connectivity is normal.

- An EIP has been bound to the BMS.

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Locate the row that contains the target BMS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list.

4. Set and confirm a new password as prompted.

Figure 2-8 Resetting the password

Reset Password   ✕

BMS Name   ecc-2-card-ecc-64-nonstop

Username   root

New Password   [Enter the password]

Confirm Password [Enter the password agai]

It takes about 10 minutes to reset the password. After the password is reset, the system
will automatically restart the BMS.
If the new password does not take effect, reset the password again.

[OK]   [Cancel]

The new password must meet the complexity rules in **Table 2-4**.

Table 2-4 Password requirements

Parameter	Requirement	Example Value
Password	- Consists of 8 characters to 26 characters. - Must contain at least three of the following character types: - Uppercase letters - Lowercase letters - Digits - Special characters Windows: !@$%-_=+[]:./? Linux: !@%^-_=+[]{}:,./? - Cannot contain the username or the username spelled backwards. - Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows BMSs.)	Test12@#

5.   Click **OK**.

     It takes about 10 minutes for the system to reset the password. Do not
     repeatedly perform this operation. During the process, the BMS will be
     restarted automatically. After the BMS is restarted, use the new password to
     log in to the BMS to check whether the password is reset successfully.

Related Operations

You can reset the BMS password using an API. For details, see **Resetting the BMS Password by a Few Clicks**.

You can also **change the login password in the BMS OS**. After changing the password, you must restart the BMS on the management console to make the new password take effect. It is recommended that you change the password on the console.

Helpful Links

What Do I Do If a Service Port Is Used by a One-Click Password Reset Plug-in?

2.5.3 Stopping a BMS

Scenarios

You can stop BMSs in **Running** state.

Stopping a BMS charged in yearly/monthly mode does not affect the BMS fees. If other service products, such as EVS disks, EIPs, and bandwidths are bound to the BMS, these products are billed using their own billing mode (yearly/monthly or pay-per-use).

☐ NOTE

If you choose to forcibly stop a BMS, services running on the BMS will be stopped. Before performing this operation, ensure that you have saved files on the BMS.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

    The BMS console is displayed.
3. Locate the row that contains the target BMS, click **More** in the **Operation** column, and select **Stop** from the drop-down list. To stop multiple BMSs, select them and click **Stop** at the top of the BMS list.
4. In the displayed dialog box, click **Yes**.

After a BMS is stopped, its status becomes **Stopped**.

You can perform the following operations only when the BMS is stopped:

- **Detaching the System Disk**
- **Creating an Image**
- **Rebuilding a BMS**

2.5.4 Restarting a BMS

Scenarios

You can restart BMSs on the console. Only BMSs in running state can be restarted.

📖 NOTE

> Restarting a BMS will interrupt your services. Exercise caution when performing this operation.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Locate the row that contains the target BMS, click **More** in the **Operation** column, and select **Restart** from the drop-down list. To restart multiple BMSs, select them and click **Restart** at the top of the BMS list.
4. In the displayed dialog box, click **Yes**.

2.5.5 Reinstalling the OS

Scenarios

If the OS of a BMS fails to start, suffer from viruses, or requires optimization, reinstall the OS.

The original image is used to reinstall the BMS OS. BMSs provisioned on local disks and quickly provisioned BMSs both support OS reinstallation.

After the OS is reinstalled:

- The system disk type of the quickly provisioned BMS does not change.
- The IP address and MAC address of the BMS do not change.

Precautions

Reinstalling the OS is a mission-critical operation. Before performing this operation, read the following precautions carefully:

- To reinstall the OS, you must stop the BMS, which will interrupt your services.
- Reinstalling the OS clears the data in all partitions of the system disk. Back up data before performing this operation.
- Do not power off or restart the BMS during the OS reinstallation. Otherwise, the reinstallation may fail.
- After the OS is reinstalled, custom configurations, such as DNS and hostname of the original OS will be reset. You must reconfigure the OS.

Constraints

- The reinstalled OS must be the same as the original OS.
- During the OS reinstallation, the system disk capacity of a BMS provisioned using a local disk is not displayed.
- If the EVS disk where the BMS OS is installed is deleted during the OS reinstallation, the reinstallation will fail.
- During the OS reinstallation, you cannot inject user data.

- The OS of a BMS in maintenance state cannot be reinstalled.

Prerequisites

- The BMS must be in **Stopped** or **Reinstalling OS failed** state.
- If the boot device of the BMS is the EVS disk, the EVS disk quota must be greater than 0.
- If it is a quick-provisioning BMS, ensure that the BMS has a system disk.
- If the BMS is created using a private image, ensure that the image is still available.
- The OS reinstallation depends on the bms-network-config and Cloud-Init plug-ins in the BMS image.
  - If the BMS is created using a public image, ensure that the image has the bms-network-config and Cloud-Init plug-ins.
  - If the BMS is created using a private image, check whether bms-network-config and Cloud-Init are installed by following the instructions in **Bare Metal Server Private Image Creation Guide**.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Locate the row containing the target BMS, click **More** in the **Operation** column, and select **Reinstall OS** from the drop-down list.

   The **Reinstall OS** dialog box is displayed.

   Figure 2-9 Reinstalling the BMS OS
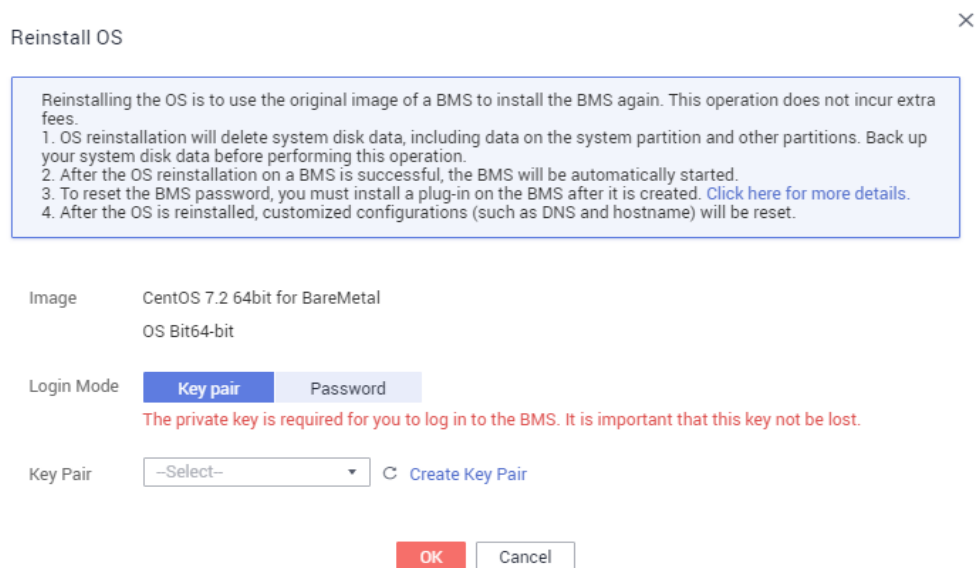
4. Set **Login Mode**.
   - **Key pair**: You can select an existing key pair or click **Create Key Pair** and create a private key used to log in to the BMS.

- **Password**: You can set the initial password for logging in to the BMS OS. The new password must meet the password complexity requirements listed in **Table 2-7**.

5. Click **OK**.

6. On the **BMS OS Reinstallation** page, confirm the OS configuration and click **Submit**.

   After the application is submitted, the BMS status changes to **Reinstalling OS**. The reinstallation is complete when the BMS status changes to **Running**. After the OS is reinstalled, the BMS will start automatically.

   ◻ **NOTE**

   Do not perform any operation on the temporary BMS during the reinstallation process.

Follow-up Operations

If the QinQ network has been configured for the BMS before the OS reinstallation, configure the network by following the instructions in sections **Configuring a User-defined VLAN (SUSE Linux Enterprise Server 12)** to **Configuring a User-defined VLAN (Windows Server)** after the OS is reinstalled.

2.5.6 Rebuilding a BMS

Scenarios

If the BMS cannot work properly due to hardware or SDI card damage, you can rebuild the BMS. This section describes how to rebuild a BMS.

◻ **NOTE**

The BMS is not automatically rebuilt. You need to contact the administrator to have your BMS rebuilt.

Notes

- Currently, only BMSs that are quickly provisioned can be rebuilt.
- After a BMS is rebuilt, it will start automatically.
- If the BMS uses an IB NIC, record the IP address of the IB NIC rebuilding the BMS.
- If the BMS uses a QinQ network, record the IP address of the QinQ network before rebuilding the BMS.

Restrictions and Limitations

- A BMS can only be rebuilt in the same POD.
- A BMS to be rebuilt must use an EVS disk as its system disk.
- Data on local disks cannot be migrated after a BMS is rebuilt.

Prerequisites

- The BMS to be rebuilt must be stopped.
- The BMS to be rebuilt must have a system disk.
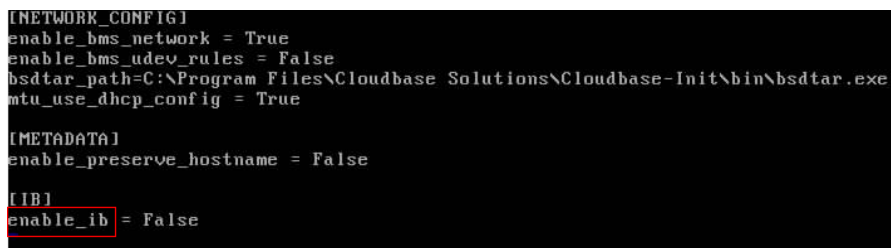
Procedure

1. If your BMS uses a QinQ network, delete configurations of the original QinQ network before rebuilding the BMS. For example, if eth3 and eth5 form port group bond1 for the QinQ network, delete the following configuration files:

   rm /etc/udev/rules.d/80-persistent-net.rules

   rm /etc/sysconfig/network-scripts/ifcfg-eth3

   rm /etc/sysconfig/network-scripts/ifcfg-eth5

   rm /etc/sysconfig/network-scripts/ifcfg-bond1

2. Contact the administrator and apply for rebuilding the BMS.

   – If your BMS uses the QinQ network, reconfigure the QinQ network based on the original QinQ network configuration and by following the instructions in **Configuring a User-defined VLAN (SUSE Linux Enterprise Server 12)** to **Configuring a User-defined VLAN (Windows Server)** after the BMS is rebuilt.

   – If your BMS uses the IB network and the IB NIC IP address assignment mode is DHCP, the IP address of the BMS will change after it is rebuilt. Therefore, if your service heavily depends on the IP address, you need to reconfigure the IP address of the IB network using the static configuration method. The operations describe how to set the IP address of the IB NIC to the original IP address.

      i. Log in to the BMS OS.

      ii. Create the **/etc/sysconfig/network-scripts/ifcfg-ib0** configuration file. The following uses CentOS as an example. Set **IPADDR** to the IP address of the BMS before it is rebuilt.
```
#/etc/sysconfig/network-scripts/ifcfg-ib0
DEVICE=ib0
ONBOOT=yes
BOOTPROTO=none
IPADDR=172.31.0.254
NETWORK=172.31.0.0
BROADCAST=172.31.0.255
NETMASK=255.255.255.0
```

      iii. Change the value of **enable_ib** in the **/opt/huawei/network_config/bms-network-config.conf** file to **False**.

      Figure 2-10 Changing the parameter value

      iv. Save the configuration and exit. Then restart the NIC.

      ifdown ib0

      ifup ib0

      v. Run the following command to check whether the configured IP address takes effect:

      ifconfig ib0

2.5.7 Backing Up a BMS

Scenarios

To ensure data security, you can back up all EVS system and data disks of a BMS. This backup mode prevents data inconsistency caused by the difference in the backup creation time. The Cloud Server Backup Service (CSBS) offers the backup service for BMSs. It works based on the consistent snapshot technology for Elastic Volume Service (EVS) disks. With CSBS, you can use backup data to restore BMS data, ensuring data security and correctness.

Restrictions and Limitations

- BMS backups cannot be used to create images.
- BMSs with shared EVS disks cannot be backed up.
- When the BMS is restored using backup, the BMS will automatically stop, which will interrupt tenant services. After the BMS is stopped, it is locked for a specified time period during which tenants cannot perform operations on the BMS.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Locate the row that contains the target BMS, click **More** in the **Operation** column, and select **Create Backup**.

   The **Create CSBS Backup** page is displayed.
4. Perform the following operations as prompted:
   - Select a BMS: By default, the BMS to be backed up is selected in the BMS list. Retain the default.
   - Configure the backup: Select **Auto Backup** and select a backup policy.

     ☐ NOTE

     After the selected BMS is associated with the backup policy, the BMS will be automatically backed up based on the backup policy.

     If the selected BMS has been associated with other policy, it will be disassociated from the original policy automatically and then associated with the new policy.

   You can also select **back up now**. The selected BMS will be backed up immediately.

   For more information, see **Cloud Backup and Recovery Getting Started**.

2.5.8 Releasing a BMS

Scenarios

You can delete BMSs you no longer need. Once the status of a BMS becomes **Deleted**, no fees will be incurred for the BMS.

After a BMS is deleted, it is still displayed in the BMS list for a short period of time, after which it will be deleted from the BMS list. Tags and disks of the BMS will be disassociated from the BMS, and data on the disks will be deleted.

📖 NOTE

You can manually release a BMS billed in yearly/monthly mode after the validity period ends. If you do not renew the BMS, it will be released automatically. Before the BMS expires, you can apply for unsubscription to release the BMS in advance.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

    The BMS console is displayed.
3. Locate the row that contains the target BMS, click **More** in the **Operation** column, and select **Unsubscribe** from the drop-down list.
4. On the **Unsubscribes** page, select a reason and click **Unsubscribe**.
5. Enter the verification code you receive and click **OK**.

2.6 Using User Data and Metadata

2.6.1 Injecting User Data into BMSs

Application Scenarios

Use the user data injection function to inject user data into BMSs to:

- Use scripts to simplify BMS configuration.
- Use scripts to initialize the BMS OS configuration.
- Upload your scripts to BMSs during BMS creation.
- Use scripts to perform other operations.

Restrictions

- For Linux:
  - The image that is used to create BMSs must have Cloud-Init installed.
  - The user data to be injected must be less than or equal to 32 KB.
  - User data uploaded as text can contain only ASCII characters. User data uploaded as a file can contain any characters, and the file size must be less than or equal to 32 KB.
  - The image that is used to create ECSs must be a public image, a private image created based on a public image, or a private image with Cloud-Init installed.
  - The format of the customized scripts must comply with user data script specifications.
  - DHCP must be enabled for the VPC, and port 80 must be enabled for the security group in the outbound direction.

- If the password login mode is selected, user data injection is not supported.
- For Windows:
  - The image that is used to create BMSs must have Cloudbase-Init installed.
  - The user data to be injected must be less than or equal to 32 KB.
  - User data uploaded as text can contain only ASCII characters. User data uploaded as a file can contain any characters, and the file size must be less than or equal to 32 KB.
  - The image that is used to create ECSs must be a public image, a private image created based on a public image, or a private image with Cloudbase-Init installed.
  - DHCP must be enabled for the VPC, and port 80 must be enabled for the security group in the outbound direction.

Use Methods

1. Create a user data script, the format of which complies with user data script specifications. For details, see section **Related Links**.
2. When creating a BMS, set **Advanced Settings** to **Configure now**, and paste the content of the user data script to the **User Data Injection** text box or upload the user data file.

Figure 2-11 Injecting user data
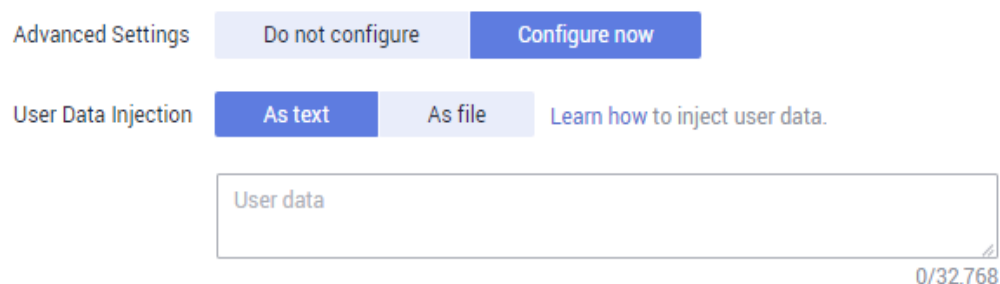
3. The created BMS automatically runs Cloud-Init or Cloudbase-Init to read the user data script upon startup.

User Data Scripts of Linux BMSs

Customized user data scripts of Linux BMSs are based on the open-source Cloud-Init architecture. This architecture uses BMS metadata as the data source for automatically configuring the BMSs. The customized script types are compatible with open-source Cloud-Init. For details about Cloud-Init, see **http://cloudinit.readthedocs.io/en/latest/topics/format.html**.

- Script execution time: A customized user data script is executed after the time when the status of the target BMS changes to **Running** and before the time when **/etc/init** is executed.

  📖 **NOTE**

  By default, the scripts are executed as user **root**.

- Script type: Both user-data scripts and Cloud-Config data scripts are supported.

Table 2-5 Linux BMS script types

-	User-Data Script	Cloud-Config Data
Description	Scripts, such as Shell and Python scripts, are used for custom configurations.	Methods pre-defined in Cloud-Init, such as the Yum source and SSH key, are used for configuring certain BMS applications.
Format	A script must be started with **#!**, for example, **#!/bin/bash** and **#!/usr/bin/env python**. When a script is started for the first time, it will be executed at the rc.local-like level, indicating a low priority in the boot sequence.	The first line must be **#cloud-config**, and no space is allowed in front of it.
Constraint	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.
Frequency	The script is executed only once when the BMS is started for the first time.	The execution frequency varies according to the applications configured on the BMS.

- How can I view the customized user data injected into a Linux BMS?

    a. Log in to the BMS.

    b. Run the following command to view the customized user data as user **root**:

        curl http://169.254.169.254/openstack/latest/user_data

- Script usage examples

    This section describes how to inject scripts in different formats into Linux BMSs and view script execution results.

    Example 1: Inject a User-Data script.

    When creating a BMS, set **User Data Injection** to **Text** and enter the customized user data script.

    ```
    #!/bin/bash
    echo "Hello, the time is now $(date -R)" | tee /root/output.txt
    ```

    After the BMS is created, start it and run the **cat *[file]*** command to check the script execution result.

    ```
    [root@XXXXXXXX ~]# cat /root/output.txt
    Hello, the time is now Mon, 16 Jul 2016 16:03:18+0800
    ```

    Example 2: Inject a Cloud-Config Data script.

    When creating a BMS, set **User Data Injection** to **Text** and enter the customized user data script.

```
#cloud-config
bootcmd:
- echo 192.168.1.130 us.archive.ubuntu.com >> /etc/hosts
```

After the BMS is created, start it and run the **cat /etc/hosts** command to check the script execution result.

Figure 2-12 Viewing the running result

User Data Scripts of Windows BMSs

Customized user data scripts of Windows BMSs are based on the open-source Cloudbase-Init architecture. This architecture uses BMS metadata as the data source for initializing and automatically configuring the BMSs. The customized script types are compatible with open-source Cloudbase-Init. For details about Cloudbase-Init, see **https://cloudbase-init.readthedocs.io/en/latest/userdata.html**.

- Script type: Both batch-processing program scripts and PowerShell scripts are supported.

Table 2-6 Windows BMS script types

-	Batch-Processing Program Script	PowerShell Script
Format	The script must be started with **rem cmd**, which is the first line of the script. No space is allowed at the beginning of the first line.	The script must be started with **#ps1**, which is the first line of the script. No space is allowed at the beginning of the first line.
Constraint	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.

- How can I view the customized user data injected into a Windows BMS?
  a. Log in to the BMS.
  b. Access the following URL in the address box of the browser and view the injected user data:
     http://169.254.169.254/openstack/latest/user_data
- Script usage examples

  This section describes how to inject scripts in different formats into Windows BMSs and view script execution results.

  Example 1: Inject a batch-processing program script.

  When creating a BMS, set **User Data Injection** to **Text** and enter the customized user data script.

```
rem cmd
echo "Hello, BAT Test" > C:\1111.txt
```

After the BMS is created, start it and check the script execution result. In this example, a text file named **1111** is added to disk C:\.

Figure 2-13 Creating text file **1111**
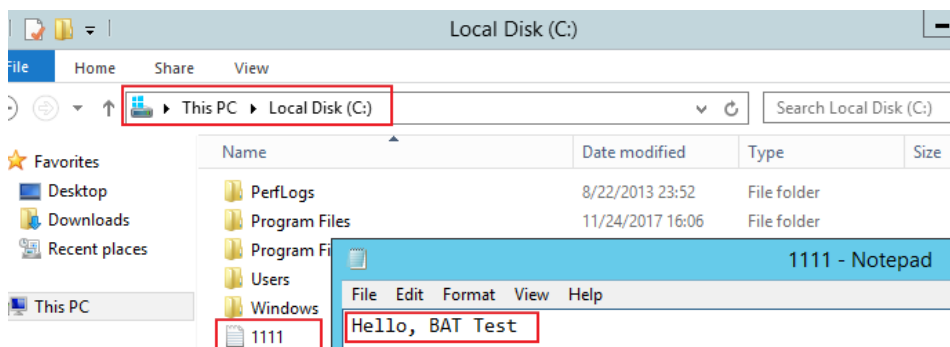
To view the user data injected into the Windows BMS, log in at http://169.254.169.254/openstack/latest/user_data.

Figure 2-14 Viewing user data 1111

Example 2: Inject a PowerShell script.

When creating a BMS, set **User Data Injection** to **Text** and enter the customized user data script.

```
#ps1
echo "Hello, Powershell Test" > C:\aaaa.txt
```

After the BMS is created, start it and check the script execution result. In this example, a text file named **aaaa** is added to disk C:\.

Figure 2-15 Creating text file **aaaa**

To view the user data injected into the Windows BMS, log in at http://169.254.169.254/openstack/latest/user_data.

Figure 2-16 Viewing user data aaaa

Case 1

This case illustrates how to use the user data injection function to simplify BMS configuration.

In this example, vim is configured to enable syntax highlighting, display line numbers, and set the tab stop to **4**. Configuration file **.vimrc** is created and injected into the **/root/.vimrc** directory during BMS creation. After the BMS is created, vim is automatically configured based on your requirements. This helps to improve BMS configuration efficiency, especially in batch ECS creation scenarios.

The content of the script file to be injected is as follows:

```
#cloud-config
write_files:
 - path: /root/.vimrc
   content: |
     syntax on
     set tabstop=4
     set number
```

Case 2

This case illustrates how to use the user data injection function to reset the password for logging in to a Linux BMS.

In this example, the password of user **root** is reset to "******".

📖 **NOTE**

The new password must meet the password complexity requirements listed in **Table 2-7**.

Table 2-7 Password requirements

Parameter	Requirements	Example Value
Password	- Consists of 8 characters to 26 characters. - Must contain at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters !@$%^-_=+[]{}:,./? - Cannot contain the username or the username spelled backwards. - Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows BMSs.)	Test12$@

The content of the script file to be injected is as follows. (Retain the indentation in the following script.)

```
#cloud-config
chpasswd:
  list: |
    root:******
  expire: False
```

After the BMS is created, you can use the reset password to log in to it. To ensure system security, change the password of user **root** after logging in to the BMS for the first time.

Case 3

This case illustrates how to use the user data injection function to create a user on a Windows BMS and configure the password for the user.

In this example, the user's username is **abc**, its password is **********, and the user is added to the **administrators** user group.

📖 NOTE

The new password must meet the password complexity requirements listed in **Table 2-8**.

Table 2-8 Password requirements

Parameter	Requirements	Example Value

Password	• Consists of 8 characters to 26 characters.
	• Must contain at least three of the following character types:
	  – Uppercase letters
	  – Lowercase letters
	  – Digits
	  – Special characters !@$%^-_=+[]{}:,./?
	• Cannot contain the username or the username spelled backwards.
	• Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows BMSs.)
		Test12$@

The content of the script file to be injected is as follows:

```
rem cmd
net user abc ****** /add
net localgroup administrators abc /add
```

After the BMS is created, you can use the created username and password to log in to it.

Case 4

This case illustrates how to use the user data injection function to update system software packages for a Linux BMS and enable the HTTPd service. After the user data is injected, you can use the HTTPd service.

The content of the script file to be injected is as follows:

```
#!/bin/bash
yum update -y
service httpd start
chkconfig httpd on
```

Case 5

This case illustrates how to use the user data injection function to assign user **root** permission for remotely logging in to a Linux BMS. After injecting the file, you can log in to the BMS as user **root** in SSH key authentication mode.

The content of the script file to be injected is as follows:

```
#cloud-config
disable_root: false
runcmd:
- sed -i 's/^PermitRootLogin.*$/PermitRootLogin without-password/' /etc/ssh/sshd_config
- sed -i '/^KexAlgorithms.*$/d' /etc/ssh/sshd_config
- service sshd restart
```

Related Links

For more information about user data injection cases, visit the official Cloud-init/
Cloudbase-init website:

- **https://cloudinit.readthedocs.io/en/latest/**
- **https://cloudbase-init.readthedocs.io/en/latest/**

2.6.2 Metadata

Introduction

The BMS metadata includes BMS basic information on the cloud platform, such as the BMS ID, hostname, and network information. The BMS metadata can be obtained using compatible OpenStack and EC2 APIs listed in **Table 2-9**.

Table 2-9 BMS metadata types

Metadata Type	Metadata Item	Description
OpenStack type	/meta_data.json	This interface is used to query BMS metadata. For the key fields in the BMS metadata, see **Table 2-10**.
	/password	This interface is used to query the BMS password. If a key pair is selected during the creation of a Windows BMS, Cloudbase-Init is used to save the ciphertext password when the BMS is initialized.
	/user_data	This interface is used to query BMS user data. This metadata allows you to specify scripts and configuration files for initializing BMSs. For details, see **Injecting User Data into BMSs**. For password-authenticated Linux BMSs, save the password injection script.
	/network_data.json	This interface is used to query network information of BMSs.
	/securitykey	This interface is used to obtain temporary AKs and SKs. Before obtaining a temporary AK and SK on a BMS, ensure that the **op_svc_ecs** account has been authorized on IAM and that the desired BMS resources have been authorized for management.

Metadata Type	Metadata Item	Description
EC2 type	/meta-data/ hostname	This interface is used to query the host name of a BMS. To remove the suffix **.novalocal** from a BMS, see: **Is the BMS Host Name with Suffix novalocal Normal?**
	/meta-data/ instance-type	This interface is used to query the flavor name of a BMS.
	/meta-data/local-ipv4	This interface is used to query the fixed IP address of a BMS. If there are multiple NICs, only the IP address of the primary NIC is displayed.
	/meta-data/ placement/ availability-zone	This interface is used to query AZ information about a BMS.
	/meta-data/public-ipv4	This interface is used to query the EIP of a BMS. If there are multiple NICs, only the EIP of the primary NIC is displayed.
	/meta-data/public-keys/0/openssh-key	This interface is used to query the public key of a BMS.
	/user-data	This interface is used to query BMS user data.
	/meta-data/ security-groups	This interface is used to query the name of the security group of the BMS.

Table 2-10 Metadata key fields

Parameter	Type	Description
uuid	String	Specifies the BMS ID.
availability_z one	String	Specifies the AZ where the BMS is located.
meta	Dict	Specifies the metadata information, including the image name, image ID, and VPC ID.

Parameter	Type	Description
hostname	String	Specifies the hostname of the BMS. To remove the suffix **.novalocal** from a BMS, see: **Is the BMS Host Name with Suffix novalocal Normal?**
vpc_id	String	Specifies the ID of the VPC where the BMS is located.

The following describes the URI and methods of using the supported BMS metadata.

Prerequisites

- You have logged in to the BMS.
- Security group rules in the outbound direction meet the following requirements:
  - Protocol: TCP
  - Port Range: 80
  - Remote End: 169.254.0.0/16

  ☐ NOTE

  If you use the default security group rules in the outbound direction, the preceding requirements are met, and the metadata can be accessed. The default outbound security group rule is as follows:
  - Protocol: Any
  - Port Range: Any
  - Remote End: 0.0.0.0/16

Metadata (OpenStack Metadata API)

This interface is used to query BMS metadata.

- URI

  /169.254.169.254/openstack/latest/meta_data.json

- Method

  Supports GET requests.

- Example

  The following describes how to use the cURL tool to query the BMS metadata:

  curl http://169.254.169.254/openstack/latest/meta_data.json

```
{
    "random_seed": "rEocCViRS+dNwlYdGIxJHUp+00poeUsAdBFkbPbYQTmpNwpoEb43k9z+96TyrekNKS
+iLYDdRNy4kKGoNPEVBCc05Hg1TcDblAPfJwgJS1okqEtlcofUhKmL3K0fto
+5KXEDU3GNuGwyZXjdVb9HQWU+E1jztAJjjqsahnU+g/tawABTVySLBKlAT8fMGax1mTGgArucn/
WzDcy19DGioKPE7F8ILtSQ4Ww3VClK5VYB/h0x+4r7IVHrPmYX/
bi1Yhm3Dc4rRYNaTjdOV5gUOsbO3oAeQkmKwQ/
NO0N8qw5Ya4l8ZUW4tMav4mOsRySOOB35v0bvaJc6p
```

+50DTbWNeX5A2MLiEhTP3vsPrmvk4LRF7CLz2J2TGIM14OoVBw7LARwmv9cz532zHki/c8tlhRzLmOTXh/
wL36zFW10DeuReUGmxth7IGNmRMQKV6+miI78jm/KMPpgAdK3vwYF/
GcelOFJD2HghMUUCeMbwYnvijLTejuBpwhJMNiHA/NvlEsxJDxqBCoss/Jfe+yCmUFyxovJ
+L8oNkTzkmtCNzw3Ra0hiKchGhqK3BIeToV/kVx5DdF081xrEA
+qyoM6CVyfJtEoz1zlRRyoo9bJ65Eg6JJd8dj1UCVsDqRY1pIjgzE/
Mzsw6AaaCVhaMJL7u7YMVdyKzA6z65Xtvujz0Vo=",
    "uuid": "ca9e8b7c-f2be-4b6d-a639-f10b4d994d04",
    "availability_zone": "lt-test-1c",
    "hostname": "bms-ddd4-l00349281.novalocal",
    "launch_index": 0,
    "meta": {
      "metering.image_id": "3a64bd37-955e-40cd-ab9e-129db56bc05d",
      "metering.imagetype": "gold",
      "metering.resourcespeccode": "physical.s3.small",
      "metering.cloudServiceType": "service.type.ec2",
      "image_name": "CentOS 7.6 64bit",
      "os_bit": "64",
      "vpc_id": "3b6c201f-aeb3-4bce-b841-64756e66cb49",
      "metering.resourcetype": "1",
      "cascaded.instance_extrainfo": "pcibridge:2",
      "os_type": "Linux",
      "charging_mode": "0"
    },
    "project_id": "6e8b0c94265645f39c5abbe63c4113c6",
    "name": "ecs-ddd4-l00349281"
}

User Data (OpenStack Metadata API)

This interface is used to query BMS user data. The value is configured when you create a BMS. It cannot be changed after the configuration.

- URI

  /169.254.169.254/openstack/latest/user_data

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/openstack/latest/user_data

  ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbiBqdXN0IHN1Y2ggYSBkaXJlY3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhlIHBsYWNlIHRvIGdvIG5vdy4gQnV0IHRoZSBza2ga25vd3MgdGhlIHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJucyBiZWhpbmQgYWxsIGNsb3VkcywgYW5kIHlvdSB3aWxsIGtub3cgdG9vLCB3aGVuIHlvdSBsaWZ0IHlvdXJzZWxmIGhpZ2ggZW5vdWdoIHRvIHNlZSBiZXlvbmQgaG9yaXpvbnMuIiNCg0LVJpY2hhcmQgQmFjaA==
  =

NOTE

If user data is not injected during BMS creation, the query result is 404.

Figure 2-17 404 Not Found

Network Data (OpenStack Metadata API)

This interface is used to query network information of a BMS.

- URI

  /openstack/latest/network_data.json

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/openstack/latest/network_data.json

  ```
  {
      "services": [{
          "type": "dns",
          "address": "100.125.1.250"
      },
      {
          "type": "dns",
          "address": "100.125.21.250"
      }],
      "networks": [{
          "network_id": "67dc10ce-441f-4592-9a80-cc709f6436e7",
          "type": "ipv4_dhcp",
          "link": "tap68a9272d-71",
          "id": "network0"
      }],
      "links": [{
          "type": "cascading",
          "vif_id": "68a9272d-7152-4ae7-a138-3ef53af669e7",
          "ethernet_mac_address": "fa:16:3e:f7:c1:47",
          "id": "tap68a9272d-71",
          "mtu": null
      }]
  }
  ```

Security Key (OpenStack Metadata API)

This interface is used to obtain temporary AKs and SKs.

☐ NOTE

- To obtain a temporary AK and SK on a BMS, ensure that the **op_svc_ecs** account has been authorized on IAM and that the desired BMS resources have been authorized for management.
- Temporary AKs and SKs expire an hour later.
- When using temporary AKs and SKs, add **'X-Security-Token':securitytoken** in the message header. **securitytoken** is the value returned when a call is made to the API.

- URI

  /openstack/latest/securitykey

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/openstack/latest/securitykey

User Data (EC2 Compatible API)

This interface is used to query BMS user data. The value is configured when you create a BMS. It cannot be changed after the configuration.

- URI

  /169.254.169.254/latest/user-data

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/latest/user-data

ICAgICAgDQoiQSBjbG91ZCBkb2VsIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbiBqdXN0IHN1Y2ggYSBkaXJlY3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhlIHBsYWNlIHRvIGdvIG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhlIHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJucyBiZWhpbmQgYWxsIGNsb3VkcywgYW5kIHlvdSB3aWxsIGtub3cgdG9vLCB3aGVuIHlvdSBsaWZ0IHlvdXJzZWxmIGhpZ2ggZW5vdWdoIHRvIHNlZSBiZXlvbmQgaG9yaXpvbnMuIg0KLVJpY2hhcmQgQmFjaA==

Hostname (EC2 Compatible API)

This interface is used to query the name of the host accommodating a BMS. The **.novalocal** suffix will be added later.

- URI

  /169.254.169.254/latest/meta-data/hostname

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/latest/meta-data/hostname

  bms-test.novalocal

Instance Type (EC2 Compatible API)

This interface is used to query the flavor name of a BMS.

- URI

  /169.254.169.254/latest/meta-data/instance-type

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/latest/meta-data/instance-type

  physical.o2.medium

Local IPv4 (EC2 Compatible API)

This interface is used to query the fixed IP address of a BMS. If there are multiple NICs, only the IP address of the primary NIC is displayed.

- URI

  /169.254.169.254/latest/meta-data/local-ipv4

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/latest/meta-data/local-ipv4

  192.1.1.2

Availability Zone (EC2 Compatible API)

This interface is used to query AZ information about a BMS.

- URI

  /169.254.169.254/latest/meta-data/placement/availability-zone

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/latest/meta-data/placement/availability-zone

  az1.dc1

Public IPv4 (EC2 Compatible API)

This interface is used to query the EIP of a BMS. If there are multiple NICs, only the EIP of the primary NIC is displayed.

- URI

  /169.254.169.254/latest/meta-data/public-ipv4

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/latest/meta-data/public-ipv4

  46.1.1.2

Public Keys (EC2 Compatible API)

This interface is used to query the public key of a BMS.

- URI

  /169.254.169.254/latest/meta-data/public-keys/0/openssh-key

- Method

  Supports GET requests.

- Example

  curl http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key

ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDI5Fw5k8Fgzajn1zJwLoV3+wMP+6CyvsSiIc/
hioggSnYu/AD0Yqm8vVO0kWlun1rFbdO+QUZKyVr/OPUjQSw4SRh4qsTKf/+eFoWTjplFvd1WCBZzS/
WRenxlwR00KkczHSJro763+wYcwKieb4eKRxaQoQvoFgVjLBULXAjH4eKoKTVNtMXAvPP9aMy2SLgsJNt
Mb9ArfziAiblQynq7UIfLnN3VclzPeiWrqtzjyOp6CPUXnL0lVPTvbLe8sUteBsJZwlL6K4i
+Y0lf3ryqnmQgC21yW4Dzu+kwk8FVT2MgWkCwiZd8gQ/+uJzrJFyMfUOBIklOBfuUENIJUhAB
Generated-by-Nova

2.7 Installing Drivers and Toolkits

2.7.1 Installing the NVIDIA GPU Driver and CUDA Toolkit on a P1 BMS

Scenarios

After a GPU-accelerated P1 BMS (using the physical.p1.large flavor) is created, the NVIDIA GPU driver and CUDA Toolkit must be installed on it for computing acceleration.

Prerequisites

- An EIP has been bound to the BMS.
- You have obtained the required driver installation packages.

Table 2-11 Download paths for the NVIDIA GPU driver and CUDA Toolkit

OS	Driver	How to Obtain
Ubuntu 16.04 and CentOS 7.4	NVIDIA GPU driver installation package: **NVIDIA-Linux-x86_64-375.66.run**	**https://www.nvidia.com/download/driverResults.aspx/118955/en-us**
	CUDA Toolkit installation package: **cuda_8.0.61_375.26_linux.run**	**https://developer.nvidia.com/compute/cuda/8.0/Prod2/local_installers/cuda_8.0.61_375.26_linux-run**

The procedure of installing the NVIDIA GPU driver and CUDA Toolkit varies depending on the OS.

CentOS 7.4

Step 1 Log in to the target BMS and run the following command to switch to user **root**:

su root

Step 2 (Optional) If the **gcc**, **gcc-c++**, **make**, and **kernel-devel** dependency packages do not exist, run the following commands to install the gcc, gcc-c++, make, and kernel-devel tools:

yum install gcc

yum install gcc-c++

yum install make

yum install kernel-devel-`uname -r`

Step 3   (Optional) Add the Nouveau driver to the blacklist.

If the Nouveau driver has been installed and loaded, perform the following operations to add the Nouveau driver to the blacklist to avoid conflicts:

1.   Add **blacklist nouveau** to the end of the **/etc/modprobe.d/blacklist.conf** file.

2.   Run the following commands to back up and reconstruct initramfs:

mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.bak

dracut -v /boot/initramfs-$(uname -r).img $(uname -r)

3.   Run the **reboot** command to restart the BMS.

Step 4   (Optional) If the X service is running, run the **systemctl set-default multi-user.target** command and restart the BMS to enter multi-user mode.

Step 5   (Optional) Install the NVIDIA GPU driver.

If you selected a specified version of NVIDIA GPU driver rather than a version contained in the CUDA Toolkit, perform this step.

1.   Download NVIDIA GPU driver installation package **NVIDIA-Linux-x86_64-**_xxx.yy_**.run** from **https://www.nvidia.com/Download/index.aspx?lang=en**, and upload this package to the **/tmp** directory on the BMS.

Figure 2-18 Searching for the NVIDIA GPU driver package (CentOS 7.4)

2.   Run the following command to install the NVIDIA GPU driver:

sh ./NVIDIA-Linux-x86_64-_xxx.yy_**.run**

3.   Run the following command to delete the installation package:

rm -f NVIDIA-Linux-x86_64-_xxx.yy_**.run**

Step 6   Install the CUDA Toolkit.

1.   Download CUDA Toolkit installation package **cuda_**_a.b.cc_xxx.yy_**linux.run** from **https://developer.nvidia.com/cuda-downloads**, and upload this package to the **/tmp** directory on the BMS.

2.   Run the following command to change the permission to the installation package:

chmod +x cuda__a.b.cc_xxx.yy_**linux.run**

3. Run the following command to install the CUDA Toolkit:

   ./cuda_*a.b.cc_xxx.yy*_linux.run -toolkit -samples -silent -override -- tmpdir=/tmp/

4. Run the following command to delete the installation package:

   rm -f cuda_*a.b.cc_xxx.yy*_linux.run

5. Run the following commands to check whether the installation is successful:

   cd /usr/local/cuda/samples/1_Utilities/deviceQueryDrv/

   make

   ./deviceQueryDrv

   If the command output contains "Result = PASS", the CUDA Toolkit and the NVIDIA GPU driver have been installed successfully.

   ----End

Ubuntu 16.04

Step 1 Log in to the target BMS and run the following command to switch to user **root**:

   sudo root

Step 2 (Optional) If the **gcc**, **g++**, and **make** dependency packages do not exist, run the following commands to install the gcc, g++, and make tools:

   apt-get install gcc

   apt-get install g++

   apt-get install make

Step 3 (Optional) Add the Nouveau driver to the blacklist.

   If the Nouveau driver has been installed and loaded, perform the following operations to add the Nouveau driver to the blacklist to avoid conflicts:

   1. Add the following information to the end of the **/etc/modprobe.d/blacklist.conf** file:
      ```
      blacklist nouveau
      options nouveau modeset=0
      ```

   2. Run the following commands to back up and reconstruct initramfs:

      mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.bak

      sudo update-initramfs -u

   3. Run the **sudo reboot** command to restart the BMS.

Step 4 (Optional) If the X service is running, run the **systemctl set-default multi-user.target** command and restart the BMS to enter multi-user mode.

Step 5 (Optional) Install the NVIDIA GPU driver.

   If you selected a specified version of NVIDIA GPU driver rather than a version contained in the CUDA Toolkit, perform this step.

   1. Download NVIDIA GPU driver installation package **NVIDIA-Linux-x86_64-*xxx.yy*.run** from **https://www.nvidia.com/Download/index.aspx?lang=en**, and upload this package to the **/tmp** directory on the BMS.

Figure 2-19 Searching for the NVIDIA GPU driver package (Ubuntu 16.04)

2. Run the following command to install the NVIDIA GPU driver:

   sh ./NVIDIA-Linux-x86_64-*xxx.yy***.run**

3. Run the following command to delete the installation package:

   rm -f NVIDIA-Linux-x86_64-*xxx.yy***.run**

Step 6 Install the CUDA Toolkit.

1. Download CUDA Toolkit installation package **cuda_***a.b.cc_xxx.yy***_linux.run** from **https://developer.nvidia.com/cuda-downloads**, and upload this package to the **/tmp** directory on the BMS.

2. Run the following command to change the permission to the installation package:

   chmod +x cuda_*a.b.cc_xxx.yy***_linux.run**

3. Run the following command to install the CUDA Toolkit:

   ./cuda_*a.b.cc_xxx.yy***_linux.run -toolkit -samples -silent -override --tmpdir=/tmp/**

4. Run the following command to delete the installation package:

   rm -f cuda_*a.b.cc_xxx.yy***_linux.run**

5. Run the following commands to check whether the installation is successful:

   cd /usr/local/cuda/samples/1_Utilities/deviceQueryDrv/

   make

   ./deviceQueryDrv

   If the command output contains "Result = PASS", the CUDA Toolkit and the NVIDIA GPU driver have been installed successfully.

----End

2.7.2 Installing the NVIDIA GPU Driver and CUDA Toolkit on a P2 BMS

Scenarios

After a GPU-accelerated P2 BMS (using the physical.p2.large flavor) is created, the NVIDIA GPU driver and CUDA Toolkit must be installed on it for computing acceleration.

Prerequisites

- An EIP has been bound to the BMS.
- You have obtained the required driver installation packages.

Table 2-12 Download paths for the NVIDIA GPU driver and CUDA Toolkit

OS	Driver	How to Obtain
Ubuntu 16.04 and CentOS 7.4	NVIDIA GPU driver installation package: **NVIDIA-Linux-x86_64-384.81.run**	**https://www.nvidia.com/download/driverResults.aspx/124722/en-us**
	CUDA Toolkit installation package: **cuda_9.0.176_384.81_linux.run**	**https://developer.nvidia.com/compute/cuda/9.0/Prod/local_installers/cuda_9.0.176_384.81_linux-run**

The procedure of installing the NVIDIA GPU driver and CUDA Toolkit varies depending on the OS.

CentOS 7.4

Step 1 Log in to the target BMS and run the following command to switch to user **root**:

su root

Step 2 (Optional) If the **gcc**, **gcc-c++**, **make**, and **kernel-devel** dependency packages do not exist, run the following commands to install the gcc, gcc-c++, make, and kernel-devel tools:

yum install gcc

yum install gcc-c++

yum install make

yum install kernel-devel-`uname -r`

Step 3 (Optional) Add the Nouveau driver to the blacklist.

If the Nouveau driver has been installed and loaded, perform the following operations to add the Nouveau driver to the blacklist to avoid conflicts:

1. Add **blacklist nouveau** to the end of the **/etc/modprobe.d/blacklist.conf** file.
2. Run the following commands to back up and reconstruct initramfs:

   mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.bak

   dracut -v /boot/initramfs-$(uname -r).img $(uname -r)

3. Run the **reboot** command to restart the BMS.

Step 4 (Optional) If the X service is running, run the **systemctl set-default multi-user.target** command and restart the BMS to enter multi-user mode.

Step 5 (Optional) Install the NVIDIA GPU driver.

If you selected a specified version of NVIDIA GPU driver rather than a version contained in the CUDA Toolkit, perform this step.

1. Download NVIDIA GPU driver installation package **NVIDIA-Linux-x86_64-*xxx.yy*.run** from **https://www.nvidia.com/Download/index.aspx?lang=en**, and upload this package to the **/tmp** directory on the BMS.

   Figure 2-20 Searching for the NVIDIA GPU driver package (CentOS 7.4)

2. Run the following command to install the NVIDIA GPU driver:

   sh ./NVIDIA-Linux-x86_64-*xxx.yy*.run

3. Run the following command to delete the installation package:

   rm -f NVIDIA-Linux-x86_64-*xxx.yy*.run

Step 6 Install the CUDA Toolkit.

1. Download CUDA Toolkit installation package **cuda_*a.b.cc_xxx.yy*_linux.run** from **https://developer.nvidia.com/cuda-downloads**, and upload this package to the **/tmp** directory on the BMS.

2. Run the following command to change the permission to the installation package:

   chmod +x cuda_*a.b.cc_xxx.yy*_linux.run

3. Run the following command to install the CUDA Toolkit:

   ./cuda_*a.b.cc_xxx.yy*_linux.run -toolkit -samples -silent -override -- tmpdir=/tmp/

4. Run the following command to delete the installation package:

   rm -f cuda_*a.b.cc_xxx.yy*_linux.run

5. Run the following commands to check whether the installation is successful:

   cd /usr/local/cuda/samples/1_Utilities/deviceQueryDrv/

   make

   ./deviceQueryDrv

   If the command output contains "Result = PASS", the CUDA Toolkit and the NVIDIA GPU driver have been installed successfully.

----End

Ubuntu 16.04

Step 1 Log in to the target BMS and run the following command to switch to user **root**:

sudo root

Step 2 (Optional) If the **gcc**, **g++**, and **make** dependency packages do not exist, run the following commands to install the gcc, g++, and make tools:

apt-get install gcc

apt-get install g++

apt-get install make

Step 3 (Optional) Add the Nouveau driver to the blacklist.

If the Nouveau driver has been installed and loaded, perform the following operations to add the Nouveau driver to the blacklist to avoid conflicts:

1. Add the following information to the end of the **/etc/modprobe.d/blacklist.conf** file:
   ```
   blacklist nouveau
   options nouveau modeset=0
   ```
2. Run the following commands to back up and reconstruct initramfs:

   mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.bak

   sudo update-initramfs -u
3. Run the **sudo reboot** command to restart the BMS.

Step 4 (Optional) If the X service is running, run the **systemctl set-default multi-user.target** command and restart the BMS to enter multi-user mode.

Step 5 (Optional) Install the NVIDIA GPU driver.

If you selected a specified version of NVIDIA GPU driver rather than a version contained in the CUDA Toolkit, perform this step.

1. Download NVIDIA GPU driver installation package **NVIDIA-Linux-x86_64-***xxx.yy***.run** from **https://www.nvidia.com/Download/index.aspx?lang=en**, and upload this package to the **/tmp** directory on the BMS.

   Figure 2-21 Searching for the NVIDIA GPU driver package (Ubuntu 16.04)

2. Run the following command to install the NVIDIA GPU driver:

   sh ./NVIDIA-Linux-x86_64-*xxx.yy***.run**

3. Run the following command to delete the installation package:

   rm -f NVIDIA-Linux-x86_64-*xxx.yy***.run**

Step 6 Install the CUDA Toolkit.

1. Download CUDA Toolkit installation package **cuda_***a.b.cc_xxx.yy***_linux.run** from **https://developer.nvidia.com/cuda-downloads**, and upload this package to the **/tmp** directory on the BMS.

2. Run the following command to change the permission to the installation package:

   chmod +x cuda_*a.b.cc_xxx.yy***_linux.run**

3. Run the following command to install the CUDA Toolkit:

   ./cuda_*a.b.cc_xxx.yy***_linux.run -toolkit -samples -silent -override --tmpdir=/tmp/**

4. Run the following command to delete the installation package:

   rm -f cuda_*a.b.cc_xxx.yy***_linux.run**

5. Run the following commands to check whether the installation is successful:

   cd /usr/local/cuda/samples/1_Utilities/deviceQueryDrv/

   make

   ./deviceQueryDrv

   If the command output contains "Result = PASS", the CUDA Toolkit and the NVIDIA GPU driver have been installed successfully.

   ----End

2.7.3 Installing the NVIDIA GPU Driver and CUDA Toolkit on a P3 BMS

Scenarios

After a GPU-accelerated P3 BMS (using the physical.p3.large flavor) is created, the NVIDIA GPU driver and CUDA Toolkit must be installed on it for computing acceleration.

Prerequisites

- An EIP has been bound to the BMS.

- You have obtained the required driver installation packages.

Table 2-13 Download paths for the NVIDIA GPU driver and CUDA Toolkit

OS	Driver	How to Obtain
Ubuntu 16.04 and CentOS 7.4	NVIDIA GPU driver installation package: **NVIDIA-Linux-x86_64-384.81.run**	**http://www.nvidia.com/download/driverResults.aspx/124722/en-us**

OS	Driver	How to Obtain
	CUDA Toolkit installation package: **cuda_9.0.176_384.81_lin ux.run**	**https:// developer.nvidia.com/ compute/cuda/9.0/Prod/ local_installers/ cuda_9.0.176_384.81_linux- run**

The procedure of installing the NVIDIA GPU driver and CUDA Toolkit varies depending on the OS.

CentOS 7.4

Step 1 Log in to the target BMS and run the following command to switch to user **root**:

su root

Step 2 (Optional) If the **gcc**, **gcc-c++**, **make**, and **kernel-devel** dependency packages do not exist, run the following commands to install the gcc, gcc-c++, make, and kernel-devel tools:

yum install gcc

yum install gcc-c++

yum install make

yum install kernel-devel-`uname -r`

Step 3 (Optional) Add the Nouveau driver to the blacklist.

If the Nouveau driver has been installed and loaded, perform the following operations to add the Nouveau driver to the blacklist to avoid conflicts:

1. Add **blacklist nouveau** to the end of the **/etc/modprobe.d/blacklist.conf** file.
2. Run the following commands to back up and reconstruct initramfs:

   mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.bak

   dracut -v /boot/initramfs-$(uname -r).img $(uname -r)
3. Run the **reboot** command to restart the BMS.

Step 4 (Optional) If the X service is running, run the **systemctl set-default multi-user.target** command and restart the BMS to enter multi-user mode.

Step 5 (Optional) Install the NVIDIA GPU driver.

If you selected a specified version of NVIDIA GPU driver rather than a version contained in the CUDA Toolkit, perform this step.

1. Download NVIDIA GPU driver installation package **NVIDIA-Linux-x86_64-***xxx.yy***.run** from **https://www.nvidia.com/Download/index.aspx?lang=en**, and upload this package to the **/tmp** directory on the BMS.

2. Run the following command to install the NVIDIA GPU driver:

   sh ./NVIDIA-Linux-x86_64-*xxx.yy***.run**

3. Run the following command to delete the installation package:

   rm -f NVIDIA-Linux-x86_64-*xxx.yy***.run**

Step 6 Install the CUDA Toolkit.

1. Download CUDA Toolkit installation package **cuda_***a.b.cc_xxx.yy***_linux.run** from **https://developer.nvidia.com/cuda-downloads**, and upload this package to the **/tmp** directory on the BMS.

2. Run the following command to change the permission to the installation package:

   chmod +x cuda_*a.b.cc_xxx.yy***_linux.run**

3. Run the following command to install the CUDA Toolkit:

   ./cuda_*a.b.cc_xxx.yy***_linux.run -toolkit -samples -silent -override --tmpdir=/tmp/**

4. Run the following command to delete the installation package:

   rm -f cuda_*a.b.cc_xxx.yy***_linux.run**

5. Run the following commands to check whether the installation is successful:

   cd /usr/local/cuda/samples/1_Utilities/deviceQueryDrv/

   make

   ./deviceQueryDrv

   If the command output contains "Result = PASS", the CUDA Toolkit and the NVIDIA GPU driver have been installed successfully.

   ----End

Ubuntu 16.04

Step 1 Log in to the target BMS and run the following command to switch to user **root**:

   sudo root

Step 2 (Optional) If the **gcc**, **g++**, and **make** dependency packages do not exist, run the following commands to install the gcc, g++, and make tools:

   apt-get install gcc

apt-get install g++

apt-get install make

Step 3 (Optional) Add the Nouveau driver to the blacklist.

If the Nouveau driver has been installed and loaded, perform the following operations to add the Nouveau driver to the blacklist to avoid conflicts:

1. Add the following information to the end of the **/etc/modprobe.d/blacklist.conf** file:
   ```
   blacklist nouveau
   options nouveau modeset=0
   ```

2. Run the following commands to back up and reconstruct initramfs:

   mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.bak

   sudo update-initramfs -u

3. Run the **sudo reboot** command to restart the BMS.

Step 4 (Optional) If the X service is running, run the **systemctl set-default multi-user.target** command and restart the BMS to enter multi-user mode.

Step 5 (Optional) Install the NVIDIA GPU driver.

If you selected a specified version of NVIDIA GPU driver rather than a version contained in the CUDA Toolkit, perform this step.

1. Download NVIDIA GPU driver installation package **NVIDIA-Linux-x86_64-*xxx.yy*.run.** from **https://www.nvidia.com/Download/index.aspx?lang=en**, and upload this package to the **/tmp** directory on the BMS.

   Figure 2-23 Searching the NVIDIA GPU driver package

2. Run the following command to install the NVIDIA GPU driver:

   sh ./NVIDIA-Linux-x86_64-*xxx.yy*.run

3. Run the following command to delete the installation package:

   rm -f NVIDIA-Linux-x86_64-*xxx.yy*.run

Step 6 Install the CUDA Toolkit.

1. Download CUDA Toolkit installation package **cuda_*a.b.cc_xxx.yy*_linux.run** from **https://developer.nvidia.com/cuda-downloads**, and upload this package to the **/tmp** directory on the BMS.

2. Run the following command to change the permission to the installation package:

   chmod +x cuda_*a.b.cc_xxx.yy***_linux.run**

3. Run the following command to install the CUDA Toolkit:

   ./cuda_*a.b.cc_xxx.yy***_linux.run -toolkit -samples -silent -override --tmpdir=/tmp/**

4. Run the following command to delete the installation package:

   rm -f cuda_*a.b.cc_xxx.yy***_linux.run**

5. Run the following commands to check whether the installation is successful:

   cd /usr/local/cuda/samples/1_Utilities/deviceQueryDrv/

   make

   ./deviceQueryDrv

   If the command output contains "Result = PASS", the CUDA Toolkit and the NVIDIA GPU driver have been installed successfully.

6. Run the following command to check whether the driver is running properly:

   nvidia-smi topo -m

   If GPU information is displayed in the command output, the driver is running properly.

   ----End

3 Image

3.1 Private Image Overview

A private image is an image available only to the user who created it. It contains an OS, preinstalled public applications, and a user's private applications. You can create a private image in the following ways:

- **Creating a Private Image from a BMS**

  ☐ NOTE

  Currently, only a BMS that supports quick provisioning (the OS is installed on an EVS disk) can be used to create a private image.

- **Creating a Private Image from an External Image File**

  You can upload external image files to the cloud platform and register them as your private images. Supported external image formats include VMDK, VHD, QCOW2, RAW, VHDX, QED, VDI, QCOW, ZVHD2, and ZVHD.

  ☐ NOTE

  Images of other formats must be converted using the image conversion tool before they can be used on BMSs. For details about how to convert the image format, see **Image Management Service Best Practices**.

After a private image is created successfully, the image status becomes **Normal**. You can use the image to create BMSs or share the image with other tenants. You can also replicate the image to your other regions. The following figure shows how to use private images.

Figure 3-1 Using private images
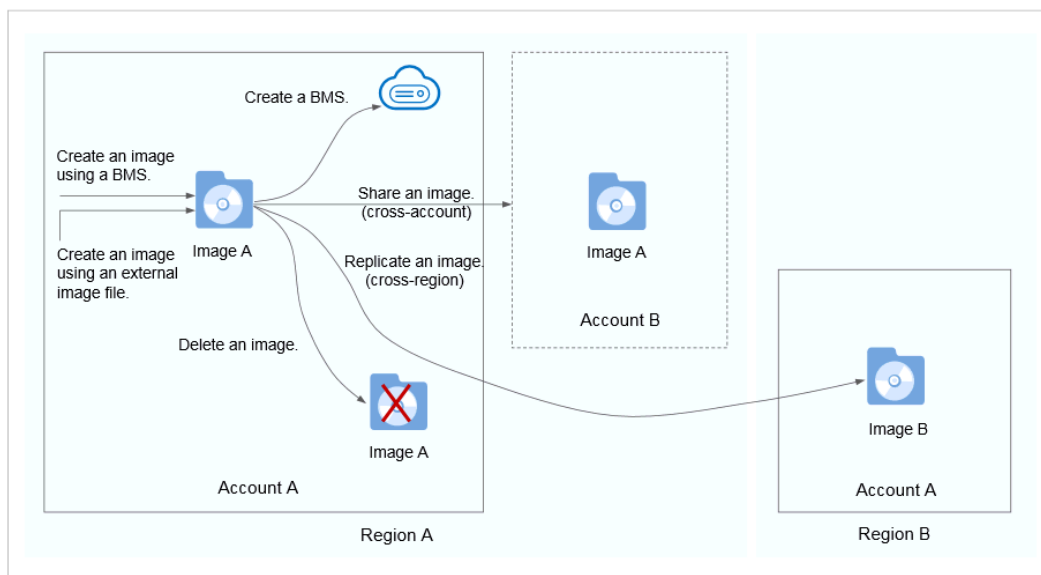
3.2 Creating a Private Image from a BMS

Scenarios

You can create a private image from a BMS and copy the system disk data of the BMS to the private image. The system disk contains an OS and pre-installed applications for running services.

Constraints

- This operation is supported only when the system disk is an EVS disk.

- Data disks of a BMS cannot be exported as images.

- The BMS must be stopped.

- This operation depends on the bms-network-config and Cloud-Init plug-ins in the BMS image.

  – If the BMS is created using a public image, the image has the bms-network-config and Cloud-Init plug-ins by default.

  – If the BMS is created using a private image, check whether bms-network-config and Cloud-Init are installed by following the instructions in **Bare Metal Server Private Image Creation Guide**.

Precautions

- Delete sensitive data from the BMS before using it to creating a private image to prevent data leak.

- Delete residual files from the OS. For details, see **Deleting Files**.

- During the image creation process, do not change the BMS status. Otherwise, the image will fail to be created.

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Locate the row that contains the target BMS, click **More** in the **Operation** column, and select **Stop** from the drop-down list.

   Only a BMS in stopped state can be used to create a private image.

4. After the BMS status changes to **Stopped**, click **More** in the **Operation** column and select **Make Image**.

   The page for creating an image is displayed.

Figure 3-2 Creating a private image

5. Enter the image name, select an enterprise project, set a tag, and enter description as needed.

   Click **Next**.

6. On the displayed **Details** page, confirm the configuration and click **Submit**.

7. Return to the image list. If the status of the private image changes to **Normal**, the private image is created successfully.

Follow-up Operations

If you want to create BMSs using the private image, see **Creating a BMS Using a Private Image**. On the page for creating BMSs, select the private image you have created.

3.3 Creating a Private Image from an External Image File

Scenarios

You can create and register a private image using an external image file. **Figure 3-3** shows the procedure.

Figure 3-3 Creating a private image from an external image file

The procedure contains the following steps:

1. Prepare an image file. For details, see **Bare Metal Server Private Image Creation Guide**.
2. Upload the image file to your OBS bucket. For details, see **Upload an External Image File**.
3. On the management console, select the uploaded image file and register it as a private image. For details, see **Register a Private Image**.

Upload an External Image File

You can import an image file in VHD, VMDK, QCOW2, RAW, VHDX, QCOW, VDI, QED, ZVHD, or ZVHD2 format to create a private image.

Use OBS Browser+ to upload external image files. For details, see **OBS Browser+ Best Practices**.

When uploading the external image file, you must select an OBS bucket with standard storage.

Download OBS Browser+ from the following link:

https://support.huaweicloud.com/intl/en-us/browsertg-obs/obs_03_1003.html

Register a Private Image

1. Log in to the management console.
2. Under **Computing**, click **Image Management Service**.

   The IMS console is displayed.
3. Click **Create Image** in the upper right corner.
4. Configure the following information:

   Image Type and Source

- **Type**: Select **System disk image**.
- **Source**: Select **Image file**.

  In the bucket list, select the bucket that stores the image file and select the image file.

Image Information

- **Function**: Select **BMS system disk image**.

  Ensure that you have completed initialization configuration on the image file by following the instructions in **Bare Metal Server Private Image Creation Guide**.

- **OS**: (Optional) Select the OS of the image file.

  To ensure that the image can be created and used properly, select the OS consistent with that of the image file.

- **System Disk (GB)**: Set the system disk size. You are advised to set the value to the image system disk size plus 2 GB.

- **Name**: Enter a name for the image to be created. The value can contain only letters, digits, spaces, hyphens (-), underscores (_), and periods (.), and cannot start or end with a space.

- **Enterprise Project**: Select an enterprise project from the drop-down list. This parameter is available only when you have enabled the enterprise project function, or your account is an enterprise account. To enable this function, contact your customer manager.

- **Description**: (Optional) Enter description of the image.

5. Click **Next**.

   On the displayed **Details** page, confirm the configuration and click **Submit**.

6. Return to the image list. If the status of the private image changes to **Normal**, the private image is registered successfully.

   📖 **NOTE**

   The time required for registering a private image varies depending on the size of the image file.

Follow-up Operations

You can use the private image to create a BMS by following the instructions in **Creating a BMS Using a Private Image**.

4 Disk

4.1 Disk Types

HUAWEI CLOUD provides various storage products for your BMSs, including block storage based on the distributed storage architecture, dedicated storage based on Huawei enterprise storage architecture, and local disks.

- Block storage refers to EVS disks, which are block-based storage products and adopt a three-copy distributed mechanism. EVS disks provide high reliability, performance, and scalability. You can create or release them at any time.

- Dedicated Distributed Storage Service (DSS) provides dedicated physical storage resources and adopts a three-copy distributed mechanism similar to block storage. It provides high availability and durability, and stable and low latency using multiple technologies, such as data redundancy and cache acceleration.

- Local disks include NVMe SSDs, SATA disks, and others. They provide a low latency, high throughput, and high cost-effectiveness and are applicable to scenarios that have large volumes of data and require high storage I/O performance and real-time performance.

  Because local disks of a single physical server may encounter a single point of failure (SPOF), you are advised to configure data redundancy at the application layer to ensure data availability.

Table 4-1 Comparison of storage products

Storage Product	Storage Type	Typical Application Scenarios	Process
Block storage	Shared storage pools	- Enterprise daily work - Development and testing - Enterprise applications, including SAP, Microsoft Exchange, and Microsoft SharePoint - Distributed file systems - Various databases, including MongoDB, Oracle, SQL Server, MySQL, and PostgreSQL	Create a disk and then attach the disk to the BMS.
DSS	Physically isolated storage pools and dedicated resources	- Hybrid load: DSS supports hybrid deployment of HPC, database, email, OA, and web applications. - High-performance computing - OLAP applications	DSS can be used with BMSs in DeCs or those not in DeCs. - DeC scenario: Enable DeC, apply for a storage pool, create a disk in the storage pool, and attach the disk to the BMS. - Non-DeC scenario: Apply for a storage pool, create a disk in the storage pool, and attach the disk to the BMS.
Local disks	Local disks of servers	- Big data - Distributed cache	Create a BMS and use its local disks directly.

4.2 Attaching Data Disks

Scenarios

If the existing disks of a BMS fail to meet service requirements, for example, due to insufficient disk space or poor disk performance, you can attach more available disks to the BMS, or create more disks and attach them to the BMS.

Constraints

- 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.

Prerequisites

Disks are available.

For details about how to create disks, see **Purchasing an EVS Disk** in *Elastic Volume Service Quick Start*.

📖 **NOTE**

> If DSS is used, see **Dedicated Distributed Storage Service Getting Started**.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. In the upper right corner of the BMS list, enter the name, private IP address, ID, or flavor of a BMS and click   to search for the desired BMS.
4. Click the name of the target BMS.

   The page showing details of the BMS is displayed.
5. Click the **Disks** tab. Then, click **Attach Disk**.

   The **Attach Disk** dialog box is displayed.
6. Select the disk type and target disk, and set the mount point as prompted.

   📖 **NOTE**

   > If no EVS disks are available, click **Create Disk** in the lower part of the list.
7. Click **OK**.

After the disk is attached, you can view the information about it on the **Disks** tab.

Follow-up Operations

If the attached disk is newly created, the disk can be used only after it is initialized (formatted). For details about how to initialize data disks, see **Initializing Data Disks**.

📖 NOTE

After the BMS is restarted, the drive letter of the EVS disk may change. For the mapping between the EVS disk device and drive letter, see **How Do I Obtain the Drive Letter of an EVS Disk?**

4.3 Initializing Data Disks

4.3.1 Introduction to Data Disk Initialization Scenarios and Partition Styles

Scenarios

After a disk is attached to a BMS, you need to log in to the BMS to initialize (format) the disk before you can use the disk properly.

- System disk

  A system disk does not need to be initialized because it is automatically created and initialized during the BMS creation. The default disk partition style is master boot record (MBR).

- Data disk

  - If a data disk is created during the BMS creation, it will be automatically attached to the BMS.

  - If a data disk is created explicitly, you need to manually attach the data disk to the BMS.

  In both cases, the data disk can only be used after it is initialized. Choose a proper disk partition style based on your service plans.

Disk Partition Style

Table 4-2 lists the common disk partition styles. For Linux OSs, different disk partition styles require different partitioning tools.

Table 4-2 Disk partition styles

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Linux Partitioning Tool
Master Boot Record (MBR)	2 TB	4 primary partitions 3 primary partitions and 1 extended partition With the MBR partition style, primary partitions and an extended partition can be included, where the extended partition can contain several logical partitions. For example, if 6 partitions need to be created, you can create the partitions in the following two ways: 3 primary partitions and 1 extended partition, with the extended partition containing 3 logical partitions 1 primary partition and 1 extended partition, with the extended partition containing 5 logical partitions	- fdisk - parted
Guid Partition Table (GPT)	18 EB 1 EB = 1048576 TB	Unlimited Disk partitions allocated using GPT are not categorized.	parted

⚠️ CAUTION

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Currently, an EVS data disk supports up to 32 TB. Therefore, use the GPT partition style if your disk capacity is greater than 2 TB.

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

Partitioning Operation Guide

For a disk with less than 2 TB capacity, see one of the following topics:

- **Initializing a Windows Data Disk (Windows Server 2016)**
- **Initializing a Linux Data Disk (fdisk)**
- **Initializing a Linux Data Disk (parted)**

For a disk with greater than 2 TB capacity, see one of the following topics:

- **Initializing a Windows Data Disk Greater Than 2 TB (Windows Server 2012)**
- **Initializing a Linux Data Disk Greater Than 2 TB (parted)**

4.3.2 Initializing a Windows Data Disk (Windows Server 2016)

Scenarios

This section uses Windows Server 2016 Standard 64bit to describe how to initialize a data disk attached to a BMS running Windows.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is greater than 2 TB. For details about disk partition styles, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

The method for initializing a disk varies depending on the OSs running on the BMS. This document is for reference only. For detailed operations and differences, see the product documents of the OSs running on the corresponding BMSs.

⚠️ CAUTION

When using an EVS disk for the first time, if you have not initialized the disk, including creating partitions and file systems, the additional capacity added to the disk in a later expansion operation may not be normally used.

Prerequisites

- You have logged in to the BMS.
- A data disk has been attached to the BMS and has not been initialized.

Procedure

Step 1 On the BMS desktop, click the start icon in the lower left corner.

The **Windows Server** window is displayed.

Step 2 Click **Server Manager**.

The **Server Manager** window is displayed.

Figure 4-1 Server Manager

Step 3 In the navigation tree on the left, choose **File and Storage Services**.

The **Servers** page is displayed.

Figure 4-2 Servers

Step 4 In the navigation pane, choose **Disks**.

The **Disks** page is displayed.

Figure 4-3 Disks

Step 5 Disks are listed in the right pane. If the new disk is in the offline state, bring it online before initialize it.

1. Right-click the new disk and choose **Bring Online** from the shortcut menu.

   The **Bring Disk Online** dialog box is displayed.

Figure 4-4 Bring Disk Online

2. Click **Yes** to confirm the operation.

3. Click in the upper area of the page to refresh the disk information. When the disk status changes from **Offline** to **Online**, the disk has been brought online.

Figure 4-5 Bring online succeeded
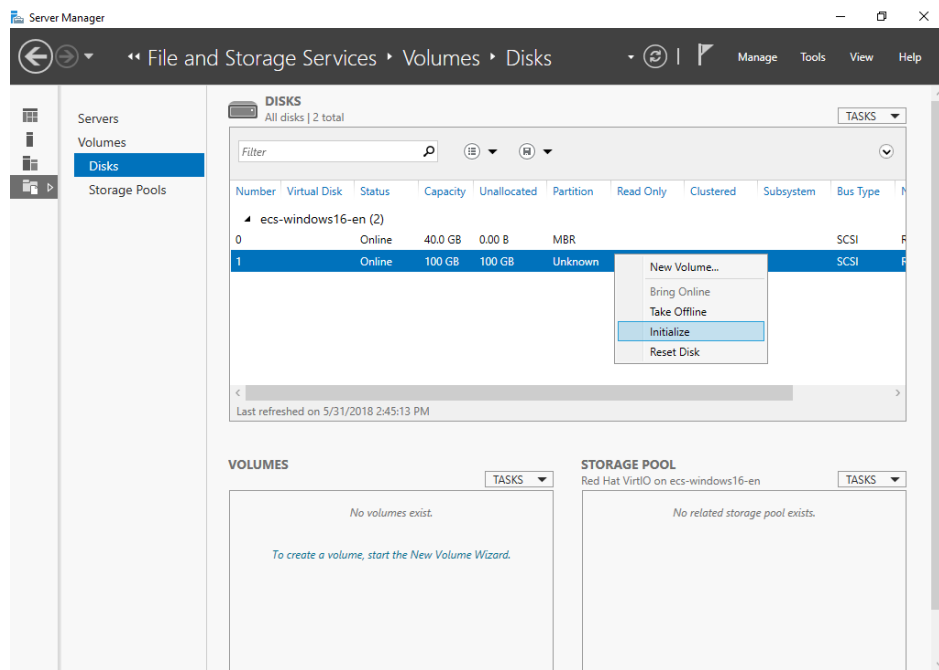
Step 6 After the disk has been brought online, initialize the disk.

1. Right-click the new disk and choose **Initialize** from the shortcut menu. The **Initialize Disk** dialog box is displayed.

Figure 4-6 Initialize Disk (Windows 2016)

2. Click **Yes** to confirm the operation.

3. Click [icon] in the upper area of the page to refresh the disk information.

   When the disk partition changes from **Unknown** to **GPT**, the initialization is complete.

Figure 4-7 Completing initialization
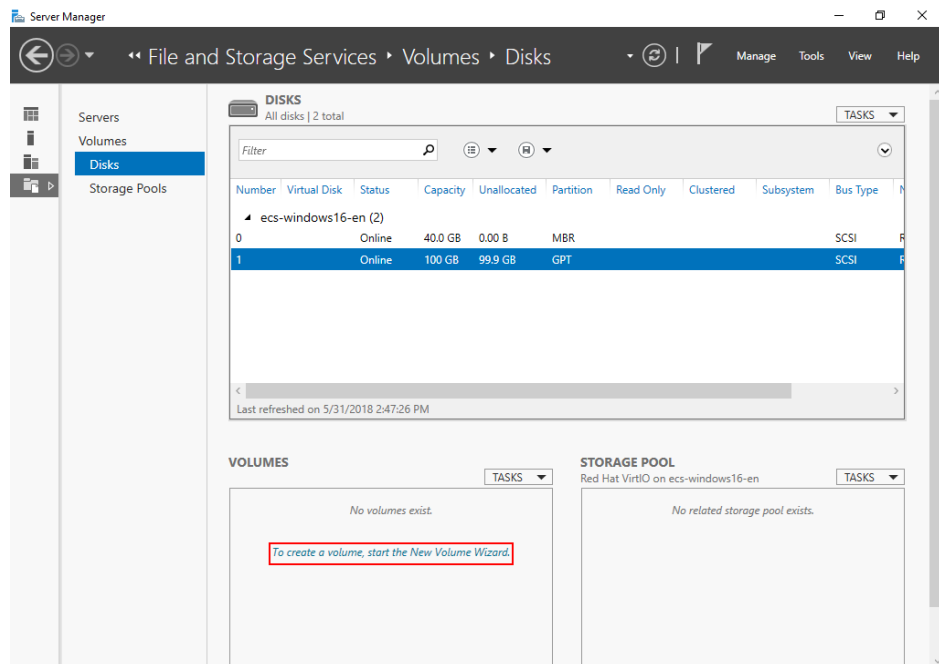
Step 7 In the lower left area of the page, click **To create a volume, start the New Volume Wizard.** to create a new volume.

The **New Volume Wizard** window is displayed.

Figure 4-8 New Volume Wizard

Step 8 Follow the prompts and click **Next**.

The **Select the server and disk** page is displayed.

Figure 4-9 Select the server and disk
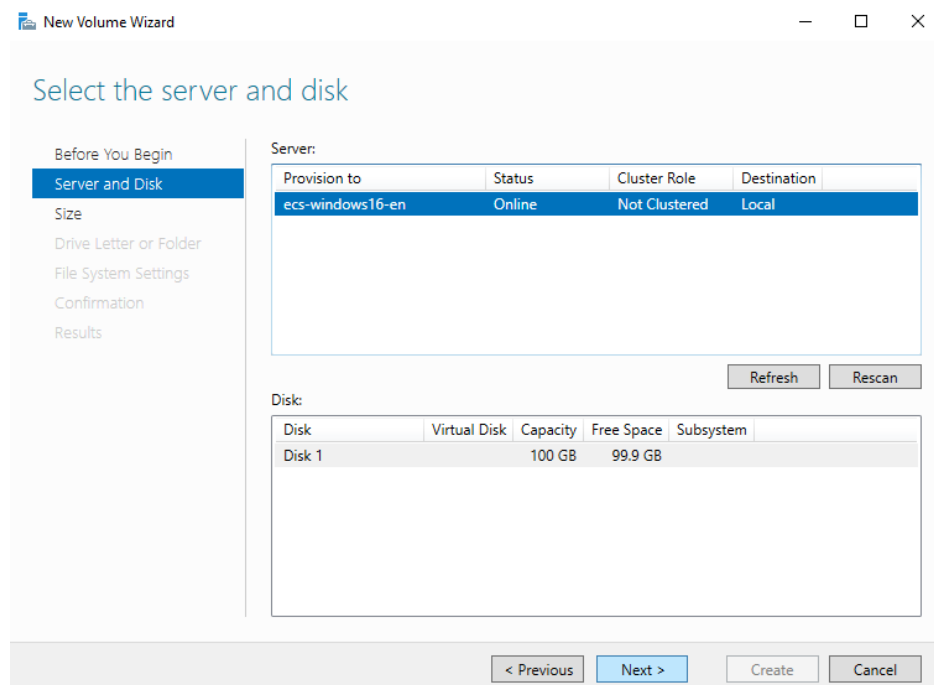
Step 9 Select the server and disk, and then click **Next**. The system selects the server to which the disk is attached by default. You can specify the server based on your requirements. In this example, the default setting is used.

The **Specify the size of the volume** page is displayed.

Figure 4-10 Specify Volume Size (Windows 2016)
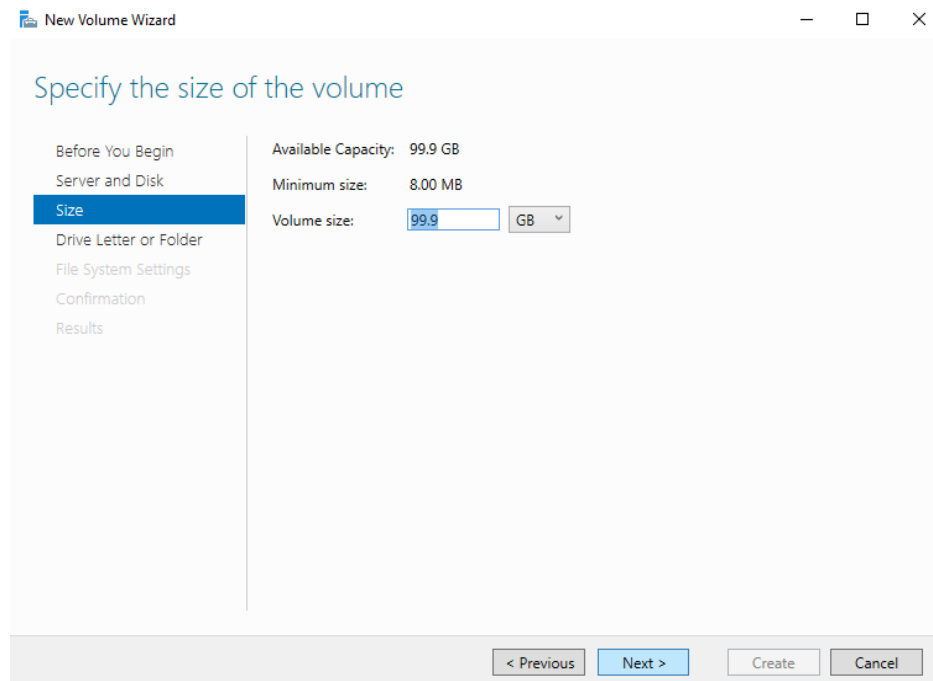
Step 10 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The **Assign to a drive letter or folder** page is displayed.

Figure 4-11 Assign to a drive letter or folder
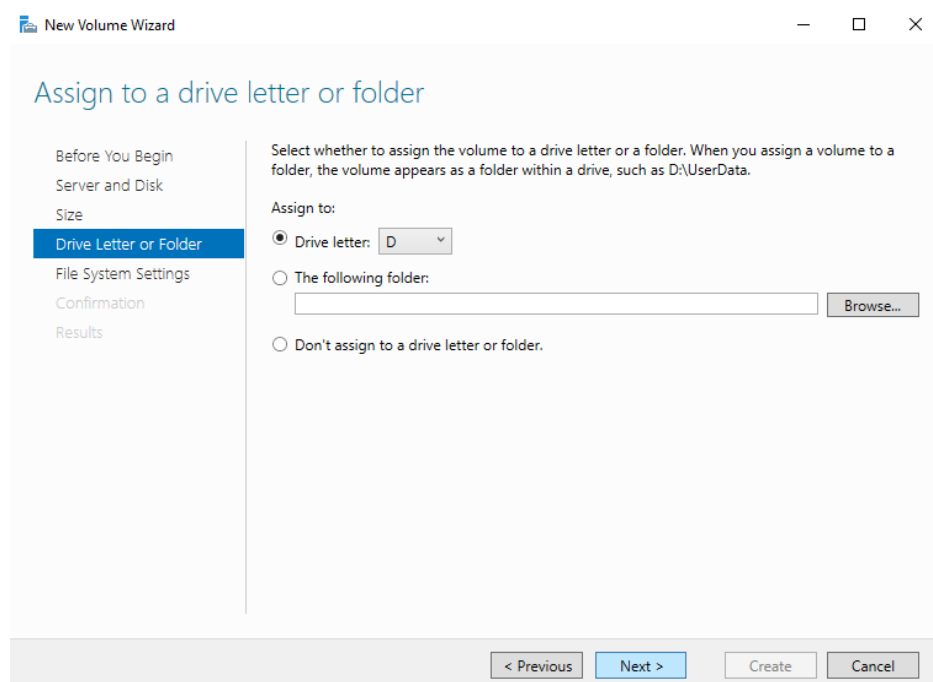
Step 11 Assign the volume to a drive letter or folder and click **Next**. The system assigns the volume to drive letter D by default. In this example, the default setting is used.

The **Select file system settings** page is displayed.

Figure 4-12 Select file system settings
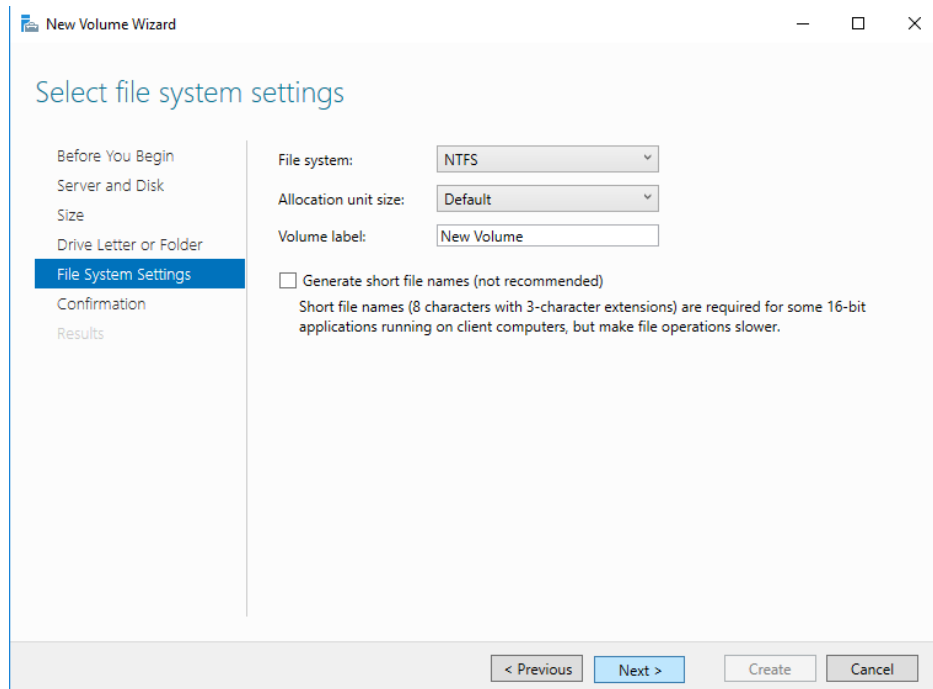
Step 12 Specify file system settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type based on the actual condition. In this example, the default setting is used.

📖 **NOTE**

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

The **Confirm selections** page is displayed.

Figure 4-13 Confirm selections

Step 13 Confirm the volume location, volume properties, and file system settings. Then, click **Create** to create a volume.

If the page shown in **Figure 4-14** is displayed, the volume is successfully created.

Figure 4-14 Completion
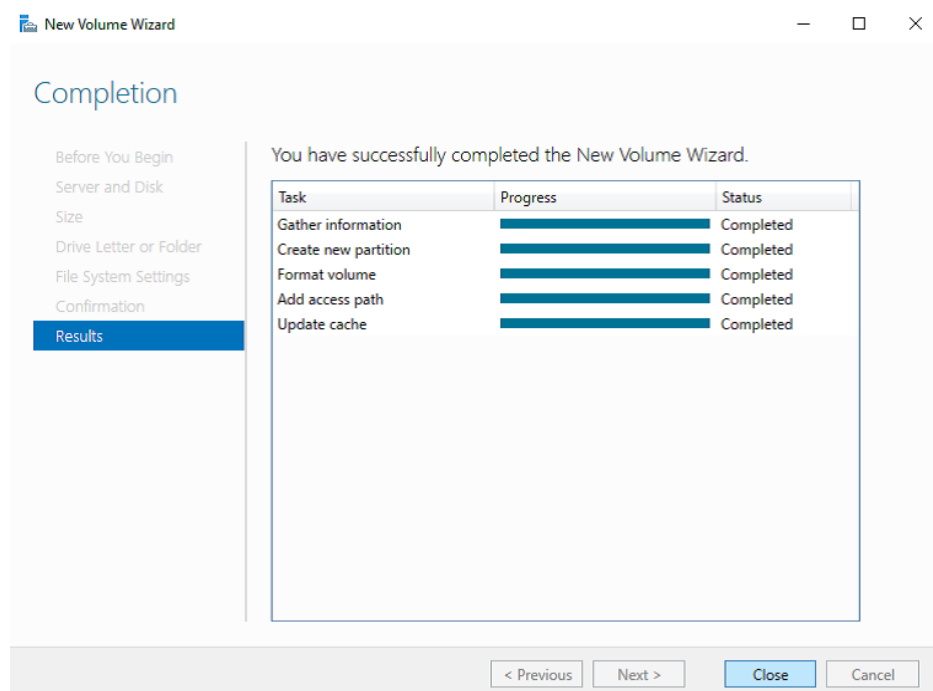
Step 14 After the volume is created, click and check whether a new volume appears in File Explorer. In this example, New Volume (D:) is the new volume.

- If New Volume (D:) appears, the disk is successfully initialized and no further action is required.

Figure 4-15 File Explorer

- If New Volume (D:) does not appear, perform the following operations to assign the volume to another drive letter or folder:

  a. Click , enter **cmd**, and press **Enter**.

  The **Administrator: Command Prompt** window is displayed.

  b. Run the **diskmgmt** command.

  The **Disk Management** page is displayed.

Figure 4-16 Disk Management (Windows 2016)

c. In the right pane of **Disk 1**, right-click and choose **Change Drive Letter and Paths**.

The **Change Drive Letter and Paths for New Volume** dialog box is displayed.

Figure 4-17 Change Drive Letter and Paths for New Volume

d. Click **Add**.

The **Add Drive Letter or Path** dialog box is displayed.

e. Select **Assign the following drive letter** to re-assign the volume to a drive letter. Then, click **OK**. Drive letter D is used in this example.

After assigning the drive letter, you can view New Volume (D:) in File Explorer.

□ **NOTE**

The drive letter selected here must be the same as that set in **Step 11**.

----End

4.3.3 Initializing a Linux Data Disk (fdisk)

Scenarios

This section uses CentOS 7.0 64-bit as an example.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is greater than 2 TB. In Linux OSs, if the GPT partition style is used, the fdisk partitioning tool cannot be used. The parted partitioning tool must be used. For details about disk partition styles, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

The method for initializing a disk varies depending on the OSs running on the BMS. This document is for reference only. For detailed operations and differences, see the product documents of the OSs running on the corresponding BMSs.

⚠ **CAUTION**

When using an EVS disk for the first time, if you have not initialized the disk, including creating partitions and file systems, the additional capacity added to the disk in a later expansion operation may not be normally used.

Prerequisites

- You have logged in to the BMS.
- A data disk has been attached to the BMS and has not been initialized.

Create Partitions and Attach a Disk

The following example shows you how to use fdisk to create a primary partition on a data disk that has been attached to the BMS. The default partitioning style is MBR and the default file system format is **ext4**. Mount the file system to **/mnt/sdc**, and configure automatic mounting upon system start.

Step 1 Run the following command to query information about the added data disk:

fdisk -l

Information similar to the following is displayed:

```
[root@bms-b656 test]# fdisk -l

Disk /dev/sda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1   *        2048     2050047     1024000   83  Linux
/dev/xvda2         2050048    22530047    10240000   83  Linux
/dev/xvda3        22530048    24578047     1024000   83  Linux
/dev/xvda4        24578048    83886079    29654016    5  Extended
/dev/xvda5        24580096    26628095     1024000   82  Linux swap / Solaris

Disk /dev/sdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

The command output shows that the BMS has two disks, system disk **/dev/sda** and data disk **/dev/sdb**.

Step 2 Run the following command to use fdisk to perform the partitioning operations for the added data disk:

fdisk *Newly added data disk*

For example, run the following command to use fdisk to perform the partitioning operations for the **/dev/sdb** data disk:

fdisk /dev/sdb

Information similar to the following is displayed:

```
[root@ecs-b656 test]# fdisk /dev/sdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0xb00005bd.
Command (m for help):
```

Step 3 Enter **n** and press **Enter** to create a new partition.

Information similar to the following is displayed:

```
Command (m for help): n
Partition type:
   p   primary (0 primary, 0 extended, 4 free)
   e   extended
```

There are two types of disk partitions:

- Choosing **p** creates a primary partition.

- Choosing **e** creates an extended partition.

Step 4 Recreate the partition with the same partition type as before. In this example a primary partition is used. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed:

```
Select (default p): p
Partition number (1-4, default 1):
```

Partition number indicates the serial number of the primary partition. The value can be **1** to **4**.

Step 5 Enter the same partition number as the partition had before and press **Enter**. Primary partition number **1** is used in this example.

Information similar to the following is displayed:

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

First sector indicates the start cylinder number. The value can be **2048** to **20971519**, and the default value is **2048**.

Step 6 Ensure that you enter the same first cylinder as the partition had before. In this example, we previously noted down **2048**, so we type in **2048** here and press **Enter**.

Information similar to the following is displayed:

```
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Last sector indicates the end cylinder number. The value can be **2048** to **20971519**, and the default value is **20971519**.

Step 7 In this example, select the default end cylinder number **20971519** and press **Enter**.

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
Using default value 20971519
Partition 1 of type Linux and of size 10 GiB is set
Command (m for help):
```

A primary partition has been created for a 10-GB data disk.

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed:

```
Command (m for help): p

Disk /dev/sdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0xb00005bd

  Device Boot      Start         End      Blocks   Id  System
```

/dev/sdb1   2048   20971519   10484736   83 Linux

Command (m for help):

Details about the **/dev/sdb1** partition are displayed.

Step 9 Enter **w** and press **Enter** to write the partition result into the partition table.

Information similar to the following is displayed:

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

Calling ioctl() to re-read partition table.
Syncing disks.

The partition is successfully created.

☐ NOTE

In case that you want to discard the changes made before, you can exit fdisk by entering **q**.

Step 10 Run the following command to synchronize the new partition table to the OS:

partprobe

Step 11 Run the following command to set the format for the file system of the newly created partition:

mkfs -t *File system format* **/dev/sdb1**

For example, run the following command to set the **ext4** file system for the **/dev/sdb1** partition:

mkfs -t ext4 /dev/sdb1

Information similar to the following is displayed:

[root@bms-b656 test]# mkfs -t ext4 /dev/sdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 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

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

The formatting takes a period of time. Observe the system running status and do not exit.

☐ NOTE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

Step 12 Run the following command to create a mount point:

mkdir *Mount point*

For example, run the following command to create the **/mnt/sdc** mount point:

mkdir /mnt/sdc

Step 13 Run the following command to mount the new partition on the mount point created in **Step 12**:

mount /dev/sdb1 *Mount point*

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

mount /dev/sdb1 /mnt/sdc

Step 14 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed:

```
[root@bms-b656 test]# df -TH
Filesystem     Type      Size  Used Avail Use% Mounted on
/dev/xvda2     xfs       11G  7.4G  3.2G  71% /
devtmpfs       devtmpfs  4.1G    0  4.1G   0% /dev
tmpfs          tmpfs     4.1G  82k  4.1G   1% /dev/shm
tmpfs          tmpfs     4.1G  9.2M  4.1G   1% /run
tmpfs          tmpfs     4.1G    0  4.1G   0% /sys/fs/cgroup
/dev/sda3      xfs       1.1G  39M  1.1G   4% /home
/dev/sda1      xfs       1.1G  131M  915M  13% /boot
/dev/sdb1      ext4      11G   38M  9.9G   1% /mnt/sdc
```

The newly created **/dev/sdb1** is mounted on **/mnt/sdc**.

----End

Set Automatic Disk Attachment Upon BMS Start

To automatically attach a disk when a BMS starts, you should not specify its partition, for example **/dev/sdb1**, in **/etc/fstab**. This is because the sequence of cloud devices may change during the server start or stop process, for example, from **/dev/sdb** to **/dev/sdc**. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach a disk at system start.

📖 **NOTE**

The universally unique identifier (UUID) is the unique character string for disk partitions in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of **/dev/sdb1**:

blkid /dev/sdb1

Information similar to the following is displayed:

```
[root@bms-b656 test]# blkid /dev/sdb1
/dev/sdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

The UUID of **/dev/sdb1** is displayed.

Step 2 Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc   ext4 defaults   0 2

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configurations and exits the vi editor.

----End

4.3.4 Initializing a Linux Data Disk (parted)

Scenarios

This section uses CentOS 7.0 64-bit as an example to describe how to initialize a data disk attached to a BMS running Linux and use parted to partition the data disk.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is greater than 2 TB. In Linux OSs, if the GPT partition style is used, the fdisk partitioning tool cannot be used. The parted partitioning tool must be used. For details about disk partition styles, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

The method for initializing a disk varies depending on the OSs running on the BMS. This document is for reference only. For detailed operations and differences, see the product documents of the OSs running on the corresponding BMSs.

⚠ **CAUTION**

When using an EVS disk for the first time, if you have not initialized the disk, including creating partitions and file systems, the additional capacity added to the disk in a later expansion operation may not be normally used.

Prerequisites

- You have logged in to the BMS.
- A data disk has been attached to the BMS and has not been initialized.

Creating Partitions and Attaching a Disk

The following example shows you how to use parted to create a partition on a new data disk that has been attached to the BMS. The default partitioning style is GPT and the default file system format is **ext4**. Mount the file system to **/mnt/sdc**, and configure automatic mounting upon system start.

Step 1 Run the following command to query information about the added data disk:

lsblk

Information similar to the following is displayed:

```
[root@bms-centos-70 linux]# 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 command output shows that the BMS has two disks, system disk **/dev/sda** and data disk **/dev/sdb**.

Step 2 Run the following command to enter parted to partition the added data disk:

parted *Added data disk*

For example, run the following command to use fdisk to perform the partitioning operations for the **/dev/sdb** data disk:

parted /dev/sdb

Information similar to the following is displayed:

```
[root@bms-centos-70 linux]# parted /dev/sdb
GNU Parted 3.1
Using /dev/sdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 3 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) p
Error: /dev/sdb: unrecognised disk label
Model: Xen Virtual Block Device (xvd)
Disk /dev/sdb: 10.7GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
```

In the command output, the **Partition Table** value is **unknown**, indicating that the disk partition style is unknown.

Step 4 Run the following command to set the disk partition style:

mklabel *Disk partition style*

For example, run the following command to set the partition style to GPT: (Disk partition styles include MBR and GPT.)

mklabel gpt

⚠ **CAUTION**

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

Step 5 Enter **p** and press **Enter** to view the disk partition style.

Information similar to the following is displayed:

```
(parted) mklabel gpt
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/sdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End  Size  File system  Name  Flags
```

Step 6 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 7 Enter **mkpart opt** *2048s 100%* and press **Enter**.

In this example, one partition is created for the added data disk. Variable *2048s* indicates the disk start capacity, and variable *100%* indicates the disk end capacity. The two values are used for reference only. You can determine the number of partitions and the partition capacity based on your service requirements.

Information similar to the following is displayed:

```
(parted) mkpart opt 2048s 100%
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? Ignore
```

If the preceding warning message is displayed, enter **Ignore** to ignore the performance warning.

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed:

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/sdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End         Size       File system  Name  Flags
1       2048s  20969471s   20967424s               opt
```

Details about the **/dev/sdb1** partition are displayed.

Step 9 Enter **q** and press **Enter** to exit parted.

Step 10 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@bms-centos-70 linux]# 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  100G  0 disk
└─sdb1 202:17   0  100G  0 part
```

In the command output, **/dev/sdb1** is the partition you created.

Step 11 Run the following command to set the format for the file system of the newly created partition:

mkfs -t *File system format* **/dev/sdb1**

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

mkfs -t ext4 /dev/sdb1

Information similar to the following is displayed:

```
[root@bms-centos-70 linux]# mkfs -t ext4 /dev/sdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2620928 blocks
131046 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677925
80 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

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

The formatting takes a period of time. Observe the system running status and do not exit.

☐ NOTE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

Step 12 Run the following command to create a mount point:

mkdir *Mount point*

For example, run the following command to create the **/mnt/sdc** mount point:

mkdir /mnt/sdc

Step 13 Run the following command to mount the new partition on the created mount point:

mount /dev/sdb1 *Mount point*

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

mount /dev/sdb1 /mnt/sdc

Step 14 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed:

```
[root@bms-centos-70 linux]# df -TH
Filesystem     Type      Size  Used Avail Use% Mounted on
/dev/sda2      xfs       39G  4.0G  35G  11% /
devtmpfs       devtmpfs 946M    0 946M   0% /dev
tmpfs          tmpfs    954M    0 954M   0% /dev/shm
tmpfs          tmpfs    954M 9.1M 945M   1% /run
tmpfs          tmpfs    954M    0 954M   0% /sys/fs/cgroup
/dev/sdb1      ext4      11G  38M 101G   1% /mnt/sdc
```

The newly created **/dev/sdb1** is mounted on **/mnt/sdc**.

----End

Set Automatic Disk Attachment Upon BMS Start

To automatically attach a disk when a BMS starts, you should not specify its partition, for example **/dev/sdb1**, in **/etc/fstab**. This is because the sequence of cloud devices may change during the server start or stop process, for example, from **/dev/sdb** to **/dev/sdc**. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach a disk at system start.

📖 **NOTE**

The universally unique identifier (UUID) is the unique character string for disk partitions in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of **/dev/sdb1**:

blkid /dev/sdb1

Information similar to the following is displayed:

```
[root@bms-b656 test]# blkid /dev/sdb1
/dev/sdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

The UUID of **/dev/sdb1** is displayed.

Step 2 Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc     ext4 defaults     0  2
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configurations and exits the vi editor.

----End

4.3.5 Initializing a Windows Data Disk Greater Than 2 TB (Windows Server 2012)

Scenarios

This section uses Windows Server 2012 R2 Standard 64bit to describe how to initialize a data disk whose capacity is greater than 2 TB. In the following operations, the capacity of the example disk is 3 TB.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is greater than 2 TB. For details about disk partition styles, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

The method for initializing a disk varies depending on the OSs running on the BMS. This document is for reference only. For detailed operations and differences, see the product documents of the OSs running on the corresponding BMSs.

⚠ CAUTION

When using an EVS disk for the first time, if you have not initialized the disk, including creating partitions and file systems, the additional capacity added to the disk in a later expansion operation may not be normally used.

Prerequisites

- You have logged in to the BMS.
- A data disk has been attached to the BMS and has not been initialized.

Procedure

Step 1 On the BMS desktop, click in the lower left corner.

The **Server Manager** window is displayed.

Figure 4-19 Server Manager (Windows 2012)

Step 2 In the upper right corner of the **Server Manager** page, choose **Tools** > **Computer Management**.

The **Computer Management** page is displayed.

Figure 4-20 Computer Management

Step 3 Choose **Storage** > **Disk Management**.

The disk list is displayed.

Figure 4-21 Disk list

Step 4 Disks are listed in the right pane. If the new disk is in the offline state, bring it online before initialize it.

In the **Disk 1** area, right-click and choose **Online** from the shortcut menu.

When the Disk 1 status changes from **Offline** to **Not Initialized**, the disk has been brought online.

Figure 4-22 Bring online succeeded (Windows 2012)
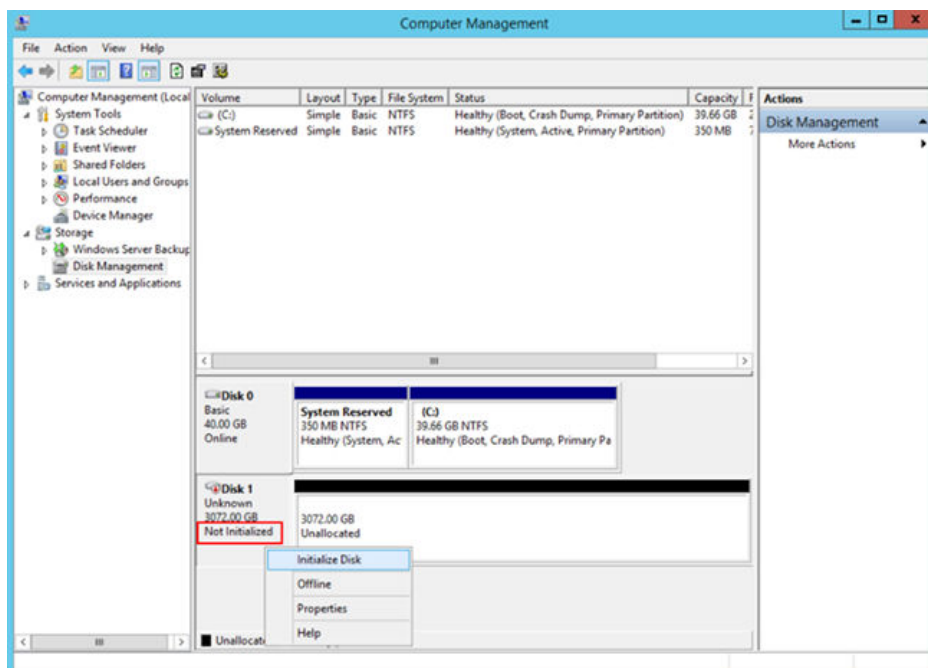
Step 5 In the **Disk 1** area, right-click and choose **Initialize Disk** from the shortcut menu.

The **Initialize Disk** dialog box is displayed.

Figure 4-23 Initialize Disk (Windows 2012)
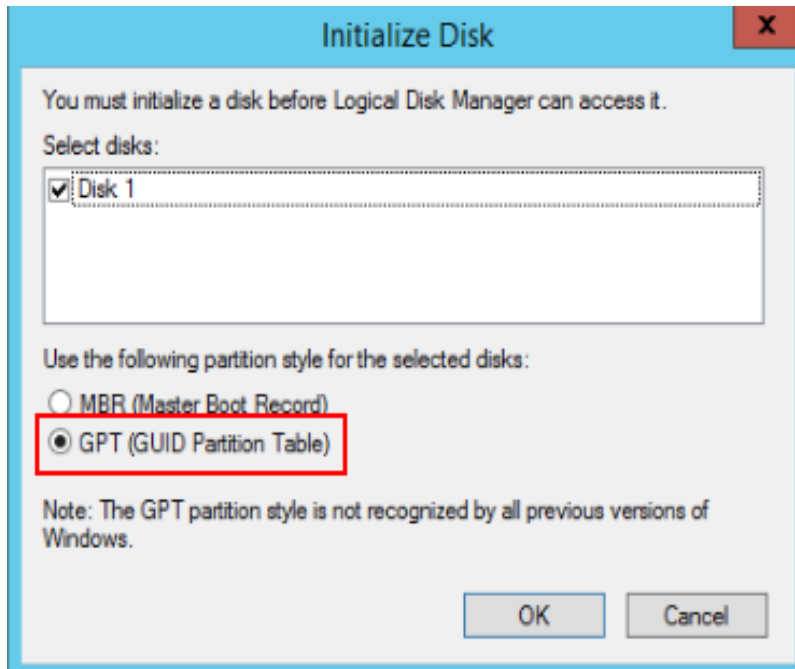
Step 6 The **Initialize Disk** dialog box displays the disk to be initialized. If the disk capacity is greater than 2 TB, select **GPT (GUID Partition Table)** and click **OK**.

The **Computer Management** page is displayed.

Figure 4-24 Computer Management (Windows 2012)
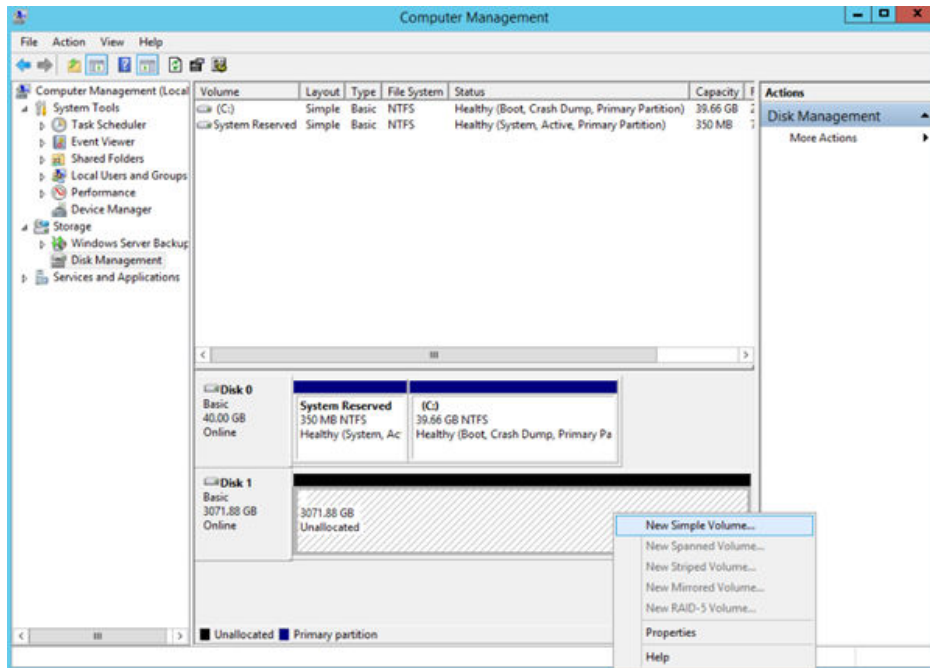
⚠ CAUTION

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

Step 7 Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.

The **New Simple Volume Wizard** window is displayed.

Figure 4-25 New Simple Volume Wizard (Windows 2012)

Step 8 Follow the prompts and click **Next**.

The **Specify Volume Size** page is displayed.

Figure 4-26 Specify Volume Size (Windows 2012)

Step 9 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The **Assign Drive Letter or Path** page is displayed.

Figure 4-27 Assign Driver Letter or Path (Windows 2012)

Step 10 Assign the volume to a drive letter or folder and click **Next**. The system assigns the volume to drive letter D by default. In this example, the default setting is used.

The **Format Partition** page is displayed.

Figure 4-28 Format Partition (Windows 2012)

Step 11 Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type based on the actual condition. In this example, the default setting is used.

The **Completing the New Simple Volume Wizard** page is displayed.

Figure 4-29 Completing the New Simple Volume Wizard (Windows 2012)

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

Step 12 Click **Finish**.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in **Figure 4-30**.

Figure 4-30 Disk initialization succeeded (Windows 2012)
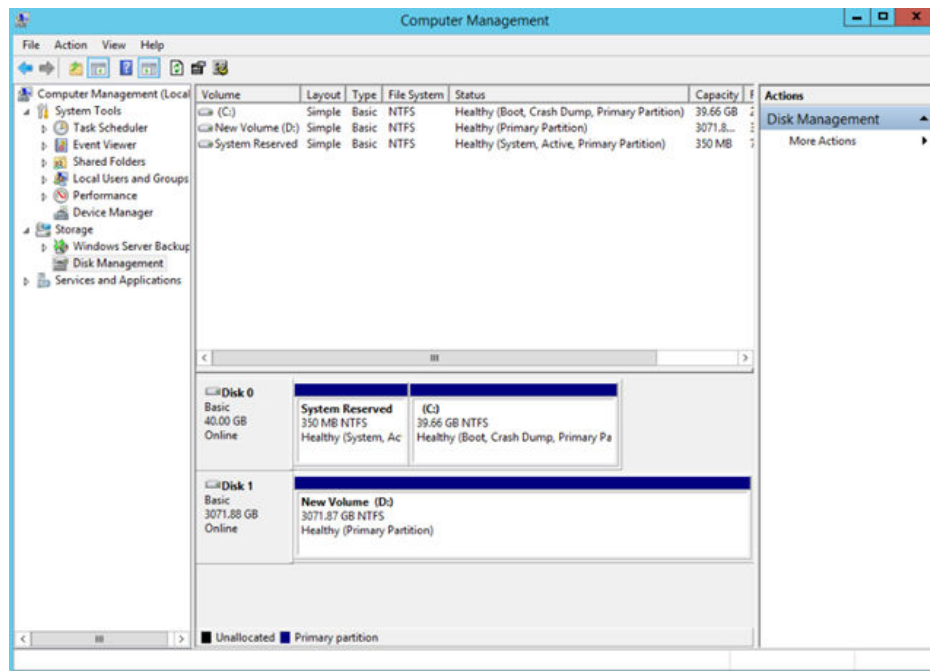
Step 13 After the volume is created, click and check whether a new volume appears in **This PC**. In this example, New Volume (D:) is the new volume.

If New Volume (D:) appears, the disk is successfully initialized and no further action is required.

Figure 4-31 This PC (Windows 2012)

----End

4.3.6 Initializing a Linux Data Disk Greater Than 2 TB (parted)

Scenarios

This section uses CentOS 7.4 64bit to describe how to use parted to initialize a data disk whose capacity is greater than 2 TB. In the following operations, the capacity of the example disk is 3 TB.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is greater than 2 TB. In Linux OSs, if the GPT partition style is used, the fdisk partitioning tool cannot be used. The parted partitioning tool must be used. For details about disk partition styles, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

The method for initializing a disk varies depending on the OSs running on the BMS. This document is for reference only. For detailed operations and differences, see the product documents of the OSs running on the corresponding BMSs.

> ⚠ **CAUTION**
>
> When using an EVS disk for the first time, if you have not initialized the disk, including creating partitions and file systems, the additional capacity added to the disk in a later expansion operation may not be normally used.

Prerequisites

- You have logged in to the BMS.
- A data disk has been attached to the BMS and has not been initialized.

Creating Partitions and Attaching a Disk

The following example shows you how to use parted to create a partition on a new data disk that has been attached to the BMS. The default partitioning style is GPT and the default file system format is **ext4**. Mount the file system to **/mnt/sdc**, and configure automatic mounting upon system start.

Step 1 Run the following command to query information about the added data disk:

lsblk

Information similar to the following is displayed:

```
[root@bms-centos74 ~]# lsblk
NAME   MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda    253:0    0  40G  0 disk
├─vda1 253:1    0   1G  0 part /boot
└─vda2 253:2    0  39G  0 part /
vdb    253:16   0   3T  0 disk
```

The command output shows that the BMS has two disks, system disk **/dev/vda** and data disk **/dev/vdb**.

Step 2 Run the following command to enter parted to partition the added data disk:

parted *Added data disk*

In this example, **/dev/vdb** is the newly added data disk.

parted /dev/vdb

Information similar to the following is displayed:

```
[root@bms-centos74 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)
```

Step 3 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) p
Error: /dev/vdb: unrecognised disk label
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
(parted)
```

In the command output, the **Partition Table** value is **unknown**, indicating that the disk partition style is unknown.

Step 4 Run the following command to set the disk partition style:

mklabel *Disk partition style*

The disk partition style can be MBR or GPT. If the disk capacity is greater than 2 TB, choose the GPT partition style.

mklabel gpt

⚠ CAUTION

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

Step 5 Enter **p** and press **Enter** to view the disk partition style.

Information similar to the following is displayed:

```
(parted) mklabel gpt
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End  Size  File system  Name  Flags

(parted)
```

Step 6 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 7 Enter **mkpart opt** *2048s 100%* and press **Enter**.

In this example, one partition is created for the added data disk. Variable *2048s* indicates the disk start capacity, and variable *100%* indicates the disk end capacity. The two values are used for reference only. You can determine the number of partitions and the partition capacity based on your service requirements.

Information similar to the following is displayed:
```
(parted) mkpart opt 2048s 100%
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? Cancel
```

If the preceding warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk. In this example, the first sector with the best disk performance is **2048s**. Therefore, the system does not display the warning message.

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed:

```
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 6442450944s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End         Size        File system  Name  Flags
 1      2048s  6442448895s  6442446848s              opt
```

Details about the **dev/vdb1** partition are displayed.

Step 9 Enter **q** and press **Enter** to exit parted.

Step 10 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@bms-centos74 ~]# lsblk
NAME   MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda    253:0    0  40G  0 disk
├─vda1 253:1    0   1G  0 part /boot
└─vda2 253:2    0  39G  0 part /
vdb    253:16   0   3T  0 disk
└─vdb1 253:17   0   3T  0 part
```

In the command output, **/dev/vdb1** is the partition you created.

Step 11 Run the following command to set the format for the file system of the newly created partition:

mkfs -t *File system format* **/dev/vdb1**

For example, run the following command to set the **ext4** file system for the **/dev/vdb1** partition:

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@bms-centos74 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
201326592 inodes, 805305856 blocks
40265292 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2952790016
24576 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
     32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
     4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
     102400000, 214990848, 512000000, 550731776, 644972544

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

The formatting takes a period of time. Observe the system running status and do not exit.

◻ NOTE

> The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

Step 12 Run the following command to create a mount point:

mkdir *Mount point*

For example, run the following command to create the **/mnt/sdc** mount point:

mkdir /mnt/sdc

Step 13 Run the following command to mount the new partition on the created mount point:

mount /dev/vdb1 *Mount point*

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

mount /dev/vdb1 /mnt/sdc

Step 14 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed:

```
[root@bms-centos74 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda2       ext4       42G  1.5G   38G   4% /
devtmpfs        devtmpfs  2.0G     0  2.0G   0% /dev
tmpfs           tmpfs     2.0G     0  2.0G   0% /dev/shm
tmpfs           tmpfs     2.0G  8.9M  2.0G   1% /run
tmpfs           tmpfs     2.0G     0  2.0G   0% /sys/fs/cgroup
/dev/vda1       ext4      1.1G  153M  801M  17% /boot
tmpfs           tmpfs     398M     0  398M   0% /run/user/0
/dev/vdb1       ext4      3.3T   93M  3.1T   1% /mnt/sdc
```

In the command output, the newly created **dev/vdb1** partition has been mounted on **/mnt/sdc**.

----End

Setting Automatic Disk Mounting at System Start

To automatically attach a disk when a BMS starts, you should not specify its partition, for example **/dev/vdb1**, in **/etc/fstab**. This is because the sequence of cloud devices may change during the BMS stop and start, for example, **/dev/vdb1** may change to **/dev/vdb2**. You are advised to use the UUID in **/etc/fstab** to automatically attach a disk at system start.

☐ NOTE

The universally unique identifier (UUID) is the unique character string for disk partitions in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of **/dev/vdb1**:

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@bms-centos74 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="bdd29fe6-9cee-4d4f-a553-9faad281f89b" TYPE="ext4" PARTLABEL="opt"
PARTUUID="c7122c92-ed14-430b-9ece-259920d5ee74"
```

In the command output, the UUID of **/dev/vdb1** is displayed.

Step 2 Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

UUID=bdd29fe6-9cee-4d4f-a553-9faad281f89b /mnt/sdc   ext4 defaults   0 2

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configurations and exits the vi editor.

----End

4.4 Detaching a Disk

Scenarios

A disk attached to a BMS can be detached.

- A disk mounted to **/dev/sda** functions as the system disk. You can only detach the system disk from a stopped BMS.

- Disks mounted to a mount point other than **/dev/sda** function as data disks and can be detached from a running or stopped BMS.

Constraints

- Detaching the system disk is a mission-critical operation. A BMS without the system disk cannot start. Exercise caution when performing this operation.

- Before detaching a data disk from a running Windows BMS, ensure that no program is reading data from or writing data to the disk. Otherwise, data will be lost.

- Before detaching a data disk from a running Linux BMS, you must log in to the BMS and run the **umount** command to cancel the association between the disk and the file system. In addition, ensure that no program is reading data from or writing data to the disk. Otherwise, detaching the disk will fail.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.
   The BMS console is displayed.
3. Click the name of the BMS from which the disk is to be detached. The page showing details of the BMS is displayed.
4. Click the **Disks** tab. Locate the row containing the disk to be detached and click **Detach**.

4.5 Expanding Disk Capacity

If a disk does not have sufficient capacity, you can expand its capacity. Both the system disk and data disk can be expanded. The maximum size of a system disk is

1 TB. For details about how to expand the disk capacity, see **Expansion Overview** in *Elastic Volume Service User Guide*.

After the capacity expansion is successful, allocate the partition for the extended space of the DSS disk.

- For details about the follow-up operations after a system disk is expanded, see **Extending Disk Partitions and File Systems (Windows)** or **Extending Partitions and File Systems for System Disks (Linux)** in *Elastic Volume Service User Guide*.

- For details about the follow-up operations after a data disk is expanded, see **Extending Disk Partitions and File Systems (Windows)** or **Extending Partitions and File Systems for Data Disks (Linux)** in *Elastic Volume Service User Guide*.

5 Key Pair and Password

5.1 Using an SSH Key Pair

Use an SSH Key Pair

To ensure system security, you are recommended to use the key authentication mode to authorize the user who attempts to log in to a BMS. Therefore, you must use an existing key pair or create a new one for remote login authentication.

- Creating a Key Pair

  If no key pair is available, create one that contains a public and a private key used for login authentication. You can use either of the following methods:

  - Create a key pair using the management console. After the creation, the public key is automatically stored in the system, and the private key is manually stored in a local directory. For details, see **Create a Key Pair on the Management Console**.

  - Use PuTTYgen to create a key pair, and save both the public and private keys to the local host. For details, see **Create a Key Pair Using PuTTYgen**. After the creation, import the key pair by following the instructions provided in **Import a Key Pair**. Then, the key pair can be used.

    ☐ NOTE

    PuTTYgen is a tool for generating public and private keys. You can obtain the tool from **https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html**.

- Using an existing key pair

  If a key pair is available locally, for example, generated using PuTTYgen, you can import the public key on the management console so that the system maintains the public key file. For details, see **Import a Key Pair**.

Create a Key Pair on the Management Console

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3.   In the navigation tree, choose **Key Pair**.

4.   On the right side of the page, click **Create Key Pair**.

5.   Enter the key name and click **OK**.

An automatically allocated key name consists of **KeyPair-** and a 4-digit random number. Change it to an easy-to-remember one, for example, **KeyPair-*xxxx*_bms**.

6.   Download the private key file. Alternatively, the system will automatically download it for you. The file name is the specified key pair name with a suffix of .pem. Securely store the private key file. In the displayed dialog box, click **OK**.

⚠ CAUTION

This is the only opportunity for you to save the private key file. Keep it secure. When creating a BMS, provide the name of the desired key pair. Each time you log in to the BMS using SSH, provide the private key.

Create a Key Pair Using PuTTYgen

Step 1   Obtain the public and private keys.

1.   Double-click **puttygen.exe**. The **PuTTY Key Generator** window is displayed.

Figure 5-1 PuTTY Key Generator
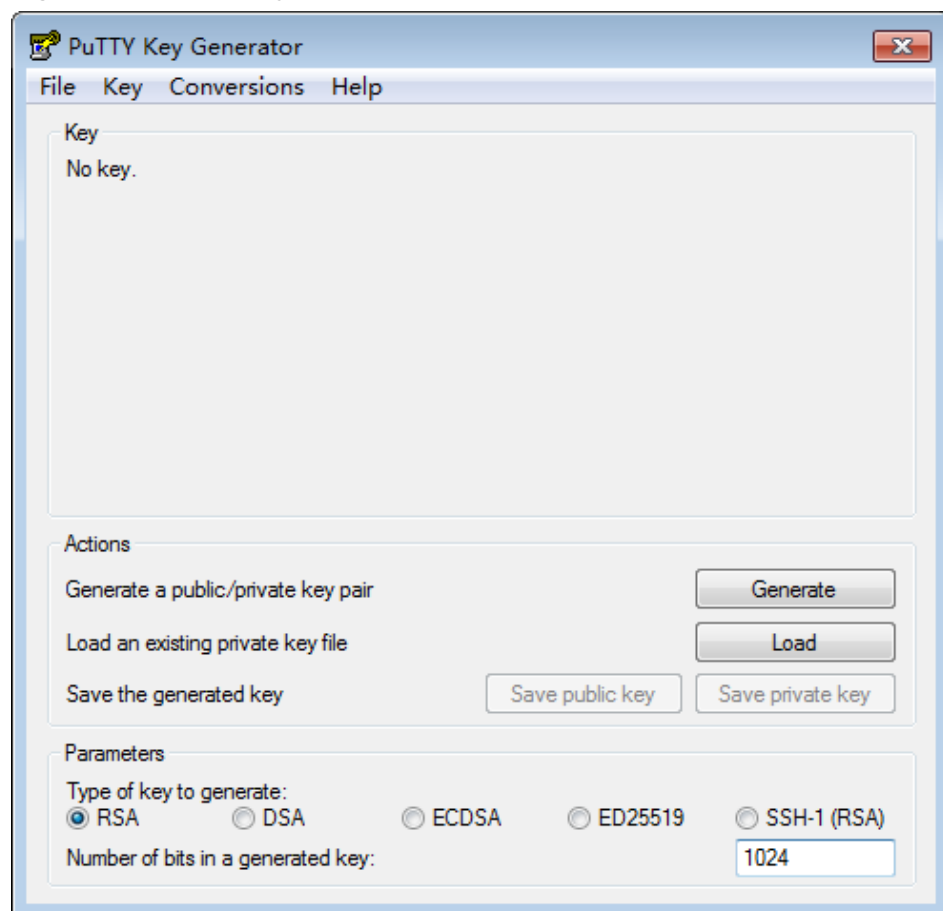
2. Click **Generate**.

    The key generator automatically generates a key pair that consists of a public key and a private key. The public key is that shown in the red box in **Figure 5-2**.

    Figure 5-2 Obtaining the public and private keys
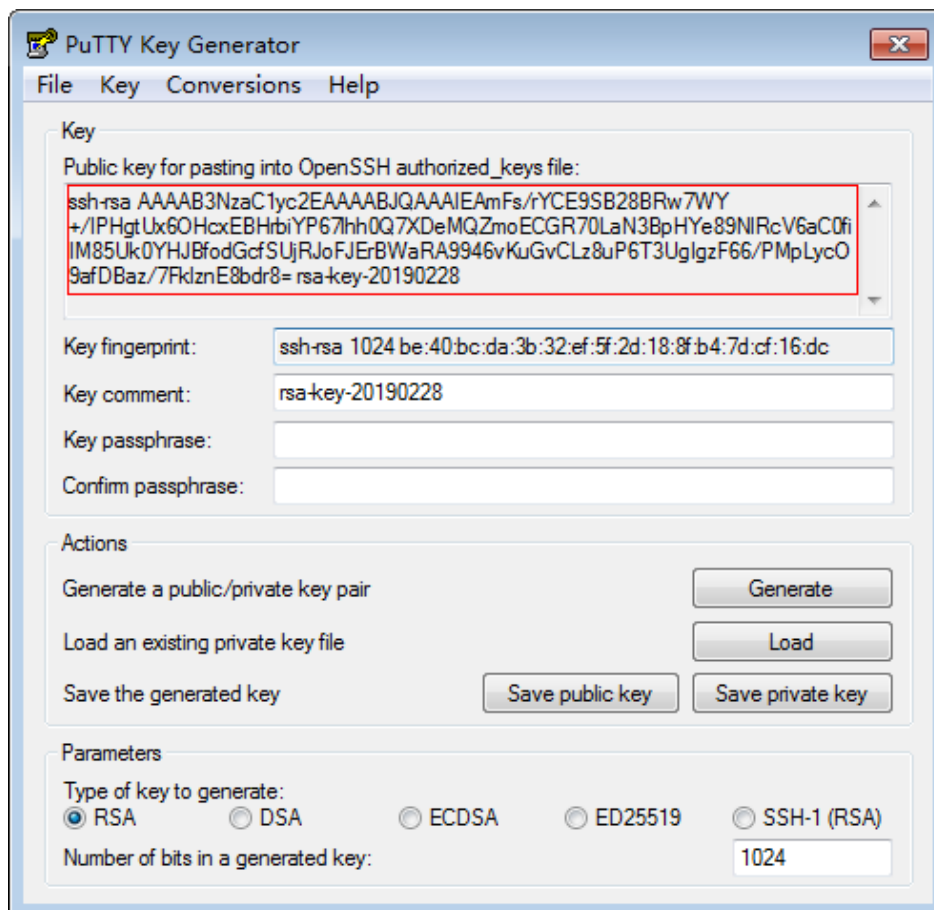
Step 2 Copy the public key content to a .txt file and save the file in a local directory.

☐ NOTE

> Do not save the public key by clicking **Save public key**. Storing a public key by clicking **Save public key** of PuTTYgen will change the format of the public key content. Such a key cannot be imported to the management console.

Step 3 Save the private key file.

The format in which to save your private key varies depending on application scenarios: To ensure BMS security, you are limited to downloading the private key only once.

●   Saving the private key in .ppk format

    When you are required to log in to a Linux BMS using PuTTY, you must use the .ppk private key. To save the private key in .ppk format, perform the following operations:

    a.   On the **PuTTY Key Generator** page, choose **File** > **Save private key**.

b. Save the private key, for example, **kp-123.ppk**, to the local PC.

- Saving the private key in .pem format

  When you are required to log in to a Linux BMS using Xshell or attempt to obtain the password for logging in to a Windows BMS, you must use the .pem private key for authentication. To save the private key in .ppk format, perform the following operations:

  a. On the **PuTTY Key Generator** page, choose **Conversions** > **Export OpenSSH key**.

  > ⚠ **CAUTION**
  >
  > If you use this private file to obtain the password for logging in to a Windows BMS, when you choose **Export OpenSSH key**, do not configure **Key passphrase**. Otherwise, obtaining the password will fail.

  b. Save the private key, for example, **kp-123.pem**, in a local directory.

Step 4 After the public key file and private key file are saved, import the public key to the system by referring to **Import a Key Pair**.

----End

Import a Key Pair

If you store a public key by clicking **Save public key** of PuTTYgen, the format of the public key content will change. Such a key cannot be imported to the management console. To resolve this issue, obtain the public key content in correct format and import the content to the management console. For details, see **What Do I Do If a Key Pair Created Using PuTTYgen Cannot Be Imported to the Management Console?**

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. In the navigation tree, choose **Key Pair**.
4. On the right side of the page, click **Import Key Pair**.

   Figure 5-3 Importing a key pair

5. Use either of the following methods to import the key pair:
   - Selecting a file

      i.   On the **Import Key Pair** page of the management console, click **Select File** and select the local public key file, for example, the .txt file saved in **Step 2**.

         ☐ NOTE

            When importing a key pair, ensure that the public key is imported. Otherwise, importing the key pair will fail.

      ii.   Click **OK**.

         After the public key is imported, you can change its name.

–   Copying the public key content

      i.   Copy the content of the public key in .txt file into the **Public Key Content** text box.

      ii.   Click **OK**.

Delete a Key Pair

If you no longer need a key pair, you can delete it. After a key pair is deleted, it cannot be restored. However, you can still use the private key saved locally to log in to the BMS, and the deleted key pair is still displayed in the BMS details.

☐ NOTE

- If your key pair has been bound to a BMS and you do not unbind the key pair from the BMS before deleting the key pair, you cannot create a key pair of the same name. When you enter this name when creating or importing a key pair, the console displays an error message indicating that the key pair already exists.

- If your key pair is not bound to any BMS or has been unbound from the BMS before it is deleted, you can create a key pair of the same name.

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. In the navigation tree, choose **Key Pair**.

4. Locate the row that contains the target key pair and click **Delete** in the **Operation** column.

5.2 (Optional) Installing the One-Click Password Reset Plug-in

Scenarios

If the password of your BMS is lost or expires and your BMS has the password reset plug-in CloudResetPwdAgent installed, you can reset the password by a few clicks.

This method is convenient and efficient. After you have created a BMS, you are advised to log in to it and install the password reset plugins.

📖 **NOTE**

> By default, the one-click password reset plug-in has been installed on the BMSs created using public images. To check whether the plug-in has been installed for the BMS, see step **Step 1** in "Install the Password Reset Plug-ins on a Linux BMS" or step **Step 1** in "Install the Password Reset Plug-ins on a Windows BMS".

Notes

1. Do not use any other password reset plug-in.
2. It is up to you to decide whether to install the CloudResetPwdAgent plugin.
3. After installing the plug-in, do not uninstall it. Otherwise, the management console may fail to reset the password.
4. After you reinstall or change the BMS OS, the one-click password resetting function will become invalid. If you still want to use this function, reinstall the CloudResetPwdAgent plug-in.
5. CloudResetPwdAgent can be automatically upgraded only if an EIP is bound to the BMS. You can also download the upgrade package and upgrade CloudResetPwdAgent manually.

Prerequisites

- The BMS must be in **Running** state.
- A Windows BMS must have larger than 600 MB remaining space and data can be written to its drive C.

  A Linux BMS must have larger than 600 MB remaining space and data can be written to its root directory.
- Ensure that DHCP is enabled in the VPC to which the BMS belongs.
- The BMS network connectivity is normal.
- Ensure that security group rules in the outbound direction meet the following requirements:
  - Protocol: TCP
  - Port Range: 80
  - Remote End: 169.254.0.0/16

  If you use the default outbound security group rule, the preceding requirements are met, and the BMS can be initialized. The default outbound security group rule is as follows:
  - Protocol: ANY
  - Port Range: ANY
  - Remote End: 0.0.0.0/16

Install the Password Reset Plug-in on a Linux BMS

Step 1 Use either of the following methods to check whether the password reset plug-in is installed on the BMS:

Method 1: Use the management console.

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

The BMS console is displayed.

3. Locate the row that contains the target BMS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list.

   – If the dialog box shown in **Figure 5-4** is displayed, the password reset plug-in is installed. No further action is required.

   Figure 5-4 Information displayed if the password reset plug-in is installed

   | Reset Password   ✕

   BMS Name   bms-m2m-maoz

   Username   root

   ★ New Password   [Enter the password]

   ★ Confirm Password [   ]

   It takes about 10 minutes to reset the password. After the password is reset, the system will
   automatically restart the BMS.
   If the new password does not take effect, reset the password again.

   [OK] [Cancel]
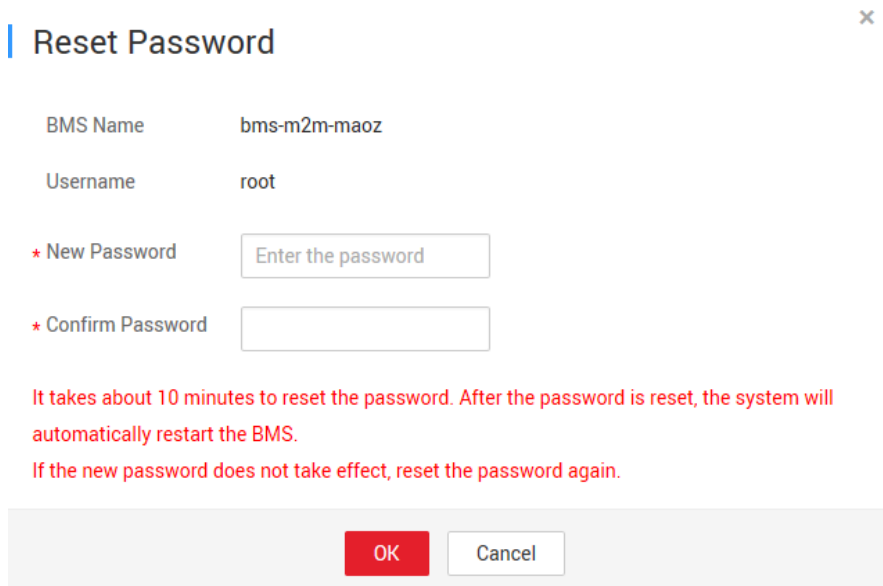
   – If the dialog box shown in **Figure 5-5** is displayed, the password reset plug-in is not installed. Perform subsequent operations to install it.

   Figure 5-5 Information displayed if the password reset plug-in is not installed

   | Reset Password   ✕

   Install the password reset plug-in and use it to reset the BMS password by following the steps
   provided in the Resetting the BMS Password.
   Resetting the BMS password adversely affects in-use services. Therefore, exercise caution when
   performing this operation. If you have any questions, contact us at ▓▓▓▓▓▓

   [OK]

Method 2: Log in to the OS.

1. Log in to the BMS as user **root**.

2. Run the following command to check whether CloudResetPwdAgent and CloudResetPwdUpdateAgent are installed:

   ls -lh /Cloud*

Figure 5-6 Checking whether the password reset plug-in is installed

```
[root@bms-1b3d test01 ~]# ls -lh /Cloud*
/CloudResetPwdUpdateAgent:
total 20K
drwxr-xr-x. 2 root root 4.0K Nov 29 03:41 bin
drwxr-xr-x. 2 root root 4.0K May 30  2018 conf
drwxr-xr-x. 3 root root 4.0K May 30  2018 depend
drwxr-xr-x. 2 root root 4.0K May 30  2018 lib
drwxr-xr-x. 2 root root 4.0K Nov 29 03:41 logs

/CloudrResetPwdAgent:
total 20K
drwxr-xr-x. 2 root root 4.0K Nov 29 03:43 bin
drwxr-xr-x. 2 root root 4.0K May 30  2018 conf
drwxr-xr-x. 2 root root 4.0K May 30  2018 depend
drwxr-xr-x. 2 root root 4.0K May 30  2018 lib
drwxr-xr-x. 2 root root 4.0K Nov 29 03:43 logs
[root@bms-1b3d_test01 ~]#
```

Check whether the information is similar to that shown in **Figure 5-6**.

- – If yes, the plugin has been installed. No further action is required.
- – If no, the plug-in has not been installed. Perform subsequent operations.

Step 2 Download the one-click password reset plug-in **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**. Customize the directory.

wget https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip

Step 3 Run the following command to decompress **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Customize the directory.

unzip -o -d *Plug-in decompression directory* **CloudResetPwdAgent.zip**

Example:

If the decompressed directory is **/home/linux/test**, run the following command:

unzip -o -d */home/linux/test* **CloudResetPwdAgent.zip**

Step 4 Install password reset plugins CloudResetPwdAgent and CloudResetPwdUpdateAgent.

1. Run the following command to open the **CloudResetPwdUpdateAgent.Linux** file:

   cd */home/linux/test***/CloudResetPwdAgent/CloudResetPwdAgent/CloudResetPwdUpdateAgent.Linux**

2. Run the following command to add the execute permission for the **setup.sh** file:

   chmod +x setup.sh

3. Run the following command to install the plugins:

sudo sh setup.sh

4. Run the following commands to check whether the installation is successful:

   service cloudResetPwdUpdateAgent status

   If the status of CoudResetPwdUpdateAgent is **running**, plugin installation is successful. Otherwise, the installation fails.

   📖 **NOTE**

   You can also check whether the password reset plug-in is installed successfully based on **Step 1**.

   If the installation fails, check whether the installation environment meets requirements and install the plugin again.
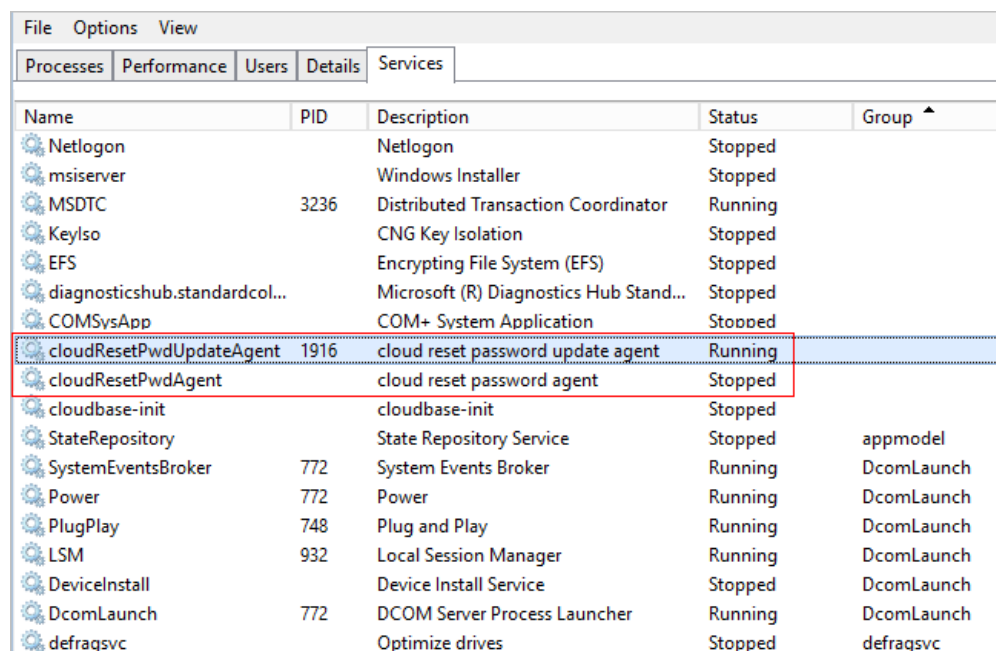
   ----End

Install the Password Reset Plug-in on a Windows BMS

Step 1 Check whether CloudResetPwdAgent and CloudResetPwdUpdateAgent have been installed on the BMS. To check this, perform the following operations:

Start the Task Manager and check whether **cloudResetPwdAgent** and **CloudResetPwdUpdateAgent** are displayed on the **Services** tab page.

Figure 5-7 Successful plug-in installation

- If yes, no further action is required.
- If no, go to the next step.

Step 2 Download the one-click password reset plug-in **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**. Customize the directory.

Download path: **https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip**

Step 3 Decompress **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Customize the directory.

Step 4 Install password reset plugins CloudResetPwdAgent and CloudResetPwdUpdateAgent.

1. Double-click **setup.bat** in both the **CloudResetPwdAgent.Windows** and **CloudResetPwdUpdateAgent.Windows** folders.

2. View the Task Manager to check whether the installation is successful.

   If you can find **cloudResetPwdAgent** and **cloudResetPwdUpdateAgent** in the Task Manager, the installation is successful. Otherwise, the installation fails.

   📖 **NOTE**

   If the installation fails, check whether the installation environment meets requirements and install the plugin again.

----End

Uninstall the Password Reset Plug-in

If you do not need the password reset function any longer, perform the following operations to uninstall the plug-in:

- **Linux BMSs**

  Log in to the BMS as user **root** and run the following commands:

  a. **sudo bash /CloudrResetPwdAgent/bin/cloudResetPwdAgent.script remove**

  b. **sudo bash /CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script remove**

  c. **sudo rm -rf /CloudrResetPwdAgent**

  d. **sudo rm -rf /CloudResetPwdUpdateAgent**

- **Windows BMSs**

  a. Log in to the BMS and go to the **C:\CloudResetPwdUpdateAgent\bin** directory.

  b. Double-click **UninstallApp-NT.bat**.

  c. Delete **C:\CloudResetPwdUpdateAgent**.

  d. Go to the **C:\CloudResetPwdAgent\bin** directory.

  e. Double-click **UninstallApp-NT.bat**.

  f. Delete **C:\CloudResetPwdAgent**.

5.3 Obtaining the Password of a Windows BMS

Scenarios

Password authentication mode is required to log in to a Windows BMS. Therefore, you must use the key file used when you created the BMS to obtain the administrator password generated when the BMS was initially installed. The administrator user is **Administrator** or the user configured using Cloudbase-Init. This password is randomly generated, offering high security.

Prerequisites

You have obtained the private key file used during BMS creation.

Procedure

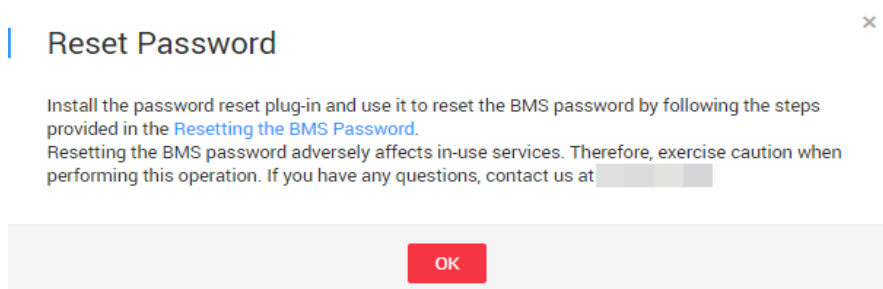1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Locate the row that contains the Windows BMS, click **More** in the **Operation** column, and select **Obtain Password**.

   Figure 5-8 Obtaining the password
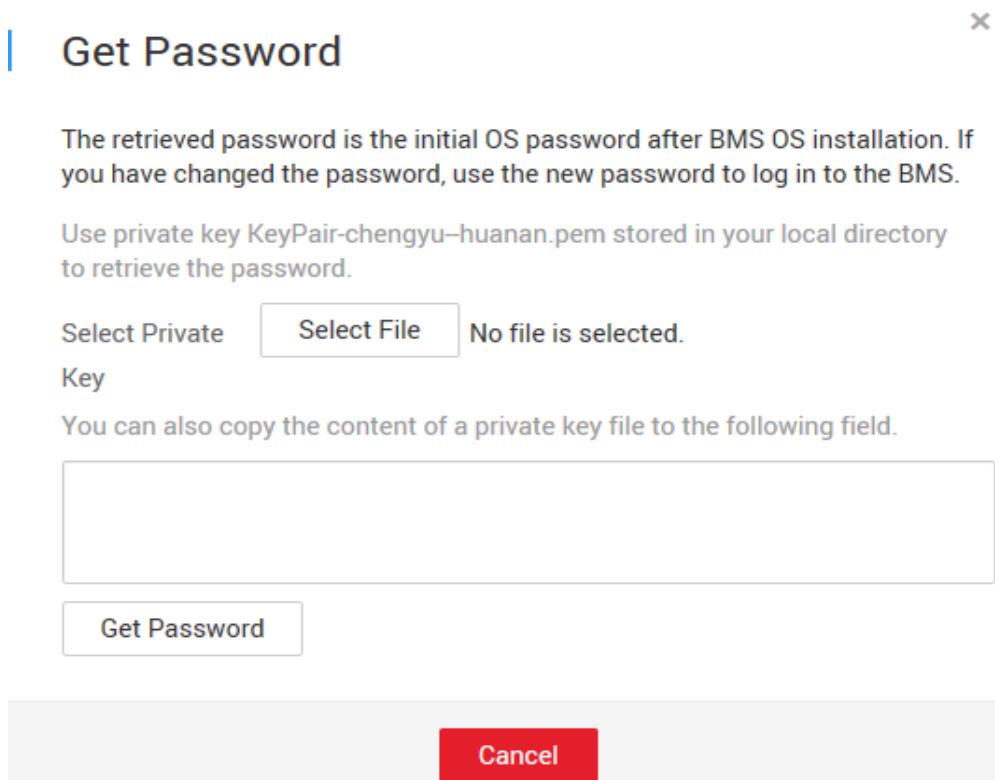
4. Use either of the following methods to obtain the password through the private key:

- – Click **Select File** and upload the private key from a local directory.
- – Copy the private key content to the text field.
5. Click **Get Password** to obtain a random password.

5.4 Deleting the Password of a Windows BMS

Scenarios

To ensure security, you are advised to delete the initial password recorded in the system.

Deleting the initial password does not affect BMS operation or login. Once deleted, the password cannot be retrieved. Before deleting a password, you are advised to record it.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Locate the target BMS in the BMS list.
4. In the **Operation** column, click **More** and select **Delete Password**.

   The following dialog box is displayed.

   Figure 5-9 Warning
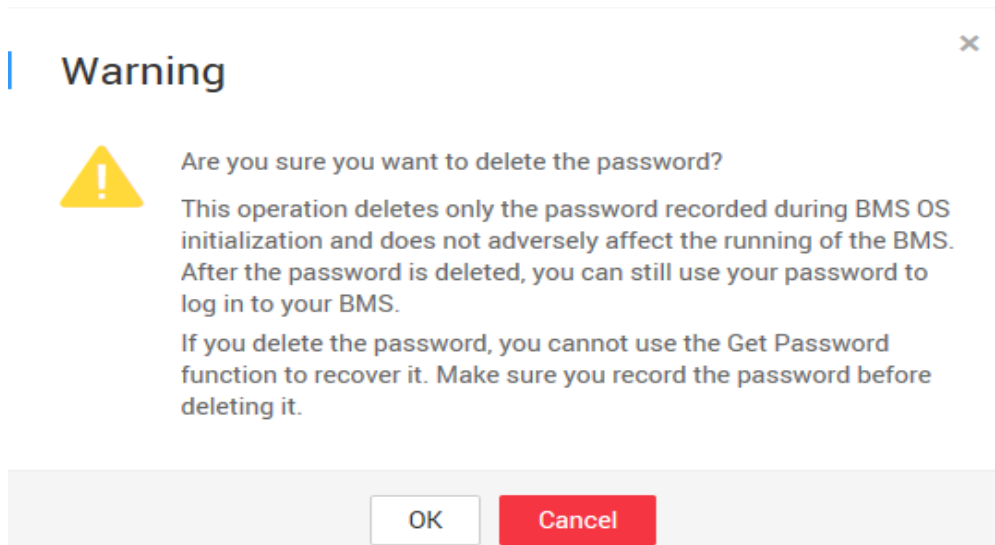
5. Click **OK** to delete the password.

6 Network

6.1 EIP

6.1.1 Overview

EIP

The Elastic IP (EIP) service provides independent public IP addresses and bandwidth for Internet access. Different from traditional static IP addresses, EIPs can be dynamically bound to or unbound from resources such as BMSs, ECSs, and NAT gateways. If a server becomes faulty, the EIP can be quickly unbound from it and bound to another healthy server to recover services.

Figure 6-1 Accessing the Internet through an EIP

Helpful Links

- **Can Multiple EIPs Be Bound to a BMS?**
- **Will I Obtain an EIP That Has Been Released?**
- **What Are the Differences Between EIPs, Private IP Addresses, and Virtual IP Addresses?**

6.1.2 Binding an EIP to a BMS

Scenarios

To allow your BMS to communicate with the Internet, bind an EIP to the BMS.

Prerequisites

An EIP is available. For details about how to apply for an EIP, see **Assigning an EIP**.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Click a BMS.

    The page showing details of the BMS is displayed.

4. Click the **EIPs** tab and then **Bind EIP**.

    The **Bind EIP** dialog box is displayed.

5. Select the EIP to be bound and click **OK**.

    📖 **NOTE**

    Only one EIP can be bound to a NIC.

6.1.3 Unbinding an EIP from a BMS

Scenarios

This section describes how to unbind an EIP from a BMS.

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

    The BMS console is displayed.

3. Click a BMS.

    The page showing details of the BMS is displayed.

4. Click the **EIPs** tab. On the displayed page, locate the target EIP and click
    Unbind. In the displayed dialog box, click **Yes**.

    📖 **NOTE**

    The EIP is still be billed after it is unbound. Release it if you do not need it any more.

6.2 VPC

6.2.1 Overview

VPC

A VPC provides a logically isolated network environment for BMSs. You can
configure EIPs, security groups, and VPNs in a VPC and use the VPC for
communication between ECSs and BMSs.

View VPC NICs

You can view the network interfaces of the VPC on the **NICs** tab page of the BMS
details page. For Linux images, you can also locate the VLAN sub-interface or
bond interface in the OS based on the allocated IP address.

Figure 6-2 Viewing NICs

Take CentOS 7.4 64-bit as an example. Log in to the OS and view the NIC configuration files **ifcfg-eth0**, **ifcfg-eth1**, **ifcfg-bond0**, **ifcfg-bond0.3030**, **ifcfg-bond0.2601**, and **ifcfg-bond0.2602** in the **/etc/sysconfig/network-scripts** directory. You need to use IP mapping to match the network.

Run the **ifconfig** command. The private IP address and MAC address of VPC NIC 1 is 192.168.0.190 and fa:16:3e:02:67:66. The private IP address and MAC address of VPC NIC 2 are 192.168.1.175 and fa:16:3e:16:45:4e. eth0 and eth1 automatically form bond0, and they have the same MAC address. In addition, it can be determined that **ifcfg-eth0**, **ifcfg-eth1**, **ifcfg-bond0**, and **ifcfg-bond0.3030** are VPC NIC configuration files.

```
[root@bms-ef79 network-scripts]# ifconfig
bond0: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST>  mtu 8888
        inet 192.168.0.190  netmask 255.255.255.0  broadcast 192.168.0.255
        inet6 fe80::f816:3eff:fe02:6766  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:02:67:66  txqueuelen 1000  (Ethernet)
        RX packets 329  bytes 105378 (102.9 KiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 328  bytes 29116 (28.4 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

bond0.2601: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 8888
        inet 192.168.5.23  netmask 255.255.255.0  broadcast 192.168.5.255
        inet6 fe80::f816:3eff:fe9d:7780  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:9d:77:80  txqueuelen 1000  (Ethernet)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 18  bytes 1068 (1.0 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

bond0.2602: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 8888
        inet 10.27.194.203  netmask 255.0.0.0  broadcast 10.255.255.255
        inet6 fe80::f816:3eff:fe5e:bbb  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:5e:0b:bb  txqueuelen 1000  (Ethernet)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 18  bytes 1068 (1.0 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

bond0.3030: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 8888
        inet 192.168.1.175  netmask 255.255.255.0  broadcast 192.168.1.255
        inet6 fe80::f816:3eff:fe16:454e  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:16:45:4e  txqueuelen 1000  (Ethernet)
        RX packets 6  bytes 880 (880.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 13  bytes 1458 (1.4 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

eth0: flags=6211<UP,BROADCAST,RUNNING,SLAVE,MULTICAST>  mtu 8888
        ether fa:16:3e:02:67:66  txqueuelen 1000  (Ethernet)
        RX packets 234  bytes 67810 (66.2 KiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 328  bytes 29116 (28.4 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

eth1: flags=6211<UP,BROADCAST,RUNNING,SLAVE,MULTICAST>  mtu 8888
        ether fa:16:3e:02:67:66  txqueuelen 1000  (Ethernet)
        RX packets 95  bytes 37568 (36.6 KiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 0  bytes 0 (0.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 127.0.0.1  netmask 255.0.0.0
        inet6 ::1  prefixlen 128  scopeid 0x10<host>
        loop  txqueuelen 1  (Local Loopback)
        RX packets 3  bytes 210 (210.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 3  bytes 210 (210.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

The following figures show the NIC and bond configuration information.

```
[root@bms-ef79 network-scripts]# cat ifcfg-eth0
USERCTL=no
MTU=8888
NM_CONTROLLED=no
BOOTPROTO=dhcp
TYPE=Ethernet

MASTER=bond0

[root@bms-ef79 network-scripts]# cat ifcfg-eth1

MTU=8888no

BOOTPROTO=dhcpno

TYPE=Ethernet

MASTER=bond0

[root@bms-ef79 network-scripts]# cat ifcfg-bond0

USERCTL=no:16:3e:02:67:66

BONDING_MASTER=yesT=1

NM_CONTROLLED=no

BONDING_OPTS="mode=1 miimon=100"

TYPE=Bondnd0

MACADDR=fa:16:3e:16:45:4eripts]# cat ifcfg-bond0.3030

PERSISTENT_DHCLIENT=1

VLAN=yesbond0

BOOTPROTO=dhcpno

TYPE=Ethernet3030
```

6.2.2 Managing Virtual IP Addresses

Virtual IP Address Overview

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.

If you want to improve service high availability and avoid single points of failure, you can use BMSs that are deployed to work in the active/standby mode or one active and multiple standby modes. These BMSs use the same virtual IP address.

Figure 6-3 Networking diagram of the HA mode

- Bind two BMSs in the same subnet to the same virtual IP address.
- Configure Keepalived for the two BMSs to work in the active/standby mode. For details about Keepalived configurations, see the common configuration methods in the industry.

📖 **NOTE**

For more information about virtual IP addresses, see **Virtual Private Cloud Service Overview**.

How Do I Bind a Virtual IP Address to a BMS?

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Click the name of the BMS to which a virtual IP address needs to be bound.

   The page showing details of the BMS is displayed.

4. Click the **NICs** tab. Then, click **Manage Virtual IP Address**.

   The page showing details of the particular VPC is displayed.

5. On the **Virtual IP Address** tab, select a desired one or click **Assign Virtual IP Address** for a new one.

6. Click **Bind to Server** in the **Operation** column and select the target BMS and the NIC to bind the virtual IP address to the NIC.

6.2.3 Setting the Source/Destination Check for a NIC

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

    The BMS console is displayed.

3. Click the name of the target BMS.

    The page showing details of the BMS is displayed.

4. Select the **NICs** tab. On the displayed page, click the icon before the NIC to be configured.

5. Enable or disable the source/destination check function.

    By default, the source/destination check function is enabled. When this function is enabled, the system checks whether source IP addresses contained in the packets sent by BMSs are correct. If the IP addresses are incorrect, the system does not allow the BMSs to send the packets. This mechanism prevents packet spoofing, thereby improving system security. If the BMS functions as a NAT server, router, or firewall, you must disable the source/destination check for the BMS.

6.3 High-Speed Network

6.3.1 Overview

High-Speed Network

A high-speed network is an internal network among BMSs and shares the same physical plane with VPC. After you create a high-speed network on the management console, the system will create a dedicated VLAN sub-interface in the BMS OS for network data communication. It uses the 10 Gbit/s port. A high-speed network has only east-west traffic and supports only communication at layer 2 because it does not support layer 3 routing.

☐ NOTE

In some regions, high-speed networks have been upgraded to enhanced high-speed networks with higher performance. For details, see **Overview**.

View High-Speed NICs

You can view the network interfaces of the high-speed network on the **NICs** tab page of the BMS details page. For Linux images, you can also locate the VLAN sub-interface or bond interface in the OS based on the allocated IP address.

Figure 6-4 Viewing high-speed NICs

Take CentOS 7.4 64-bit as an example. Log in to the OS and view the NIC configuration files **ifcfg-eth0**, **ifcfg-eth1**, **ifcfg-bond0**, **ifcfg-bond0.3441**, **ifcfg-bond0.2617**, and **ifcfg-bond0.2618** in the **/etc/sysconfig/network-scripts** directory. You need to use IP mapping to match the network.

Run the **ifconfig** command. The private IP addresses of the two high-speed NICs on the console are 192.168.5.58 and 10.34.247.26. It can be determined that **ifcfg-bond0.2617** and **ifcfg-bond0.2618** are configuration files of the high-speed NICs.

```
[root@bms-373896 network-scripts]# ifconfig
bond0: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST>  mtu 8888
        inet 192.168.0.153  netmask 255.255.255.0  broadcast 192.168.0.255
        inet6 fe80::f816:3eff:feb0:d27c  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:b0:d2:7c  txqueuelen 1000  (Ethernet)
        RX packets 8119  bytes 4222333 (4.0 MiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 459  bytes 38566 (37.6 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

bond0.2617: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 8888
        inet 192.168.5.58  netmask 255.255.255.0  broadcast 192.168.5.255
        inet6 fe80::f816:3eff:fe79:b493  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:79:b4:93  txqueuelen 1000  (Ethernet)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 18  bytes 1068 (1.0 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

bond0.2618: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 8888
        inet 10.34.247.26  netmask 255.0.0.0  broadcast 10.255.255.255
        inet6 fe80::f816:3eff:fe5f:b999  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:5f:b9:99  txqueuelen 1000  (Ethernet)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 18  bytes 1068 (1.0 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

bond0.3441: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 8888
        inet 192.168.0.49  netmask 255.255.255.0  broadcast 192.168.0.255
        inet6 fe80::f816:3eff:fe86:31f4  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:86:31:f4  txqueuelen 1000  (Ethernet)
        RX packets 219  bytes 10677 (10.4 KiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 12  bytes 1416 (1.3 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

eth0: flags=6211<UP,BROADCAST,RUNNING,SLAVE,MULTICAST>  mtu 8888
        ether fa:16:3e:b0:d2:7c  txqueuelen 1000  (Ethernet)
        RX packets 4164  bytes 2129931 (2.0 MiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 459  bytes 38566 (37.6 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

eth1: flags=6211<UP,BROADCAST,RUNNING,SLAVE,MULTICAST>  mtu 8888
        ether fa:16:3e:b0:d2:7c  txqueuelen 1000  (Ethernet)
        RX packets 3955  bytes 2092402 (1.9 MiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 0  bytes 0 (0.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 127.0.0.1  netmask 255.0.0.0
        inet6 ::1  prefixlen 128  scopeid 0x10<host>
        loop  txqueuelen 1  (Local Loopback)
        RX packets 48  bytes 2640 (2.5 KiB)
        TX packets 48  bytes 2640 (2.5 KiB) frame 0
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

The following figures show the NIC and bond configuration information.

```
[root@bms-373896 network-scripts]# cat ifcfg-bond0.2617
MACADDR=fa:16:3e:79:b4:93
USERCTL=no
PHYSDEV=bond0
VLAN=yes
IPADDR=192.168.5.58
NM_CONTROLLED=no
NETMASK=255.255.255.0
BOOTPROTO=static
DEVICE=bond0.2617
ONBOOT=yesnet
You have new mail in /var/spool/mail/root
[root@bms-373896 network-scripts]# cat ifcfg-bond0.2618
MACADDR=fa:16:3e:5f:b9:99
USERCTL=no
PHYSDEV=bond0
VLAN=yes
IPADDR=10.34.247.26
NM_CONTROLLED=no
NETMASK=255.0.0.0
BOOTPROTO=static
DEVICE=bond0.2618
TYPE=Ethernet
ONBOOT=yes
[root@bms-373896 network-scripts]#
```

6.3.2 Managing High-Speed Networks

A high-speed network is an internal network among BMSs and provides high bandwidth for connecting BMSs in the same AZ. If you want to deploy services requiring high throughput and low latency, you can create high-speed networks.

Restrictions

- When creating a BMS, the network segment used by common NICs cannot overlap with that used by high-speed NICs.

- The high-speed network does not support security groups, EIPs, DNS, VPNs, and Direct Connect connections.

- You must select different high-speed networks for different high-speed NICs of a BMS.

- After a BMS is provisioned, you cannot configure a high-speed network.

Create a High-Speed Network

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Click the **High-Speed Networks** tab and then click **Create High-Speed Network**.

4. Set the name and subnet for the high-speed network and click **OK**.

Change the Name of a High-Speed Network

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

The BMS console is displayed.

3. Click the **High-Speed Networks** tab. Locate the target high-speed network and click **Modify** in the **Operation** column.

4. Change the high-speed network name and click **OK**.

Manage Private IP Addresses

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

The BMS console is displayed.

3. Click the **High-Speed Networks** tab. Locate the target high-speed network, click **More** in the **Operation** column, and select **Manage Private IP Address** from the drop-down list.

    – To reserve a private IP address in the high-speed network for binding the IP address to a BMS during BMS creation or for other purposes, perform steps **4** to **5**.

    – To delete a private IP address, perform step **6**.

4. Click **Assign Private IP Address**.

    – If you select **Automatic Assignment**, the system automatically assigns a private IP address.

    – If you select **Manual Assignment**, you can specify a specific IP address in the high-speed network segment as the private IP address.

5. Click **OK**.

6. Locate the row that contains the target private IP address, and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK** to delete the IP address.

6.4 Enhanced High-Speed Network

6.4.1 Overview

Enhanced High-Speed Network

Figure 6-5 Enhanced high-speed network architecture

An enhanced high-speed network is a high-quality, high-speed, and low-latency internal network for BMSs to communicate with each other. It has the following features:

- Networks for high-speed internal interconnection
- Internal networks that you can customize
- A total bandwidth greater than 10 Gbit/s

Hardware and software in high-speed networks are upgraded to provide enhanced high-speed networks. **Figure 6-6** shows the architecture of the high-speed network and **Figure 6-7** shows a comparison between the architectures of the high-speed network and enhanced high-speed network.

Figure 6-6 High-speed network architecture
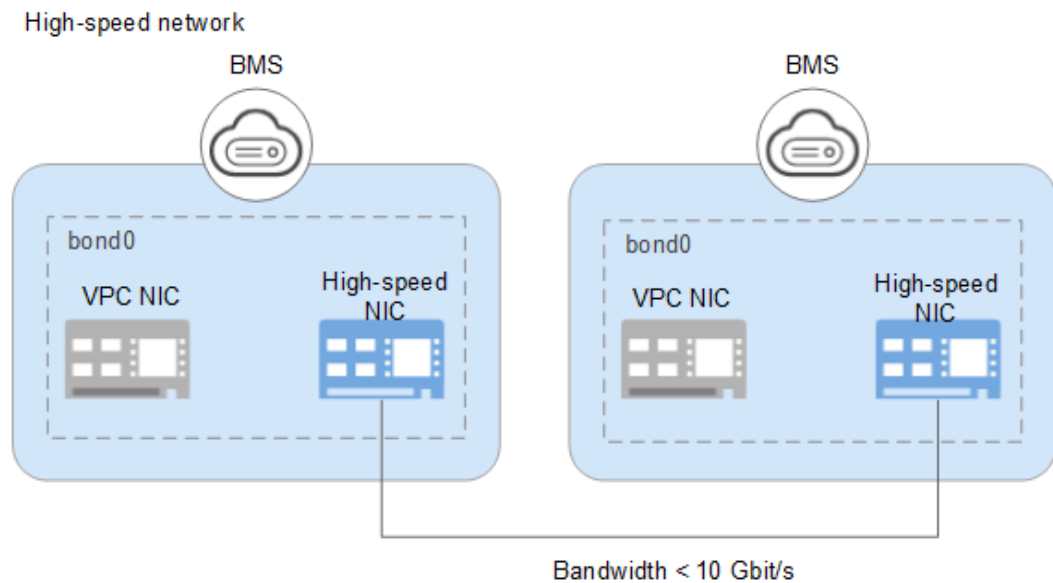
Figure 6-7 Comparison between the high-speed network and enhanced high-speed network

Compared with the high-speed network, the enhanced high-speed network has the following advantages:

- The bandwidth is increased to 10 Gbit/s or higher.
- The number of network planes can be customized and a maximum of 4000 subnets are supported.
- BMSs can be virtualized to access the Internet.

View Enhanced High-Speed NICs

You can view the network interfaces of the enhanced high-speed network on the **NICs** tab page of the BMS details page.

Figure 6-8 Viewing enhanced high-speed NICs

Application Scenarios

The enhanced high-speed NIC applies to the following scenarios:

- Scenario 1: bonding

  When bonding enhanced high-speed NICs, you can choose whether to configure VLANs based on network planning.

  - Do not configure VLANs.

    If no VLAN is required, you can configure IP addresses and subnet masks directly when bonding enhanced high-speed NICs. After the configuration is complete, enhanced high-speed NICs on the same network can communicate with each other.

  - Configure VLANs.

    If VLANs are required, you can configure VLAN sub-interfaces after bonding enhanced high-speed NICs.

- Scenario 2: no bonding

  If you use enhanced high-speed NICs directly without bonding them, you cannot configure VLANs or configure IP addresses or subnet masks. After the configuration is complete, enhanced high-speed NICs on the same network can communicate with each other.

  📖 **NOTE**

  A single enhanced high-speed NIC also supports bonding.

Configuring an Enhanced High-Speed NIC (SUSE Linux Enterprise Server 12) to **Configuring an Enhanced High-Speed NIC (Windows Server)** describe how to bond enhanced high-speed NICs in the OS. The configuration method varies depending on the OS.

6.4.2 Adding an Enhanced High-Speed NIC

Constraints

The BMS must be in **Running** state.

Add an Enhanced High-Speed NIC

To enable your BMS to communicate with resources in other PODs, add an enhanced high-speed NIC to the BMS.

📖 **NOTE**

A BMS has a maximum of two enhanced high-speed NICs and depends on the total bandwidth of the extension NICs. For example, if the total bandwidth allowed for the extension NICs is 2 x 10GE and the bandwidth of the first enhanced high-speed NIC is 2 x 10GE, you cannot add another enhanced high-speed NIC.

You can view the total bandwidth of extension NICs in the **Extended Configuration** column in **Flavor**.

- If a flavor's **Extended Configuration** contains **2*10GE** (for example, the **Extended Configuration** of flavor physical.h2.large is **1*100G IB + 2*10GE**), BMSs of this flavor has only one NIC without extension NIC, and the total bandwidth of extension NICs is 0.

- If a flavor's **Extended Configuration** contains **2 x 2*10GE** (for example, the **Extended Configuration** of flavor physical.s3.large is **2 x 2*10GE**), BMSs of this flavor has two NICs, of which one is an extension NIC, and the total bandwidth of extension NICs is 2*10GE.

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Click the name of the target BMS.

   The page showing details of the BMS is displayed.
4. Click the **NICs** tab. Then, click **Add NIC**.
5. Set the NIC type to enhanced high-speed NIC and select the bandwidth.
6. Click **OK**.

Follow-up Operations

The BMS cannot identify the newly added enhanced high-speed NIC. You must manually activate the NIC by following the instructions in sections **Configuring an Enhanced High-Speed NIC (SUSE Linux Enterprise Server 12)** to **Configuring an Enhanced High-Speed NIC (Windows Server)**.

6.4.3 Deleting an Enhanced High-Speed NIC

Constraints

The BMS must be in **Running** or **Stopped** state.

Delete an Enhanced High-Speed NIC

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Click the name of the target BMS.

   The page showing details of the BMS is displayed.

4. Click the **NICs** tab, locate the target enhanced high-speed NIC, click ⌄ to expand its details, and make a note of the MAC address.

⚲ **NOTE**

> After deleting a NIC on the console, you need to log in to the BMS OS and perform related operations to delete the device (the MAC address recorded will be used).

5. Click **Delete**.

6. Click **Yes**.

Follow-up Operations

Delete network devices by following the "Delete a NIC" part in **Configuring an Enhanced High-Speed NIC (SUSE Linux Enterprise Server 12)** to **Configuring an Enhanced High-Speed NIC (Windows Server)**.

6.4.4 Configuring an Enhanced High-Speed NIC (SUSE Linux Enterprise Server 12)

This section uses SUSE Linux Enterprise Server 12 SP3 (x86_64) as an example to describe how to configure an enhanced high-speed NIC of a BMS, including the configuration for adding and deleting a NIC.

Add a NIC

⚲ **NOTE**

> For details about how to add a NIC in other OSs, see:
> - **Add a NIC in SUSE Linux Enterprise Server 11**
> - **Add a NIC in Red Hat, CentOS, Oracle Linux, and EulerOS**
> - **Add a NIC in Ubuntu**
> - **Add a NIC in Windows Server**

1. Use a key or password to log in to the BMS as user **root**.

2. On the BMS CLI, run the following command to check the NIC information:

   ip link

   Information similar to the following is displayed:

   ```
   1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group
   default qlen 1
       link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
   mode DEFAULT group default qlen 1000
       link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
   3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
   mode DEFAULT group default qlen 1000
       link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
   4: eth2: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default
   qlen 1000
       link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
   5: eth3: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default
   qlen 1000
       link/ether 40:7d:0f:52:e3:a6 brd ff:ff:ff:ff:ff:ff
   6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
   DEFAULT group default qlen 1000
       link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
   ```

📖 **NOTE**

> eth0 and eth1 bear the VPC, and eth2 and eth3 bear the enhanced high-speed network.

3. Configure the udev rules:

   Run the following command to create the **80-persistent-net.rules** file:

   cp /etc/udev/rules.d/70-persistent-net.rules /etc/udev/rules.d/80-persistent-net.rules

   Write the NIC MAC address and name that are queried in **2** and that are not displayed in **80-persistent-net.rules** to the file. In this way, after the BMS is restarted, the NIC name and sequence will not change.

   📖 **NOTE**

   > Ensure that the NIC MAC address and name are lowercase letters.

   vim /etc/udev/rules.d/80-persistent-net.rules

   The modification result is as follows:

   ```
   SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:5d:b7:2a",
   NAME="eth0"
   SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:5d:b7:2b",
   NAME="eth1"
   SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e3:a5",
   NAME="eth2"
   SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e3:a6",
   NAME="eth3"
   ```

4. Run the following commands to create configuration files for NICs eth2 and eth3 (you can quickly create the files by copying existing NIC configuration files):

   cd /etc/sysconfig/network

   cp ifcfg-eth0 ifcfg-eth2

   cp ifcfg-eth1 ifcfg-eth3

   Run the following commands to modify the configuration files of NICs eth2 and eth3:

   vi ifcfg-eth2

   Modified configuration file of NIC eth2 is as follows.

   ```
   STARTMODE=auto
   MTU=8888
   NM_CONTROLLED=no
   BOOTPROTO=STATIC
   DEVICE=eth2
   USERCONTRL=no
   LLADDR=40:7d:0f:52:e3:a5
   TYPE=Ethernet
   ```

   📖 **NOTE**

   > In this configuration file, set **MTU** to **8888**, **BOOTPROTO** to **STATIC**, and configure **DEVICE** and **LLADDR** as required.

   vi ifcfg-eth3

   Modified configuration file of NIC eth3 is as follows:

   ```
   STARTMODE=auto
   MTU=8888
   NM_CONTROLLED=no
   BOOTPROTO=STATIC
   ```

```
DEVICE=eth3
USERCONTRL=no
LLADDR=40:7d:0f:52:e3:a6
TYPE=Ethernet
```

After the modification, save the change and exit.

5. Run the following command to bond NICs eth2 and eth3 to a NIC, for example, bond1:

Run the following commands to create the **ifcfg-bond1** file and modify the configuration file:

cp ifcfg-bond0 ifcfg-bond1

vi ifcfg-bond1

Modified configuration file of NIC bond1 is as follows.

```
BONDING_MASTER=yes
TYPE=Bond
MTU=8888
STARTMODE=auto
BONDING_MODULE_OPTS="mode=1 miimon=100"
NM_CONTROLLED=no
BOOTPROTO=STATIC
DEVICE=bond1
USERCONTRL=no
LLADDR=40:7d:0f:52:e3:a5
BONDING_SLAVE1=eth2
BONDING_SLAVE0=eth3
IPADDR=10.10.10.104
NETMASK=255.255.255.0
NETWORK=10.10.10.0
```

☐ **NOTE**

> In this configuration file, **MTU** is set to **8888**, **BONDING_MODULE_OPTS** is set to **mode=1 miimon=100**, **BOOTPROTO** is set to **STATIC**. **DEVICE**, **BONDING_SLAVE1**, **BONDING_SLAVE0**, **IPADDR**, **NETMASK**, and **NETWORK** are configured as required. **LLADDR** is set to the LLADDR value of the **BONDING_SLAVE1** NIC.

After the modification, save the change and exit.

6. Run the following command to start the added bond1 NIC:

wicked ifup bond1

7. Run the following command to query IP addresses:

ip addr show

An example is provided as follows:

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
     valid_lft forever preferred_lft forever
   inet6 ::1/128 scope host
     valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
group default qlen 1000
   link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
group default qlen 1000
   link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond1 state UP
group default qlen 1000
   link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond1 state UP
group default qlen 1000
   link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group
```

```
default qlen 1000
    link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
    inet 172.16.2.44/24 brd 172.16.2.255 scope global bond0
       valid_lft forever preferred_lft forever
    inet6 fe80::f816:ff:fe57:90c9/64 scope link
       valid_lft forever preferred_lft forever
7: bond1: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group
default qlen 1000
    link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
    inet 10.10.10.104/24 brd 10.10.10.255 scope global bond1
       valid_lft forever preferred_lft forever
    inet6 fe80::427d:fff:fe52:e3a5/64 scope link
       valid_lft forever preferred_lft forever
```

8. Repeat the preceding operations to configure other BMSs.

Delete a NIC

> **NOTE**
>
> For details about how to delete a NIC in other OSs, see:
> - **Delete a NIC in SUSE Linux Enterprise Server 11**
> - **Delete a NIC in Red Hat, CentOS, Oracle Linux, and EulerOS**
> - **Delete a NIC in Ubuntu**
> - **Delete a NIC in Windows Server**

1. Obtain the IP address of the bonded enhanced high-speed NIC to be deleted.

2. Use a key or password to log in to the BMS as user **root**.

3. Locate the bond network device and run the following command to stop and delete the device:

    wicked ifdown bond1

4. Run the following commands to delete network configuration files **/etc/sysconfig/network-scripts/ifcfg-eth2**, **/etc/sysconfig/network-scripts/ifcfg-eth3**, and **/etc/sysconfig/network-scripts/ifcfg-bond1**:

    rm -f /etc/sysconfig/network-scripts/ifcfg-eth2

    rm -f /etc/sysconfig/network-scripts/ifcfg-eth3

    rm /etc/sysconfig/network/ifcfg-bond1

6.4.5 Configuring an Enhanced High-Speed NIC (SUSE Linux Enterprise Server 11)

This section uses SUSE Linux Enterprise Server 11 SP4 as an example to describe how to configure an enhanced high-speed NIC of a BMS.

Add a NIC

1. Use a key or password to log in to the BMS as user **root**.

2. On the BMS CLI, run the following command to check the NIC information:

    ip link

    Information similar to the following is displayed.

    ```
    1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group
    default qlen 1
        link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
    ```

```
mode DEFAULT group default qlen 1000
    link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
mode DEFAULT group default qlen 1000
    link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default
qlen 1000
    link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default
qlen 1000
    link/ether 40:7d:0f:52:e3:a6 brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
```

☐ NOTE

Among the devices, eth0 and eth1 bear the VPC, and eth2 and eth3 bear the user-defined VLAN.

3. Configure the udev rules:

    Run the following command to create the **80-persistent-net.rules** file:

    cp /etc/udev/rules.d/70-persistent-net.rules /etc/udev/rules.d/80-persistent-net.rules

    Write the NIC MAC address and name that are queried in **2** and that are not displayed in **80-persistent-net.rules** to the file. In this way, after the BMS is restarted, the NIC name and sequence will not change.

    ### ☐ NOTE

    Ensure that the NIC MAC address and name are lowercase letters.

    vim /etc/udev/rules.d/80-persistent-net.rules

    The modification result is as follows:

    ```
    SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:5d:b7:2a",
    NAME="eth0"
    SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:5d:b7:2b",
    NAME="eth1"
    SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e3:a5",
    NAME="eth2"
    SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e3:a6",
    NAME="eth3"
    ```

4. Create the configuration files of NICs eth2 and eth3:

    You can copy an existing NIC configuration file and modify it to improve the creation efficiency.

    cd /etc/sysconfig/network

    cp ifcfg-eth0 ifcfg-eth2

    cp ifcfg-eth1 ifcfg-eth3

    Run the following commands to modify the configuration files of NICs eth2 and eth3:

    vi ifcfg-eth2

    Modified configuration file of NIC eth2 is as follows.

    ```
    STARTMODE=auto
    MTU=8888
    NM_CONTROLLED=no
    BOOTPROTO=STATIC
    DEVICE=eth2
    USERCONTRL=no
    ```

```
LLADDR=40:7d:0f:52:e3:a5
TYPE=Ethernet
```

◯ NOTE

> In this configuration file, set **MTU** to **8888**, **BOOTPROTO** to **STATIC**, and configure **DEVICE** and **LLADDR** as required.

vi ifcfg-eth3

Modified configuration file of NIC eth3 is as follows:

```
STARTMODE=auto
MTU=8888
NM_CONTROLLED=no
BOOTPROTO=STATIC
DEVICE=eth3
USERCONTRL=no
LLADDR=40:7d:0f:52:e3:a6
TYPE=Ethernet
```

After the modification, save the change and exit.

5. Run the following command to bond NICs eth2 and eth3 to a NIC, for example, bond1:

Run the following commands to create the **ifcfg-bond1** file and modify the configuration file:

cp ifcfg-bond0 ifcfg-bond1

vi ifcfg-bond1

Modified configuration file of NIC bond1 is as follows.

```
BONDING_MASTER=yes
TYPE=Bond
MTU=8888
STARTMODE=auto
BONDING_MODULE_OPTS="mode=1 miimon=100"
NM_CONTROLLED=no
BOOTPROTO=STATIC
DEVICE=bond1
USERCONTRL=no
LLADDR=40:7d:0f:52:e3:a5
BONDING_SLAVE1=eth2
BONDING_SLAVE0=eth3
IPADDR=10.10.10.104
NETMASK=255.255.255.0
NETWORK=10.10.10.0
```

◯ NOTE

> In this configuration file, **MTU** is set to **8888**, **BONDING_MODULE_OPTS** is set to **mode=1 miimon=100**, **BOOTPROTO** is set to **STATIC**. **DEVICE**, **BONDING_SLAVE1**, **BONDING_SLAVE0**, **IPADDR**, **NETMASK**, and **NETWORK** are configured as required. **LLADDR** is set to the LLADDR value of the **BONDING_SLAVE1** NIC.

After the modification, save the change and exit.

6. Run the following command to start the added bond1 NIC:

ifup bond1

7. Run the following command to query IP addresses:

ip addr show

An example is provided as follows:

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
      valid_lft forever preferred_lft forever
```

```
     inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
group default qlen 1000
     link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
group default qlen 1000
     link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond1 state UP
group default qlen 1000
     link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond1 state UP
group default qlen 1000
     link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group
default qlen 1000
     link/ether fa:16:00:57:90:c9 brd ff:ff:ff:ff:ff:ff
     inet 172.16.2.44/24 brd 172.16.2.255 scope global bond0
        valid_lft forever preferred_lft forever
     inet6 fe80::f816:ff:fe57:90c9/64 scope link
        valid_lft forever preferred_lft forever
7: bond1: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group
default qlen 1000
     link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
     inet 10.10.10.104/24 brd 10.10.10.255 scope global bond1
        valid_lft forever preferred_lft forever
     inet6 fe80::427d:fff:fe52:e3a5/64 scope link
        valid_lft forever preferred_lft forever
```

8. Repeat the preceding operations to configure other BMSs.

Delete a NIC

1. Obtain the IP address of the bonded enhanced high-speed NIC to be deleted.
2. Use a key or password to log in to the BMS as user **root**.
3. Locate the bond network device and run the following command to stop and delete the device:

    ifdown bond1
4. Run the following commands to delete network configuration files **/etc/sysconfig/network-scripts/ifcfg-eth2**, **/etc/sysconfig/network-scripts/ifcfg-eth3**, and **/etc/sysconfig/network-scripts/ifcfg-bond1**:

    rm -f /etc/sysconfig/network-scripts/ifcfg-eth2

    rm -f /etc/sysconfig/network-scripts/ifcfg-eth3

    rm /etc/sysconfig/network/ifcfg-bond1

6.4.6 Configuring an Enhanced High-Speed NIC (Red Hat, CentOS, Oracle Linux, and EulerOS)

This section uses CentOS 6.9 (x86_64) as an example to describe how to configure an enhanced high-speed NIC of a BMS.

☐ NOTE

The configuration methods of Red Hat, Oracle Linux, EulerOS, and CentOS are similar.

Add a NIC

Use a key or password to log in to the BMS as user **root**. Run the following command:

blkid | grep config-2

If the command output is empty, use **Method 2**. If the command output shown in the following figure is displayed, use **Method 1**.

```
[root@bms-8d3e ~]# blkid | grep config-2
/dev/sda4: UUID="2019-04-01-16-57-22-00" LABEL="config-2" TYPE="iso9660"
```

- Method 1

Step 1 Use a key or password to log in to the BMS as user **root**.

Step 2 On the BMS CLI, run the following command to check the NIC information:

ip link

Information similar to the following is displayed.

```
[root@bms-centos ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP qlen 1000
    link/ether fa:16:00:6d:80:29 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP qlen 1000
    link/ether fa:16:00:6d:80:29 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST> mtu 8888 qdisc mq state DOWN qlen 1000
    link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST> mtu 8888 qdisc mq state DOWN qlen 1000
    link/ether 40:7d:0f:52:e3:a6 brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,PROMISC,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP
    link/ether fa:16:00:6d:80:29 brd ff:ff:ff:ff:ff:ff
[root@bms-centos ~]#
```

☐ NOTE

eth0 and eth1 bear the VPC, and eth2 and eth3 bear the enhanced high-speed network.

Step 3 Run the following command to check whether the **/etc/udev/rules.d/** directory contains the **80-persistent-net.rules** file:

ll /etc/udev/rules.d/ | grep 80-persistent-net.rules

- If yes, and the file contains all NICs except **bond0** and **lo** obtained in step **Step 2** and their MAC addresses, go to step **Step 6**.
- If no, go to step **Step 4**.

Step 4 Run the following command to copy the **/etc/udev/rules.d/70-persistent-net.rules** file and name the copy as **/etc/udev/rules.d/80-persistent-net.rules**.

cp -p /etc/udev/rules.d/70-persistent-net.rules /etc/udev/rules.d/80-persistent-net.rules

☐ NOTE

If the **/etc/udev/rules.d/70-persistent-net.rules** file does not exist, create it with the content in the following format:

SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="4c:f9:5d:d9:e8:ac", NAME="eth0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="4c:f9:5d:d9:e8:ad", NAME="eth1"

Step 5 Configure the udev rules:

Write the MAC addresses and names of NICs except eth0 and eth1 obtained in step **Step 2** (those not contained in the **/etc/udev/rules.d/70-persistent-net.rules** file) to the **/etc/udev/rules.d/80-persistent-net.rules** file so that the names and sequence of NICs do not change after the BMS is restarted.

📖 **NOTE**

Ensure that NIC MAC address and name are lowercase letters.

vi /etc/udev/rules.d/80-persistent-net.rules

The modification result is as follows:

```
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:5d:b7:2a", NAME="eth0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:5d:b7:2b", NAME="eth1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e3:a5", NAME="eth2"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e3:a6", NAME="eth3"
```

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 6 Run the following commands to copy the network configuration file **/etc/sysconfig/network-scripts/ifcfg-bond0** to generate the **/etc/sysconfig/network-scripts/ifcfg-bond1** file, and copy the **/etc/sysconfig/network-scripts/ifcfg-eth0** file to generate the **/etc/sysconfig/network-scripts/ifcfg-eth2** and **/etc/sysconfig/network/ ifcfg-eth3** files:

cp -p /etc/sysconfig/network-scripts/ifcfg-bond0 /etc/sysconfig/network-scripts/ifcfg-bond1

cp -p /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-eth2

cp -p /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-eth3

Step 7 Run the following commands to edit the **/etc/sysconfig/network-scripts/ifcfg-eth2** and **/etc/sysconfig/network-scripts/ifcfg-eth3** files:

- **vi /etc/sysconfig/network-scripts/ifcfg-eth2**

  Edit the eth2 network configuration file as follows:

  ```
  USERCTL=no
  MTU=8888
  NM_CONTROLLED=no
  BOOTPROTO=static
  DEVICE=eth2
  TYPE=Ethernet
  ONBOOT=yes
  MASTER=bond1
  SLAVE=yes
  ```

  Change the value of **BOOTPROTO** to **static**, that of **DEVICE** to the network device name **eth2**, and that of **MASTER** to the port name of the enhanced high-speed NIC bond (**bond1**). Retain values of other parameters.

- **vi /etc/sysconfig/network-scripts/ifcfg-eth3**

  Edit the eth3 network configuration file as follows (similar to eth2):

  ```
  USERCTL=no
  MTU=8888
  NM_CONTROLLED=no
  BOOTPROTO=static
  DEVICE=eth3
  TYPE=Ethernet
  ONBOOT=yes
  MASTER=bond1
  SLAVE=yes
  ```

Step 8 Run the following command to edit the **/etc/sysconfig/network-scripts/ifcfg-bond1** file:

vi /etc/sysconfig/network-scripts/ifcfg-bond1

Edit the file as follows:

```
MACADDR=40:7d:0f:52:e3:a5
BONDING_MASTER=yes
USERCTL=no
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=static
BONDING_OPTS="mode=1 miimon=100"
DEVICE=bond1
TYPE=Bond
IPADDR=10.10.10.101
NETMASK=255.255.255.0
MTU=8888
```

Where,

- Change the value of **MACADDR** to the MAC address of eth2 or eth3.
- Change the value of **BOOTPROTO** to **static**.
- Change the value of **DEVICE** to **bond1**.
- Change the value of **IPADDR** to the IP address to be allocated to bond1. If the IP address planned for the enhanced high-speed network does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that BMSs communicating through the enhanced high-speed network are in the same network segment as the enhanced high-speed network. An example value is **10.10.10.101**.
- Set the value of **NETMASK** to the subnet mask of the IP address configured for enhanced high-speed network bond1.

Retain values of other parameters.

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 9 Run the following commands to enable port group bond1 of the enhanced high-speed network:

Run the following commands to start enhanced high-speed NICs eth2 and eth3:

ifup *eth2*

ifup *eth3*

ifup *bond1*

```
[root@bms-centos network-scripts]# ifup bond1
Determining if ip address 10.10.10.101 is already in use for device bond1...
```

Step 10 Perform the preceding operations to configure other BMSs.

Step 11 After all BMSs are configured, ping the IP address in the same network segment as the enhanced high-speed network of other BMSs from each BMS.

```
[root@bms-centos network-scripts]# ping 10.10.10.102 -I bond1
PING 10.10.10.102 (10.10.10.102) from 10.10.10.101 bond1: 56(84) bytes of data.
64 bytes from 10.10.10.102: icmp_seq=1 ttl=64 time=0.475 ms
64 bytes from 10.10.10.102: icmp_seq=2 ttl=64 time=0.033 ms
64 bytes from 10.10.10.102: icmp_seq=3 ttl=64 time=0.032 ms
^C
--- 10.10.10.102 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2263ms
rtt min/avg/max/mdev = 0.032/0.180/0.475/0.208 ms
```

----End

- Method 2

Step 1 Use a key or password to log in to the BMS as user **root**.

Step 2 On the BMS CLI, run the following command to check the NIC information:

ip link

Information similar to the following is displayed.

```
[root@bms-b467 ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000
    link/ether fa:16:3e:23:ec:5f brd ff:ff:ff:ff:ff:ff
    vf 0 MAC 00:00:00:00:00:00, link-state enable
    vf 1 MAC 00:00:00:00:00:00, link-state enable
    vf 2 MAC 00:00:00:00:00:00, link-state enable
    vf 3 MAC 00:00:00:00:00:00, link-state enable
    vf 4 MAC 00:00:00:00:00:00, link-state enable
    vf 5 MAC 00:00:00:00:00:00, link-state enable
    vf 6 MAC 00:00:00:00:00:00, link-state enable
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000
    link/ether fa:16:3e:29:25:45 brd ff:ff:ff:ff:ff:ff
    vf 0 MAC 00:00:00:00:00:00, link-state enable
    vf 1 MAC 00:00:00:00:00:00, link-state enable
    vf 2 MAC 00:00:00:00:00:00, link-state enable
    vf 3 MAC 00:00:00:00:00:00, link-state enable
    vf 4 MAC 00:00:00:00:00:00, link-state enable
    vf 5 MAC 00:00:00:00:00:00, link-state enable
    vf 6 MAC 00:00:00:00:00:00, link-state enable
4: eth2: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default qlen 1000
    link/ether 48:8e:ef:78:74:59 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default qlen 1000
    link/ether 48:8e:ef:78:74:5a brd ff:ff:ff:ff:ff:ff
6: eth4: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default qlen 1000
    link/ether 48:8e:ef:78:74:5b brd ff:ff:ff:ff:ff:ff
7: eth5: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default qlen 1000
    link/ether 48:8e:ef:78:74:5c brd ff:ff:ff:ff:ff:ff
8: eth6: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default qlen 1000
    link/ether 00:2e:c7:e0:b2:37 brd ff:ff:ff:ff:ff:ff
9: eth7: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default qlen 1000
    link/ether 00:2e:c7:e0:b2:38 brd ff:ff:ff:ff:ff:ff
```

NOTE

The NIC whose MAC address starts with **fa:16** is a network device that carries the VPC network, for example, eth0 and eth1. The NIC whose MAC address is that displayed in **View Enhanced High-Speed NICs** is a network device that carries the enhanced high-speed network, such as eth6 and eth7.

Step 3 Run the following commands to edit the **/etc/sysconfig/network-scripts/ifcfg-eth6** and **/etc/sysconfig/network-scripts/ifcfg-eth7** files:

- **vi /etc/sysconfig/network-scripts/ifcfg-eth6**

    Edit the eth6 network configuration file as follows:

    ```
    USERCTL=no
    MTU=8888
    NM_CONTROLLED=no
    BOOTPROTO=static
    DEVICE=eth6
    TYPE=Ethernet
    ONBOOT=yes
    MASTER=bond1
    SLAVE=yes
    ```

    Change the value of **BOOTPROTO** to **static**, that of **DEVICE** to the network device name **eth6**, and that of **MASTER** to the port name of the enhanced high-speed NIC bond (**bond1**). Retain values of other parameters.

- **vi /etc/sysconfig/network-scripts/ifcfg-eth7**

    Edit the eth7 network configuration file as follows (similar to eth6):

```
USERCTL=no
MTU=8888
NM_CONTROLLED=no
BOOTPROTO=static
DEVICE=eth7
TYPE=Ethernet
ONBOOT=yes
MASTER=bond1
SLAVE=yes
```

Step 4 Run the following command to edit the **/etc/sysconfig/network-scripts/ifcfg-bond1** file:

vi /etc/sysconfig/network-scripts/ifcfg-bond1

Edit the file as follows:

```
MACADDR=00:2e:c7:e0:b2:37
BONDING_MASTER=yes
USERCTL=no
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=static
BONDING_OPTS="mode=1 miimon=100"
DEVICE=bond1
TYPE=Bond
IPADDR=10.10.10.101
NETMASK=255.255.255.0
MTU=8888
```

Where,

- Change the value of **MACADDR** to the MAC address of eth6 or eth7.

- Change the value of **BOOTPROTO** to **static**.

- Change the value of **DEVICE** to **bond1**.

- Change the value of **IPADDR** to the IP address to be allocated to bond1. If the IP address planned for the enhanced high-speed network does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that BMSs communicating through the enhanced high-speed network are in the same network segment as the enhanced high-speed network. An example value is **10.10.10.101**.

- Set the value of **NETMASK** to the subnet mask of the IP address configured for enhanced high-speed network bond1.

Retain values of other parameters.

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 5 Run the following commands to enable port group bond1 of the enhanced high-speed network:

Run the following commands to start enhanced high-speed NICs eth6 and eth7:

ifup *eth6*

ifup *eth7*

ifup *bond1*

```
[root@bms-centos network-scripts]# ifup bond1
Determining if ip address 10.10.10.101 is already in use for device bond1...
```

Step 6 Perform the preceding operations to configure other BMSs.

Step 7 After all BMSs are configured, ping the IP address in the same network segment as the enhanced high-speed network of other BMSs from each BMS.

```
[root@bms-centos network-scripts]# ping 10.10.10.102 -I bond1
PING 10.10.10.102 (10.10.10.102) from 10.10.10.101 bond1: 56(84) bytes of data.
64 bytes from 10.10.10.102: icmp_seq=1 ttl=64 time=0.475 ms
64 bytes from 10.10.10.102: icmp_seq=2 ttl=64 time=0.033 ms
64 bytes from 10.10.10.102: icmp_seq=3 ttl=64 time=0.032 ms
^C
--- 10.10.10.102 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2263ms
rtt min/avg/max/mdev = 0.032/0.180/0.475/0.208 ms
```

----End

To configure a VLAN, perform the following steps:

Step 1 Configure the corresponding VLAN sub-interfaces based on the VLAN to be configured. Assuming that the VLAN ID is 316, run the following command to edit the **/etc/sysconfig/network-scripts/ifcfg-bond1.316** file:

vi /etc/sysconfig/network-scripts/ifcfg-bond1.316

Edit the file as follows:

```
USERCTL=no
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=static
DEVICE=bond1.316
TYPE=Ethernet
IPADDR=10.10.0.101
NETMASK=255.255.255.0
VLAN=yes
PHYSDEV=bond1
```

Where,

- Change the value of **DEVICE** to the name of the new bond sub-interface.

- Change the value of **IPADDR** to the IP address to be allocated to bond1.316. If the IP address planned for the VLAN sub-interface of the enhanced high-speed NIC does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that the BMSs communicating with each other through the VLAN sub-interface of the enhanced high-speed NIC are in the same network segment as the VLAN sub-interface of the enhanced high-speed NIC. An example value is **10.10.0.101**.

- Set the value of **NETMASK** to the subnet mask of the IP address configured for enhanced high-speed NIC bond1.316.

Retain values of other parameters.

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 2 After all BMSs are configured, ping the IP address in the same network segment as the enhanced high-speed network VLAN sub-interface of other BMSs from each BMS.

```
[root@bms-centos ~]# ping 10.10.0.102 -I bond1.316
PING 10.10.0.102 (10.10.0.102) from 10.10.0.101 bond1.316: 56(84) bytes of data.
64 bytes from 10.10.0.102: icmp_seq=1 ttl=64 time=0.681 ms
64 bytes from 10.10.0.102: icmp_seq=2 ttl=64 time=0.035 ms
64 bytes from 10.10.0.102: icmp_seq=3 ttl=64 time=0.031 ms
64 bytes from 10.10.0.102: icmp_seq=4 ttl=64 time=0.030 ms
^C
--- 10.10.0.102 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3342ms
rtt min/avg/max/mdev = 0.030/0.194/0.681/0.281 ms
```

----End

Delete a NIC

1. Obtain the IP address of the bonded enhanced high-speed NIC to be deleted.

2. Use a key or password to log in to the BMS as user **root**.

3. Locate the bond network device and run the following command to stop and delete the device: If the bond has VLAN sub-interfaces, they will be automatically deleted.

```
[root@bms-centos ~]# ifdown eth2
[root@bms-centos ~]# ifdown eth3
[root@bms-centos ~]# ifdown bond1
[root@bms-centos ~]# ip link delete bond1
[root@bms-centos ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
qlen 1000
    link/ether fa:16:00:6d:80:29 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP
qlen 1000
    link/ether fa:16:00:6d:80:29 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST> mtu 8888 qdisc mq state DOWN qlen 1000
    link/ether 40:7d:0f:52:e3:a5 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST> mtu 8888 qdisc mq state DOWN qlen 1000
    link/ether 40:7d:0f:52:e3:a6 brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,PROMISC,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state
UP
    link/ether fa:16:00:6d:80:29 brd ff:ff:ff:ff:ff:ff
```

4. Run the following commands to delete network configuration files **/etc/sysconfig/network-scripts/ifcfg-eth2**, **/etc/sysconfig/network-scripts/ifcfg-eth3**, and **/etc/sysconfig/network-scripts/ifcfg-bond1**:

   rm -f /etc/sysconfig/network-scripts/ifcfg-eth2

   rm -f /etc/sysconfig/network-scripts/ifcfg-eth3

   rm -f /etc/sysconfig/network-scripts/ifcfg-bond1

   If a VLAN sub-interface exists, delete network configuration file **/etc/sysconfig/network-scripts/ifcfg-bond1.***vlan*, where *vlan* indicates the VLAN ID of the VLAN sub-interface, for example, **316**.

   rm -f /etc/sysconfig/network-scripts/ifcfg-bond1.*316*

6.4.7 Configuring an Enhanced High-Speed NIC (Ubuntu)

This section uses Ubuntu 16.04 LTS (Xenial Xerus x86_64) as an example to describe how to bond enhanced high-speed NICs of a BMS.

☐ **NOTE**

The configuration methods of other Ubuntu OSs are similar to that of Ubuntu 16.04 LTS (Xenial Xerus x86_64).

Add a NIC

Step 1 Use a key or password to log in to the BMS as user **root**.

Step 2 On the BMS CLI, run the following command to check the NIC information:

ip link

Information similar to the following is displayed:

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
     valid_lft forever preferred_lft forever
   inet6 ::1/128 scope host
     valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP group
default qlen 1000
   link/ether fa:16:00:9b:91:c3 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP group
default qlen 1000
   link/ether fa:16:00:9b:91:c3 brd ff:ff:ff:ff:ff:ff
4: p5p1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
   link/ether 40:7d:0f:52:e4:1d brd ff:ff:ff:ff:ff:ff
5: p5p2: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
   link/ether 40:7d:0f:52:e4:1e brd ff:ff:ff:ff:ff:ff
6: p4p1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
   link/ether 40:7d:0f:52:e3:a9 brd ff:ff:ff:ff:ff:ff
7: p4p2: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
   link/ether 40:7d:0f:52:e3:aa brd ff:ff:ff:ff:ff:ff
8: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group
default qlen 1000
   link/ether fa:16:00:9b:91:c3 brd ff:ff:ff:ff:ff:ff
   inet 192.168.254.85/24 brd 192.168.254.255 scope global bond0
     valid_lft forever preferred_lft forever
   inet6 fe80::f816:ff:fe9b:91c3/64 scope link
     valid_lft forever preferred_lft forever
9: bond0.3157@bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group
default qlen 1000
   link/ether fa:16:00:9c:1e:79 brd ff:ff:ff:ff:ff:ff
   inet 192.168.100.14/24 brd 192.168.100.255 scope global bond0.3157
     valid_lft forever preferred_lft forever
   inet6 fe80::f816:ff:fe9c:1e79/64 scope link
     valid_lft forever preferred_lft forever
10: bond0.3159@bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group
default qlen 1000
   link/ether fa:16:00:0a:2e:8e brd ff:ff:ff:ff:ff:ff
   inet 192.168.101.153/24 brd 192.168.101.255 scope global bond0.3159
     valid_lft forever preferred_lft forever
   inet6 fe80::f816:ff:fe0a:2e8e/64 scope link
     valid_lft forever preferred_lft forever
```

☐ NOTE

eth0 and eth1 bear the VPC, and p5p1, p5p2, p4p1, and p4p2 bear the enhanced high-speed network. The following operations describe how to bond enhanced high-speed NICs p4p1 and p4p2.

Step 3 Run the following command to check whether the **/etc/udev/rules.d/** directory contains the **80-persistent-net.rules** file:

ll /etc/udev/rules.d/ | grep 80-persistent-net.rules

- If yes, and the file contains all NICs except **bond0** and **lo** obtained in step **Step 2** and their MAC addresses, go to step **Step 6**.
- If no, go to step **Step 4**.

Step 4 Run the following command to copy the **/etc/udev/rules.d/70-persistent-net.rules** file and name the copy as **/etc/udev/rules.d/80-persistent-net.rules**.

cp -p /etc/udev/rules.d/70-persistent-net.rules /etc/udev/rules.d/80-persistent-net.rules

Step 5 Configure the udev rules:

Add the NICs and their MAC addresses obtained in step **Step 2**, except **lo**, **eth0**, **eth1**, and **bond0**, to the **/etc/udev/rules.d/80-persistent-net.rules** file. This ensures that the names and sequence of NICs will not change after the BMS is restarted.

☐ NOTE

Ensure that NIC MAC address and name are lowercase letters.

vim /etc/udev/rules.d/80-persistent-net.rules

The modification result is as follows:

```
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:5d:b6:fc", NAME="eth0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:5d:b6:fd", NAME="eth1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e4:1d", NAME="p5p1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e4:1e", NAME="p5p2"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e3:a9", NAME="p4p1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:52:e3:aa", NAME="p4p2"
```

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 6 Run the following command to copy the **/etc/network/interfaces.d/50-cloud-init.cfg** file to generate the **/etc/network/interfaces.d/60-cloud-init.cfg** file:

cp -p /etc/network/interfaces.d/50-cloud-init.cfg /etc/network/interfaces.d/60-cloud-init.cfg

☐ NOTE

If the **/etc/network/interfaces.d/50-cloud-init.cfg** file does not exist, copy the **/etc/network/interfaces** file and run the following commands:

mkdir /etc/network/interfaces.d

cp -p /etc/network/interfaces /etc/network/interfaces.d/60-cloud-init.cfg

Step 7 Run the following command to edit the **/etc/network/interfaces.d/60-cloud-init.cfg** file of devices **p4p1** and **p4p2**:

vim /etc/network/interfaces.d/60-cloud-init.cfg

Edit the file as follows:

```
auto p4p1
iface p4p1 inet manual
bond_mode 1
bond-master bond1
bond_miimon 100
mtu 8888

auto p4p2
iface p4p2 inet manual
bond_mode 1
bond-master bond1
bond_miimon 100
mtu 8888
```

```
auto bond1
iface bond1 inet static
bond_miimon 100
bond-slaves none
bond_mode 1
address 10.10.10.103
netmask 255.255.255.0
hwaddress 40:7d:0f:52:e3:a9
mtu 8888
```

Parameters are as follows:

- **p4p1** and **p4p2** are the names of the NICs that carry the enhanced high-speed network.

- **hwaddress** is the MAC address of p4p1.

- Change the value of **address** to the IP address allocated to enhanced high-speed network bond1. If the IP address planned for the enhanced high-speed network does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that BMSs communicating through the enhanced high-speed network are in the same network segment as the enhanced high-speed network.

- Set the value of **netmask** to the subnet mask of the IP address configured for enhanced high-speed network bond1.

Set values of other parameters. For example, set **mtu** to **8888**, **bond_miimon** to **100**, and **bond_mode** to **1**.

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 8 Run the following command to enable the bond NIC:

ifup *p4p1*

ifup *p4p2*

◻ NOTE

> **p4p1** and **p4p2** are the NICs bearing the enhanced high-speed network.

Step 9 Run the following commands to check the NIC device status and whether the **bond1** configuration file takes effect:

ip link

```
root@bms-ubuntu:~# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group default qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode DEFAULT group default qlen 1000
    link/ether fa:16:00:9b:91:c3 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode DEFAULT group default qlen 1000
    link/ether fa:16:00:9b:91:c3 brd ff:ff:ff:ff:ff:ff
4: p5p1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default qlen 1000
    link/ether 40:7d:0f:52:e4:1d brd ff:ff:ff:ff:ff:ff
5: p5p2: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default qlen 1000
    link/ether 40:7d:0f:52:e4:1e brd ff:ff:ff:ff:ff:ff
6: p4p1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond1 state UP mode DEFAULT group default qlen 1000
    link/ether 40:7d:0f:52:e3:a9 brd ff:ff:ff:ff:ff:ff
7: p4p2: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond1 state UP mode DEFAULT group default qlen 1000
    link/ether 40:7d:0f:52:e3:a9 brd ff:ff:ff:ff:ff:ff
8: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode DEFAULT group default qlen 1000
    link/ether fa:16:00:9b:91:c3 brd ff:ff:ff:ff:ff:ff
12: bond1: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode DEFAULT group default qlen 1000
    link/ether 40:7d:0f:52:e3:a9 brd ff:ff:ff:ff:ff:ff
13: bond1.316@bond1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode DEFAULT group default qlen 1000
    link/ether 40:7d:0f:52:e3:a9 brd ff:ff:ff:ff:ff:ff
```

ifconfig

```
root@bms-ubuntu:~# ifconfig
bond0     Link encap:Ethernet  HWaddr fa:16:00:9b:91:c3
          inet addr:192.168.254.85  Bcast:192.168.254.255  Mask:255.255.255.0
          inet6 addr: fe80::f816:ff:fe9b:91c3/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
          RX packets:6079 errors:0 dropped:1410 overruns:0 frame:0
          TX packets:3470 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1241961 (1.2 MB)  TX bytes:801316 (801.3 KB)

bond1     Link encap:Ethernet  HWaddr 40:7d:0f:52:e3:a9
          inet addr:10.10.10.103  Bcast:10.10.10.255  Mask:255.255.255.0
          inet6 addr: fe80::427d:fff:fe52:e3a9/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
          RX packets:1285 errors:0 dropped:642 overruns:0 frame:0
          TX packets:707 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:78202 (78.2 KB)  TX bytes:32534 (32.5 KB)

bond1.316 Link encap:Ethernet  HWaddr 40:7d:0f:52:e3:a9
          inet addr:10.10.0.103  Bcast:10.10.0.255  Mask:255.255.255.0
          inet6 addr: fe80::427d:fff:fe52:e3a9/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:8888  Metric:1
          RX packets:37 errors:0 dropped:0 overruns:0 frame:0
          TX packets:55 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:2804 (2.8 KB)  TX bytes:4290 (4.2 KB)

eth0      Link encap:Ethernet  HWaddr fa:16:00:9b:91:c3
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
          RX packets:1443 errors:0 dropped:1410 overruns:0 frame:0
          TX packets:715 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:359890 (359.8 KB)  TX bytes:242442 (242.4 KB)

eth1      Link encap:Ethernet  HWaddr fa:16:00:9b:91:c3
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
          RX packets:4669 errors:0 dropped:0 overruns:0 frame:0
          TX packets:2788 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:892139 (892.1 KB)  TX bytes:568072 (568.0 KB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:54 errors:0 dropped:0 overruns:0 frame:0
          TX packets:54 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1
          RX bytes:6048 (6.0 KB)  TX bytes:6048 (6.0 KB)

p4p1      Link encap:Ethernet  HWaddr 40:7d:0f:52:e3:a9
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
          RX packets:643 errors:0 dropped:0 overruns:0 frame:0
          TX packets:738 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:39682 (39.6 KB)  TX bytes:34192 (34.1 KB)

p4p2      Link encap:Ethernet  HWaddr 40:7d:0f:52:e3:a9
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
          RX packets:663 errors:0 dropped:663 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:39780 (39.7 KB)  TX bytes:0 (0.0 B)
```

Step 10 Perform the preceding operations to configure other BMSs.

Step 11 After all BMSs are configured, ping the IP address in the same network segment as the enhanced high-speed network of other BMSs from each BMS.

For example, run the **ping 10.10.10.102** command. The command output is as follows:

```
[root@bms-ubuntu ~]# ping 10.10.10.102 -I bond1
PING 10.10.10.102 (10.10.10.102) from 10.10.10.103 bond1: 56(84) bytes of data.
64 bytes from 10.10.10.102: icmp_seq=1 ttl=64 time=0.681 ms
64 bytes from 10.10.10.102: icmp_seq=2 ttl=64 time=0.035 ms
64 bytes from 10.10.10.102: icmp_seq=3 ttl=64 time=0.031 ms
64 bytes from 10.10.10.102: icmp_seq=4 ttl=64 time=0.030 ms
^C
```

```
--- 10.10.10.102 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3342ms
```

----End

To configure a VLAN, perform the following steps:

Step 1 Configure the corresponding VLAN sub-interfaces based on the VLAN to be configured. Assuming that the VLAN ID is 316, run the following command to edit the **/etc/network/interfaces.d/60-cloud-init.cfg** file:

vim /etc/network/interfaces.d/60-cloud-init.cfg

Edit the file as follows:

```
auto p4p1
iface p4p1 inet manual
bond_mode 1
bond-master bond1
bond_miimon 100
mtu 8888

auto p4p2
iface p4p2 inet manual
bond_mode 1
bond-master bond1
bond_miimon 100
mtu 8888

auto bond1
iface bond1 inet static
bond_miimon 100
bond-slaves none
bond_mode 1
address 10.10.10.103
netmask 255.255.255.0
hwaddress 40:7d:0f:52:e3:a9
mtu 8888

auto bond1.316
iface bond1.316 inet static
bond_miimon 100
bond-slaves none
bond_mode 1
address 10.10.0.103
netmask 255.255.255.0
hwaddress 40:7d:0f:52:e3:a9
mtu 8888
```

Step 2 Run the following command to enable the VLAN sub-interface of the bond NIC:

ifup *bond1.316*

Step 3 After all BMSs are configured, ping the IP address in the same network segment as the enhanced high-speed network VLAN sub-interface of other BMSs from each BMS.

```
root@bms-ubuntu:~# ping 10.10.0.102 -I bond1.316
PING 10.10.0.102 (10.10.0.102) from 10.10.0.103 bond1.316: 56(84) bytes of data.
64 bytes from 10.10.0.102: icmp_seq=1 ttl=64 time=0.053 ms
64 bytes from 10.10.0.102: icmp_seq=2 ttl=64 time=0.053 ms
64 bytes from 10.10.0.102: icmp_seq=3 ttl=64 time=0.047 ms
64 bytes from 10.10.0.102: icmp_seq=4 ttl=64 time=0.049 ms
64 bytes from 10.10.0.102: icmp_seq=5 ttl=64 time=0.046 ms
^C
--- 10.10.0.102 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 3996ms
rtt min/avg/max/mdev = 0.046/0.049/0.053/0.008 ms
```

----End

Delete a NIC

1. Obtain the IP address of the bonded enhanced high-speed NIC to be deleted.

2. Use a key or password to log in to the BMS as user **root**.

3. Locate the bond network device and run the following command to stop and delete the device: If the bond has VLAN sub-interfaces, they will be automatically deleted.
   ```
   [root@bms-ubuntu ~]# ifdown p4p1
   [root@bms-ubuntu ~]# ifdown p4p2
   [root@bms-ubuntu ~]# ifdown bond1
   ```

4. Run the following command to delete network configuration file **/etc/network/interfaces.d/60-cloud-init.cfg**:

   rm -f /etc/network/interfaces.d/60-cloud-init.cfg

6.4.8 Configuring an Enhanced High-Speed NIC (Windows Server)

This section uses Windows Server 2012 R2 Standard as an example to describe how to configure an enhanced high-speed network bond of a BMS.

📖 **NOTE**

The configuration methods of other Windows Server OSs are similar to that of Windows Server 2012 R2 Standard.

Add a NIC

Step 1 Log in to a Windows BMS.

Step 2 On the Windows PowerShell CLI of the BMS, run the following command to check the NIC information:

Get-NetAdapter

Information similar to the following is displayed.

```
PS C:\Users\Administrator> Get-NetAdapter

Name                      InterfaceDescription                    ifIndex Status        MacAddress
----                      --------------------                    ------- ------        ----------
eth2                      Intel(R) 82599 10 Gigabit ?????              15 Up            2C-55-D3-C4-
eth0_198befdc-4480-49...6 Intel(R) 82599 10 Gigabit ?????              14 Up            2C-55-D3-C4-
eth1_198befdc-4480-49...6 Intel(R) 82599 10 Gigabit ?????              17 Up            40-7D-0F-52-
eth3                      Intel(R) 82599 10 Gigabit ?????              16 Disconnected  40-7D-0F-52-
Team1                     Microsoft Network Adapter Multiplexo...      21 Up            FA-16-00-86-


PS C:\Users\Administrator>
```

📖 **NOTE**

eth0 and eth1 bear the VPC, and eth3 and eth4 bear the enhanced high-speed network bond. The following steps use eth2 and eth3 to configure the enhanced high-speed network.

Step 3 To improve the outbound traffic on the OS, perform the operations in **Method 1**. If there is no special requirement on traffic, perform the operations in **Method 2**.

- **Method 1: Use the switch standalone mode for the bond in the OS. The outbound traffic is distributed across all active NICs, and the inbound traffic is received through one of the NICs in the team.**

1. Run the following command to create a bond port group for the enhanced high-speed network:

   New-NetLbfoTeam -Name *qinq* **-TeamMembers "***eth2***","***eth3***" - TeamingMode SwitchIndependent -LoadBalancingAlgorithm Dynamic - Confirm:$false**

   ```
   PS C:\Users\Administrator> New-NetLbfoTeam -Name qinq -TeamMembers "eth2","eth3" -TeamingMode Switc
   -LoadBalancingAlgorithm Dynamic -Confirm:$false

   Name                   : qinq
   Members                : {eth3, eth2}
   TeamNics               : qinq
   TeamingMode            : SwitchIndependent
   LoadBalancingAlgorithm : Dynamic
   Status                 : Degraded
   ```

   ☐ NOTE

   > In the command, *qinq* is the name of the port group planned for the enhanced high-speed network, and *eth2* and *eth3* are the network devices that bear the enhanced high-speed network obtained in **Step 2**.

2. Run the following command to query the network adapters:

   get-NetLbfoTeamMember

   ```
   PS C:\Users\Administrator> get-NetLbfoTeamMember

   Name                    : eth0_d7a1277d-7cd9-4fd4-a1ff-a7c4d8009361
   InterfaceDescription    : Intel(R) Ethernet Connection X722 for 10GbE SFP+
   Team                    : Team1
   AdministrativeMode      : Standby
   OperationalStatus       : Standby
   TransmitLinkSpeed(Gbps) : 10
   ReceiveLinkSpeed(Gbps)  : 10
   FailureReason           : AdministrativeDecision

   Name                    : eth1_d7a1277d-7cd9-4fd4-a1ff-a7c4d8009361
   InterfaceDescription    : Intel(R) Ethernet Connection X722 for 10GbE SFP+ #2
   Team                    : Team1
   AdministrativeMode      : Active
   OperationalStatus       : Active
   TransmitLinkSpeed(Gbps) : 10
   ReceiveLinkSpeed(Gbps)  : 10
   FailureReason           : NoFailure

   Name                    : eth3
   InterfaceDescription    : Intel(R) 82599 10 Gigabit ??????? #2
   Team                    : qinq
   AdministrativeMode      : Active
   OperationalStatus       : Active
   TransmitLinkSpeed(Gbps) : 10
   ReceiveLinkSpeed(Gbps)  : 10
   FailureReason           : NoFailure

   Name                    : eth2
   InterfaceDescription    : Intel(R) 82599 10 Gigabit ???????
   Team                    : qinq
   AdministrativeMode      : Active
   OperationalStatus       : Active
   TransmitLinkSpeed(Gbps) : 10
   ReceiveLinkSpeed(Gbps)  : 10
   FailureReason           : NoFailure
   ```

   Get-NetAdapter

   ```
   PS C:\Users\Administrator> Get-NetAdapter

   Name                        InterfaceDescription                      ifIndex Status       MacAddress
   ----                        --------------------                      ------- ------       ----------
   qinq                        Microsoft Network Adapter Multiple...#2        33 Up           DC-99-14-93-DE-C2
   eth1_d7a1277d-7...8009361   Intel(R) Ethernet Connection X722 ...#2        19 Up           2C-97-B1-D2-B4-87
   LOM4                        Intel(R) Ethernet Connection X722 fo...        17 Disconnected 2C-97-B1-D2-B4-89
   Team1                       Microsoft Network Adapter Multiplexo...        24 Up           FA-16-3E-35-C9-F3
   eth0_d7a1277d-7...8009361   Intel(R) Ethernet Connection X722 fo...        15 Up           2C-97-B1-D2-B4-86
   LOM3                        Intel(R) Ethernet Connection X722 ...#2        18 Disconnected 2C-97-B1-D2-B4-88
   eth2                        Intel(R) 82599 10 Gigabit ???????              14 Up           DC-99-14-93-DE-C3
   eth3                        Intel(R) 82599 10 Gigabit ???????              16 Up           DC-99-14-93-DE-C2
   ```

- **Method 2: Use the active/standby mode for the bond in the OS.**

1. Run the following command to create a bond port group for the enhanced high-speed network:

New-NetLbfoTeam -Name *Team2* -TeamMembers "*eth2*","*eth3*" -
TeamingMode SwitchIndependent -LoadBalancingAlgorithm IPAddresses -
Confirm:$false

```
PS C:\Users\Administrator> New-NetLbfoTeam -Name Team2 -TeamMembers "eth2","eth3" -TeamingMode SwitchIndependent
-LoadBalancingAlgorithm IPAddresses -Confirm:$false

Name                  : Team2
Members               : {eth3, eth2}
TeamNics              : Team2
TeamingMode           : SwitchIndependent
LoadBalancingAlgorithm : IPAddresses
Status                : Degraded
```

☐ NOTE

> In the command, *Team2* is the name of the port group planned for the enhanced
> high-speed network, and *eth2* and *eth3* are the network devices that bear the
> enhanced high-speed network obtained in **Step 2**.

2. Run the following command to set a network port of port group Team2
   created in **Step 3.1** to the standby port:

   Set-NetLbfoTeamMember -Name "*eth3*" -AdministrativeMode Standby -
   Confirm:$false

   ☐ NOTE

   > The port group configured for the enhanced high-speed network supports only the
   > active/standby mode. *eth3* is one of the ports of the port group. You can determine
   > which port is configured as the standby port based on your planning.

   get-NetLbfoTeamMember

```
PS C:\Users\Administrator> get-NetLbfoTeamMember

Name                   : eth1_198befdc-4480-4999-a2ab-d910f4e0d8e6
InterfaceDescription   : Intel(R) 82599 10 Gigabit ?????  #4
Team                   : Team1
AdministrativeMode     : Active
OperationalStatus      : Active
TransmitLinkSpeed(Gbps) : 10
ReceiveLinkSpeed(Gbps) : 10
FailureReason          : NoFailure

Name                   : eth0_198befdc-4480-4999-a2ab-d910f4e0d8e6
InterfaceDescription   : Intel(R) 82599 10 Gigabit ?????
Team                   : Team1
AdministrativeMode     : Standby
OperationalStatus      : Standby
TransmitLinkSpeed(Gbps) : 10
ReceiveLinkSpeed(Gbps) : 10
FailureReason          : AdministrativeDecision

Name                   : eth3
InterfaceDescription   : Intel(R) 82599 10 Gigabit ????? #3
Team                   : Team2
AdministrativeMode     : Standby
OperationalStatus      : Failed
TransmitLinkSpeed(Mbps) : 0
ReceiveLinkSpeed(Mbps) : 0
FailureReason          : PhysicalMediaDisconnected

Name                   : eth2
InterfaceDescription   : Intel(R) 82599 10 Gigabit ????? #2
Team                   : Team2
AdministrativeMode     : Active
OperationalStatus      : Active
TransmitLinkSpeed(Gbps) : 10
ReceiveLinkSpeed(Gbps) : 10
FailureReason          : NoFailure
```

   Get-NetAdapter

```
PS C:\Users\Administrator> Get-NetAdapter

Name                    InterfaceDescription              ifIndex Status       MacAddress          LinkSpeed
----                    --------------------              ------- ------       ----------          ---------
Team2                   Microsoft Network Adapter Multiple...#2      30 Up           40-7D-0F-52-E3-AD    10 Gbps
eth2                    Intel(R) 82599 10 Gigabit ?????              15 Up           2C-55-D3-C4-9C-5A    10 Gbps
eth0_198befdc-4480-49...6 Intel(R) 82599 10 Gigabit ?????            14 Up           2C-55-D3-C4-9C-59    10 Gbps
eth1_198befdc-4480-49...6 Intel(R) 82599 10 Gigabit ?????            17 Up           40-7D-0F-52-E3-AE    10 Gbps
eth3                    Intel(R) 82599 10 Gigabit ?????              16 Disconnected 40-7D-0F-52-E3-AD     0 bps
Team1                   Microsoft Network Adapter Multiplexo...      21 Up           FA-16-00-86-9B-83    10 Gbps
```

Step 4 Run the following command to enter the **Network Connections** page:

ncpa.cpl

Then enter the following page.
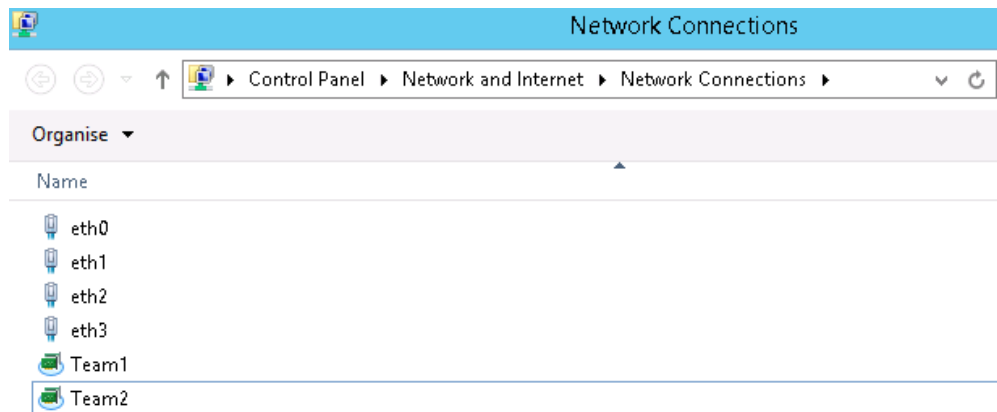
Step 5 Configure the enhanced high-speed network.

1. On the **Network Connections** page, double-click port group **Team2** created in **Step 3** to switch to the **Team2 Status** page.

2. Click **Next** to switch to the **Team2 Properties** page.

3. On the **Networking** tab page, double-click **Internet Protocol Version 4 (TCP/IPv4)** to switch to the **Internet Protocol Version 4 (TCP/IPv4) Properties** page.

4. Select **Use the following IP address**, configure the IP address and subnet mask, and click **OK**.

◻ NOTE

> If the IP address planned for the enhanced high-speed network does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that BMSs communicating through the enhanced high-speed network are in the same network segment as the enhanced high-speed network.

Step 6 Perform the preceding operations to configure other BMSs.

Step 7 After all BMSs are configured, ping the IP address in the same network segment as the enhanced high-speed network of other BMSs from each BMS.

```
PS C:\Users\Administrator> ping 10.10.10.4

Pinging 10.10.10.4 with 32 bytes of data:
Reply from 10.10.10.4: bytes=32 time<1ms TTL=128
Reply from 10.10.10.4: bytes=32 time<1ms TTL=128
Reply from 10.10.10.4: bytes=32 time<1ms TTL=128
Reply from 10.10.10.4: bytes=32 time<1ms TTL=128

Ping statistics for 10.10.10.4:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
PS C:\Users\Administrator>
```

----End

Delete a NIC

1. Log in to a Windows BMS.

2. On the Windows PowerShell CLI of the BMS, run the following command to query information about the bonded enhanced high-speed NICs to be deleted:

   Get-NetLbfoTeamNIC -Team Team2

```
PS C:\Users\Administrator> Get-NetLbfoTeamNIC -Team Team2


Name                    : Team2
InterfaceDescription    : Microsoft Network Adapter Multiplexor Driver #2
Team                    : Team2
VlanID                  :
Primary                 : True
Default                 : True
TransmitLinkSpeed(Gbps) : 10
ReceiveLinkSpeed(Gbps)  : 10
```

3. Run the following command to delete the bonded NICs:

   Remove-NetLbfoTeam -Name "Team2"

```
PS C:\Users\Administrator> Remove-NetLbfoTeam -Name Team2
```

4. Run the following commands to query the NIC information and verify that the NIC is deleted:

   Get-NetAdapter

```
PS C:\Users\Administrator> Get-NetAdapter

Name                      InterfaceDescription                    ifIndex Status       MacAddress             LinkSpeed
----                      --------------------                    ------- ------       ----------             ---------
eth2                      Intel(R) 82599 10 Gigabit ?????              15 Up           2C-55-D3-C4-9C-5A       10 Gbps
eth0_198befdc-4480-49...6 Intel(R) 82599 10 Gigabit ?????              14 Up           2C-55-D3-C4-9C-59       10 Gbps
eth1_198befdc-4480-49...6 Intel(R) 82599 10 Gigabit ?????              17 Up           40-7D-0F-52-E3-AE       10 Gbps
eth3                      Intel(R) 82599 10 Gigabit ?????              16 Up           40-7D-0F-52-E3-AD       10 Gbps
Team1                     Microsoft Network Adapter Multiplexo...      21 Up           FA-16-00-86-9B-83       10 Gbps
```

6.5 User-defined VLAN

6.5.1 Overview

User-defined VLAN

You can use the Ethernet NICs (10GE defined in BMS specifications) not used by the system to configure a user-defined VLAN. The QinQ technology is used to

isolate networks and provide additional physical planes and bandwidths. You can create VLANs to isolate network traffic as needed. User-defined VLAN NICs are in pairs. You can configure NIC bonding to achieve high availability. User-defined VLANs in different AZs cannot communicate.

Ethernet NICs not used by the system by default do not have configuration files and are in **down** state during the system startup. You can run **ifconfig -a** to view the NIC name and run **ifconfig** *eth2* **up** to configure the NIC. The configuration method varies depending on the OS.

For example, on a Linux BMS, eth0 and eth1 are automatically bonded in a VPC network, and eth2 and eth3 are used in a user-defined VLAN. You can send packets with any VLAN tags through the two network interfaces. If you want to allocate a VLAN, configure eth2 and eth3 bonding and create the target VLAN network interface on the bond device. The method is similar to that of creating a bond device and a VLAN sub-interface in a VPC.

> ☐ **NOTE**
>
> In a user-defined VLAN, ports can be bonded or not, and they can only be bonded in active/standby mode.
>
> For more information about NIC bond, visit **https://www.kernel.org/doc/Documentation/networking/bonding.txt**.

For details about how to configure a user-defined VLAN for BMSs running different OSs, see sections **Configuring a User-defined VLAN (SUSE Linux Enterprise Server 12)** to **Configuring a User-defined VLAN (Windows Server)**.

View User-defined VLANs

User-defined VLANs are presented to you through the BMS specifications shown in **Figure 6-9**.

Figure 6-9 BMS specifications

	Flavor name	CPU ⬍	Memory ⬍	Local Disk	Extended Configuration
○	physical.comtest.la...	44 core (test_io...	24*16 GB DDR4	2*600G SAS System Disk RAI...	2*10GE
○	physical.comtest01...	44 core (test_p...	24*16 GB DDR4	2*600G SAS System Disk RAI...	2*10GE
○	physical.comtest04...	44 core (test_h...	24*16 GB DDR4	2*600G SAS System Disk RAI...	2*10GE
◉	physical.d1.large	20 core Intel Xe...	128 GB DDR4	2*600G SAS System Disk RAI...	2 x 2*10GE
○	physical.s3.large	20 core Intel Xe...	128 GB DDR4	2*600G SAS System Disk RAI...	2 x 2*10GE

A BMS created using this flavor provides one two-port 10GE NIC for connecting to the VPC as well as one two-port 10GE extension NIC for a high-speed interconnection between BMSs. You can configure VLANs on the extension NIC as needed.

6.5.2 Configuring a User-defined VLAN (SUSE Linux Enterprise Server 12)

📖 **NOTE**

The network segment of the user-defined VLAN cannot overlap the network information configured on the BMS.

This section uses SUSE Linux Enterprise Server 12 SP1 (x86_64) as an example to describe how to configure a user-defined VLAN for BMSs.

Step 1 Use a key or password to log in to the BMS as user **root**.

Step 2 On the BMS CLI, run the following command to check the NIC information:

ip link

Information similar to the following is displayed.

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group
default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state DOWN mode DEFAULT group
default qlen 1000
    link/ether 38:4c:4f:89:55:8d brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state DOWN mode DEFAULT group
default qlen 1000
    link/ether 38:4c:4f:89:55:8e brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
DEFAULT group default
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
7: bond0.3133@bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
DEFAULT group default
    link/ether fa:16:3e:57:87:6e brd ff:ff:ff:ff:ff:ff
```

📖 **NOTE**

Among the devices, eth0 and eth1 bear the VPC, and eth2 and eth3 bear the user-defined VLAN.

Step 3 Configure the udev rules:

Run the following command to create the **80-persistent-net.rules** file:

cp /etc/udev/rules.d/70-persistent-net.rules /etc/udev/rules.d/80-persistent-net.rules

Write the NIC MAC address and name that are queried in **Step 2** and that are not displayed in **80-persistent-net.rules** to the file. In this way, after the BMS is restarted, the NIC name and sequence will not change.

📖 **NOTE**

Ensure that the NIC MAC address and name are lowercase letters.

vim /etc/udev/rules.d/80-persistent-net.rules

The modification result is as follows:

```
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="38:4c:4f:29:0b:e0", NAME="eth0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="38:4c:4f:29:0b:e1", NAME="eth1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="38:4c:4f:89:55:8d", NAME="eth2"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="38:4c:4f:89:55:8e", NAME="eth3"
```

After the modification, save the change and exit.

Step 4 Run the following command to check the NIC IP address:

ifconfig

Information similar to the following is displayed, where **bond0** and **bond0.313** show the NIC IP addresses automatically allocated by the system when you apply for the BMS:

```
bond0    Link encap:Ethernet  HWaddr FA:16:3E:3D:1C:E0
         inet addr:10.0.1.2  Bcast:10.0.1.255  Mask:255.255.255.0
         inet6 addr: fe80::f816:3eff:fe3d:1ce0/64 Scope:Link
         UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
         RX packets:852 errors:0 dropped:160 overruns:0 frame:0
         TX packets:1121 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:125429 (122.4 Kb)  TX bytes:107221 (104.7 Kb)

bond0.313 Link encap:Ethernet  HWaddr FA:16:3E:57:87:6E
         inet addr:10.0.3.2  Bcast:10.0.3.255  Mask:255.255.255.0
         inet6 addr: fe80::f816:3eff:fe57:876e/64 Scope:Link
         UP BROADCAST RUNNING MULTICAST  MTU:8888  Metric:1
         RX packets:169 errors:0 dropped:0 overruns:0 frame:0
         TX packets:13 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:8684 (8.4 Kb)  TX bytes:1696 (1.6 Kb)

eth0     Link encap:Ethernet  HWaddr FA:16:3E:3D:1C:E0
         UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
         RX packets:428 errors:0 dropped:10 overruns:0 frame:0
         TX packets:547 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:64670 (63.1 Kb)  TX bytes:50132 (48.9 Kb)

eth1     Link encap:Ethernet  HWaddr FA:16:3E:3D:1C:E0
         UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
         RX packets:424 errors:0 dropped:7 overruns:0 frame:0
         TX packets:574 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:60759 (59.3 Kb)  TX bytes:57089 (55.7 Kb)

lo       Link encap:Local Loopback
         inet addr:127.0.0.1  Mask:255.0.0.0
         inet6 addr: ::1/128 Scope:Host
         UP LOOPBACK RUNNING  MTU:65536  Metric:1
         RX packets:8 errors:0 dropped:0 overruns:0 frame:0
         TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:520 (520.0 b)  TX bytes:520 (520.0 b)
```

Step 5 Run the following commands to check the names of bonded NICs:

The in-service bonded NICs cannot be used on the internal communication plane. Therefore, you must obtain them by name.

cd /etc/sysconfig/network

vi ifcfg-bond0

Information similar to the following is displayed, where **bond0** is composed of NICs **eth0** and **eth1**:

```
BONDING_MASTER=yes
TYPE=Bond
STARTMODE=auto
BONDING_MODULE_OPTS="mode=4 xmit_hash_policy=layer3+4 miimon=100"
NM_CONTROLLED=no
BOOTPROTO=dhcp
DEVICE=bond0
USERCONTRL=no
LLADDR=fa:16:3e:3d:1c:e0
BONDING_SLAVE1=eth1
BONDING_SLAVE0=eth0
```

After the query, exit.

Step 6 Run the following commands to check the statuses of all NICs:

ip link

Information similar to the following is displayed.

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group
default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state DOWN mode DEFAULT group
default qlen 1000
    link/ether 38:4c:4f:89:55:8d brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state DOWN mode DEFAULT group
default qlen 1000
    link/ether 38:4c:4f:89:55:8e brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
DEFAULT group default
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
7: bond0.3133@bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
DEFAULT group default
    link/ether fa:16:3e:57:87:6e brd ff:ff:ff:ff:ff:ff
```

Step 7 Run the following commands to change the NIC status **qdisc mq state DOWN** to **qdisc mq state UP**. The following commands use NICs **eth2** and **eth3** as examples.

ip link set *eth2* **up**

ip link set *eth3* **up**

Step 8 Run the following commands to check the statuses of all NICs:

ip link

Information similar to the following is displayed.

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group
default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT group
default qlen 1000
    link/ether 38:4c:4f:89:55:8d brd ff:ff:ff:ff:ff:ff
```

```
5: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT group
default qlen 1000
    link/ether 38:4c:4f:89:55:8e brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
DEFAULT group default
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
7: bond0.3133@bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
DEFAULT group default
    link/ether fa:16:3e:57:87:6e brd ff:ff:ff:ff:ff:ff
```

Step 9 Check the statuses of the NICs in **Step 8** and obtain the names of the NICs in **qdisc mq state UP** state.

Only the NICs that are in **qdisc mq state UP** state and have not been used can be bonded. In this example, such NICs are **eth2** and **eth3**.

The LLADR values of NICs **eth2** and **eth3** are **38:4c:4f:89:55:8d** and **38:4c:4f: 89:55:8e**, respectively.

Step 10 Run the following commands to create the configuration files of NICs **eth2** and **eth3**:

You can copy an existing NIC configuration file and modify it to improve the creation efficiency.

cp *ifcfg-eth0 ifcfg-eth2*

cp *ifcfg-eth1 ifcfg-eth3*

Step 11 Run the following commands to modify the configuration files of NICs **eth2** and **eth3**:

vi *ifcfg-eth2*

vi *ifcfg-eth3*

Modified configuration file of NIC **eth2** is as follows.

In this configuration file, set **MTU** to **8888**, **BOOTPROTO** to **STATIC**, and configure **DEVICE** and **LLADDR** as required.

```
STARTMODE=auto
MTU=8888
NM_CONTROLLED=no
BOOTPROTO=STATIC
DEVICE=eth2
USERCONTRL=no
LLADDR=38:4c:4f:89:55:8d
TYPE=Ethernet
```

Modified configuration file of NIC **eth3** is as follows:

```
STARTMODE=auto
MTU=8888
NM_CONTROLLED=no
BOOTPROTO=STATIC
DEVICE=eth3
USERCONTRL=no
LLADDR=38:4c:4f:89:55:8e
TYPE=Ethernet
```

After the modification, save the change and exit.

Step 12 Run the following command to bond NICs **eth2** and **eth3** to a NIC, for example, **bond1**:

Run the following commands to create the **ifcfg-bond1** file and modify the configuration file:

cp *ifcfg-bond0 ifcfg-bond1*

vi *ifcfg-bond1*

Modified configuration file of NIC **bond1** is as follows.

In this configuration file, **MTU** is set to **8888**, **BONDING_MODULE_OPTS** is set to **mode=1 miimon=100**, **BOOTPROTO** is set to **STATIC**. **DEVICE**, **BONDING_SLAVE1**, **BONDING_SLAVE0**, **IPADDR**, **NETMASK**, and **NETWORK** are configured as required. **LLADDR** is set to the LLADDR value of the **BONDING_SLAVE1** NIC.

```
BONDING_MASTER=yes
TYPE=Bond
MTU=8888
STARTMODE=auto
BONDING_MODULE_OPTS="mode=1 miimon=100"
NM_CONTROLLED=no
BOOTPROTO=STATIC
DEVICE=bond1
USERCONTRL=no
LLADDR=38:4c:4f:89:55:8d
BONDING_SLAVE1=eth2
BONDING_SLAVE0=eth3
IPADDR=10.0.2.2
NETMASK=255.255.255.0
NETWORK=10.0.2.0
```

After the modification, save the change and exit.

Step 13 Make the configuration file take effect.

1. Run the following commands to create a temporary directory and copy the NIC configuration file to this directory:

   mkdir /opt*/tmp/*

   mkdir /opt/tmp/*xml*

   cp /etc/sysconfig/network/ifcfg* /opt/tmp/

   cp /etc/sysconfig/network/config /opt/tmp/

   cp /etc/sysconfig/network/dhcp /opt/tmp/

2. Run the following commands to stop NICs to form **bond1**:

   ip link set *eth2* **down**

   ip link set *eth3* **down**

3. Run the following command to convert the NIC configuration file to a configuration file that can be recognized by the OS:

   /usr/sbin/wicked --log-target=stderr --log-level=debug3 --debug all convert --output /opt/tmp/xml /opt/tmp/

4. Run the following commands to restart the NICs to form **bond1**:

   ip link set *eth2* **up**

   /usr/sbin/wicked --log-target=stderr --log-level=debug3 --debug all ifup --ifconfig /opt/tmp/xml/*eth2***.xml** *eth2*

   ip link set *eth3* **up**

   /usr/sbin/wicked --log-target=stderr --log-level=debug3 --debug all ifup --ifconfig /opt/tmp/xml/*eth3***.xml** *eth3*

> **/usr/sbin/wicked --log-target=stderr --log-level=debug3 --debug all ifup --ifconfig /opt/tmp/xml/bond1.xml** *bond1*

Step 14 Run the following command to query IP addresses:

ip addr show

An example is provided as follows:

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP group
default qlen 1000
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP group
default qlen 1000
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond1 state UP group
default qlen 1000
    link/ether 38:4c:4f:89:55:8d brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond1 state UP group
default qlen 1000
    link/ether 38:4c:4f:89:55:8d brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group
default
    link/ether fa:16:3e:3d:1c:e0 brd ff:ff:ff:ff:ff:ff
    inet 10.0.1.2/24 brd 10.0.1.255 scope global bond0
       valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fe3d:1ce0/64 scope link
       valid_lft forever preferred_lft forever
7: bond0.3133@bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group
default
    link/ether fa:16:3e:57:87:6e brd ff:ff:ff:ff:ff:ff
    inet 10.0.3.2/24 brd 10.0.2.255 scope global bond0.3133
       valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fe57:876e/64 scope link
       valid_lft forever preferred_lft forever
8: bond1: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP group
default
    link/ether 38:4c:4f:89:55:8d brd ff:ff:ff:ff:ff:ff
    inet 10.0.2.2/24 brd 10.0.2.255 scope global bond1
       valid_lft forever preferred_lft forever
    inet6 fe80::3a4c:4fff:fe29:b36/64 scope link
       valid_lft forever preferred_lft forever
```

Step 15 Run the following commands to delete the temporary directory:

cd /opt

rm -rf tmp/

Step 16 Repeat the preceding operations to configure other BMSs.

----End

6.5.3 Configuring a User-defined VLAN (SUSE Linux Enterprise Server 11)

This section uses SUSE Linux Enterprise Server 11 SP4 as an example to describe how to configure a user-defined VLAN for BMSs.

Step 1 Use a key or password to log in to the BMS as user **root**.

Step 2 On the BMS CLI, run the following command to check the NIC information:

ip link

Information similar to the following is displayed:

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP qlen
1000
    link/ether fa:16:3e:0d:13:7c brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP qlen
1000
    link/ether fa:16:3e:0d:13:7c brd ff:ff:ff:ff:ff:ff
4: eth4: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN qlen 1000
    link/ether 40:7d:0f:f4:ff:5c brd ff:ff:ff:ff:ff:ff
5: eth5: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN qlen 1000
    link/ether 40:7d:0f:f4:ff:5d brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP
    link/ether fa:16:3e:0d:13:7c brd ff:ff:ff:ff:ff:ff
```

📖 **NOTE**

Among the devices, eth0 and eth1 bear the VPC, and eth4 and eth5 bear the user-defined VLAN.

Step 3 Run the following command to check whether the **/etc/udev/rules.d/** directory contains the **80-persistent-net.rules** file:

ll /etc/udev/rules.d/ | grep 80-persistent-net.rules

- If yes, and the file contains all NICs except **bond0** and **lo** obtained in step **Step 2** and their MAC addresses, go to step **Step 6**.
- If no, go to step **Step 4**.

Step 4 Run the following command to copy the **/etc/udev/rules.d/70-persistent-net.rules** file and name the copy as **/etc/udev/rules.d/80-persistent-net.rules**.

cp -p /etc/udev/rules.d/70-persistent-net.rules /etc/udev/rules.d/80-persistent-net.rules

Step 5 Configure the udev rules:

Add the NICs and their MAC addresses obtained in step **Step 2**, except **lo**, **eth0**, **eth1**, and **bond0**, to the **/etc/udev/rules.d/80-persistent-net.rules** file. This ensures that the names and sequence of NICs will not change after the BMS is restarted.

📖 **NOTE**

Ensure that NIC MAC addresses and names are lowercase letters.

vim /etc/udev/rules.d/80-persistent-net.rules

The modification result is as follows:

```
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="e8:4d:d0:c8:99:67", NAME="eth0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="e8:4d:d0:c8:99:68", NAME="eth1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:f4:ff:5c", NAME="eth4"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="40:7d:0f:f4:ff:5d", NAME="eth5"
```

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 6 Run the following commands to copy the network configuration file **/etc/sysconfig/network/ifcfg-bond0** to generate the **/etc/sysconfig/network/ifcfg-bond1** file, and copy the **/etc/sysconfig/network/ifcfg-eth0** file to generate the **/etc/sysconfig/network/ifcfg-eth4** and **/etc/sysconfig/network/ifcfg-eth5** files:

cp -p /etc/sysconfig/network/ifcfg-bond0 /etc/sysconfig/network/ifcfg-bond1

cp -p /etc/sysconfig/network/ifcfg-eth0 /etc/sysconfig/network/ifcfg-eth4

cp -p /etc/sysconfig/network/ifcfg-eth0 /etc/sysconfig/network/ifcfg-eth5

Step 7 Run the following commands to edit the **/etc/sysconfig/network/ifcfg-eth4** and **/etc/sysconfig/network/ifcfg-eth5** files:

- **vim /etc/sysconfig/network/ifcfg-eth4**

  Edit the eth4 network configuration file as follows:

  ```
  STARTMODE=auto
  MTU=8888
  NM_CONTROLLED=no
  BOOTPROTO=static
  DEVICE=eth4
  USERCONTRL=no
  LLADDR=40:7d:0f:f4:ff:5c
  TYPE=Ethernet
  ```

  Change the value of **BOOTPROTO** to **static**, that of **DEVICE** to **eth4**, and that of **LLADDR** to the MAC address of eth4, which you can obtain in step **Step 2**. Retain values of other parameters.

- **vim /etc/sysconfig/network/ifcfg-eth5**

  Edit the eth5 network configuration file as follows (similar to eth4):

  ```
  STARTMODE=auto
  MTU=8888
  NM_CONTROLLED=no
  BOOTPROTO=static
  DEVICE=eth5
  USERCONTRL=no
  LLADDR=40:7d:0f:f4:ff:5d
  TYPE=Ethernet
  ```

Step 8 Run the following command to edit the **/etc/sysconfig/network/ifcfg-bond1** file:

vim /etc/sysconfig/network/ifcfg-bond1

Edit the file as follows:

```
BONDING_MASTER=yes
TYPE=Bond
STARTMODE=auto
BONDING_MODULE_OPTS="mode=1 miimon=100"
NM_CONTROLLED=no
BOOTPROTO=static
DEVICE=bond1
USERCONTRL=no
LLADDR=40:7d:0f:f4:ff:5c
BONDING_SLAVE1=eth4
BONDING_SLAVE0=eth5
IPADDR=10.10.10.4
NETMASK=255.255.255.0
MTU=8888
```

Where,

- Change the value of **BOOTPROTO** to **static**.

- Change the value of **DEVICE** to **bond1**.

- Change the value of **LLADDR** to the MAC address of a network device in step **Step 7**, for example, **40:7d:0f:f4:ff:5c**.

- Change the values of **BONDING_SLAVE1** and **BONDING_SLAVE0** to the device names in step **Step 7**, that is, **eth4** and **eth5**.

- Change the value of **IPADDR** to the IP address to be allocated to bond1. If the IP address planned for the user-defined VLAN does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that BMSs communicating through the user-defined VLAN are in the same network segment as the user-defined VLAN. An example value is **10.10.10.4**.

- Set the value of **NETMASK** to the subnet mask of the IP address allocated to bond1.

- Change the value of **MTU** to **8888**.

Retain values of other parameters.

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 9 Run the following commands to restart the network:

ifup *eth4*

ifup *eth5*

ifup *bond1*

⬜ **NOTE**

eth4 and eth5 are the network ports bear the user-defined VLAN and bond1 is the port group of the user-defined VLAN.

Step 10 Run the following commands to check the NIC device status and whether the **bond1** configuration file takes effect:

ip link

ifconfig

```
bms-multinics-test-0002:/etc/sysconfig/network # ifconfig
bond0    Link encap:Ethernet  HWaddr FA:16:3E:0D:13:7C
         inet addr:192.168.20.143  Bcast:192.168.20.255  Mask:255.255.255.0
         inet6 addr: fe80::f816:3eff:fe0d:137c/64 Scope:Link
         UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
         RX packets:5300 errors:0 dropped:1627 overruns:0 frame:0
         TX packets:1926 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:392043 (382.8 Kb)  TX bytes:424419 (414.4 Kb)

bond1    Link encap:Ethernet  HWaddr 40:7D:0F:F4:FF:5C
         inet addr:10.10.10.4  Bcast:10.10.10.255  Mask:255.255.255.0
         inet6 addr: fe80::427d:fff:fef4:ff5c/64 Scope:Link
         UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
         TX packets:15 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:0 (0.0 b)  TX bytes:1194 (1.1 Kb)

eth0     Link encap:Ethernet  HWaddr FA:16:3E:0D:13:7C
         UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
         RX packets:3673 errors:0 dropped:0 overruns:0 frame:0
         TX packets:1926 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:293157 (286.2 Kb)  TX bytes:424419 (414.4 Kb)

eth1     Link encap:Ethernet  HWaddr FA:16:3E:0D:13:7C
         UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
         RX packets:1627 errors:0 dropped:1627 overruns:0 frame:0
         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:98886 (96.5 Kb)  TX bytes:0 (0.0 b)

eth4     Link encap:Ethernet  HWaddr 40:7D:0F:F4:FF:5C
         UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
         TX packets:11 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:0 (0.0 b)  TX bytes:866 (866.0 b)

eth5     Link encap:Ethernet  HWaddr 40:7D:0F:F4:FF:5C
         UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
         TX packets:4 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:1000
         RX bytes:0 (0.0 b)  TX bytes:328 (328.0 b)

lo       Link encap:Local Loopback
         inet addr:127.0.0.1  Mask:255.0.0.0
```

Step 11 Perform the preceding operations to configure other BMSs.

Step 12 After all BMSs are configured, ping the IP addresses of other BMSs from each BMS.

```
bms-multinics-test-0001:/etc/sysconfig/network # tcpdump -i bond1 -nne host 10.10.10.4
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on bond1, link-type EN10MB (Ethernet), capture size 96 bytes
18:51:55.196928 40:7d:0f:f4:ff:5c > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806), length 60: arp who-has 10.10.10.3 tel
l 10.10.10.4
18:51:55.196951 f4:4c:7f:3f:da:07 > 40:7d:0f:f4:ff:5c, ethertype ARP (0x0806), length 42: arp reply 10.10.10.3 is-at
 f4:4c:7f:3f:da:07
18:51:55.197005 40:7d:0f:f4:ff:5c > f4:4c:7f:3f:da:07, ethertype IPv4 (0x0800), length 98: 10.10.10.4 > 10.10.10.3:
ICMP echo request, id 25888, seq 1, length 64
18:51:55.197031 f4:4c:7f:3f:da:07 > 40:7d:0f:f4:ff:5c, ethertype IPv4 (0x0800), length 98: 10.10.10.3 > 10.10.10.4:
ICMP echo reply, id 25888, seq 1, length 64
18:51:56.196847 40:7d:0f:f4:ff:5c > f4:4c:7f:3f:da:07, ethertype IPv4 (0x0800), length 98: 10.10.10.4 > 10.10.10.3:
ICMP echo request, id 25888, seq 2, length 64
18:51:56.196852 f4:4c:7f:3f:da:07 > 40:7d:0f:f4:ff:5c, ethertype IPv4 (0x0800), length 98: 10.10.10.3 > 10.10.10.4:
```

```
bms-multinics-test-0002:/etc/sysconfig/network # ping 10.10.10.3
PING 10.10.10.3 (10.10.10.3) 56(84) bytes of data.
64 bytes from 10.10.10.3: icmp_seq=1 ttl=64 time=0.546 ms
64 bytes from 10.10.10.3: icmp_seq=2 ttl=64 time=0.047 ms
64 bytes from 10.10.10.3: icmp_seq=3 ttl=64 time=0.040 ms
64 bytes from 10.10.10.3: icmp_seq=4 ttl=64 time=0.038 ms
64 bytes from 10.10.10.3: icmp_seq=5 ttl=64 time=0.036 ms
64 bytes from 10.10.10.3: icmp_seq=6 ttl=64 time=0.035 ms
64 bytes from 10.10.10.3: icmp_seq=7 ttl=64 time=0.038 ms
64 bytes from 10.10.10.3: icmp_seq=8 ttl=64 time=0.036 ms
^C
--- 10.10.10.3 ping statistics ---
8 packets transmitted, 8 received, 0% packet loss, time 7000ms
rtt min/avg/max/mdev = 0.035/0.102/0.546/0.167 ms
```

----End

6.5.4 Configuring a User-defined VLAN (Red Hat, CentOS, Oracle Linux, and EulerOS)

This section uses CentOS 6.8 (x86_64) as an example to describe how to configure a user-defined VLAN for BMSs.

📖 NOTE

The configuration methods of Red Hat, Oracle Linux, EulerOS, and CentOS are similar.

Step 1 Use a key or password to log in to the BMS as user **root**.

Step 2 On the BMS CLI, run the following command to check the NIC information:

ip link

Information similar to the following is displayed.

```
[root@bms-qinq-demo ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP qlen 1000
    link/ether fa:16:3e:e5:ec:6a brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP qlen 1000
    link/ether fa:16:3e:e5:ec:6a brd ff:ff:ff:ff:ff:ff
4: eth3: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN qlen 1000
    link/ether f4:4c:7f:3f:da:07 brd ff:ff:ff:ff:ff:ff
5: eth5: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN qlen 1000
    link/ether f4:4c:7f:3f:da:08 brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,PROMISC,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP
    link/ether fa:16:3e:e5:ec:6a brd ff:ff:ff:ff:ff:ff
[root@bms-qinq-demo ~]#
```

📖 NOTE

Among the devices, eth0 and eth1 bear the VPC, and eth3 and eth5 bear the user-defined VLAN.

Step 3 Run the following command to check whether the **/etc/udev/rules.d/** directory contains the **80-persistent-net.rules** file:

ll /etc/udev/rules.d/ | grep 80-persistent-net.rules

- If yes, and the file contains all NICs except **bond0** and **lo** obtained in step **Step 2** and their MAC addresses, go to step **Step 6**.
- If no, go to step **Step 4**.

Step 4 Run the following command to copy the **/etc/udev/rules.d/70-persistent-net.rules** file and name the copy as **/etc/udev/rules.d/80-persistent-net.rules**.

cp -p /etc/udev/rules.d/70-persistent-net.rules /etc/udev/rules.d/80-persistent-net.rules

Step 5 Configure the udev rules:

Write the MAC addresses and names of NICs except eth0 and eth1 obtained in step **Step 2** (those not contained in the **/etc/udev/rules.d/70-persistent-net.rules** file) to the **/etc/udev/rules.d/80-persistent-net.rules** file so that the names and sequence of NICs do not change after the BMS is restarted.

📖 **NOTE**

Ensure that the NIC MAC address and name are lowercase letters.

vim /etc/udev/rules.d/80-persistent-net.rules

The modification result is as follows:

```
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="e8:4d:d0:c8:99:5b", NAME="eth0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="e8:4d:d0:c8:99:5c", NAME="eth1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:3f:da:07", NAME="eth3"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:3f:da:08", NAME="eth5"
~
```

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 6 Run the following commands to copy the network configuration file **/etc/sysconfig/network-scripts/ifcfg-bond0** to generate the **/etc/sysconfig/network-scripts/ifcfg-bond1** file, and copy the **/etc/sysconfig/network-scripts/ifcfg-eth0** file to generate the **/etc/sysconfig/network-scripts/ifcfg-eth3** and **/etc/sysconfig/network/ ifcfg-eth5** files:

cp -p /etc/sysconfig/network-scripts/ifcfg-bond0 /etc/sysconfig/network-scripts/ifcfg-bond1

cp -p /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-eth3

cp -p /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-eth5

Step 7 Run the following commands to edit the **/etc/sysconfig/network-scripts/ifcfg-eth3** and **/etc/sysconfig/network-scripts/ifcfg-eth5** files:

- **vim /etc/sysconfig/network-scripts/ifcfg-eth3**

  Edit the eth3 network configuration file as follows:
  ```
  USERCTL=no
  MTU=8888
  NM_CONTROLLED=no
  BOOTPROTO=static
  DEVICE=eth3
  TYPE=Ethernet
  ONBOOT=yes
  MASTER=bond1
  SLAVE=yes
  ```

Change the value of **BOOTPROTO** to **static**, that of **DEVICE** to the network device name **eth3**, and that of **MASTER** to the port name of the user-defined VLAN (**bond1**). Retain values of other parameters.

- **vim /etc/sysconfig/network-scripts/ifcfg-eth5**

  Edit the eth5 network configuration file as follows (similar to eth3):

  ```
  USERCTL=no
  MTU=8888
  NM_CONTROLLED=no
  BOOTPROTO=static
  DEVICE=eth5
  TYPE=Ethernet
  ONBOOT=yes
  MASTER=bond1
  SLAVE=yes
  ```

Step 8 Run the following command to edit the **/etc/sysconfig/network-scripts/ifcfg-bond1** file:

vim /etc/sysconfig/network-scripts/ifcfg-bond1

Edit the file as follows:

```
MACADDR=f4:4c:7f:3f:da:07
BONDING_MASTER=yes
USERCTL=no
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=static
BONDING_OPTS="mode=1 miimon=100"
DEVICE=bond1
TYPE=Bond
IPADDR=10.10.10.3
NETMASK=255.255.255.0
MTU=8888
```

Where,

- Change the value of **MACADDR** to the MAC address of eth3 or eth5.
- Change the value of **BOOTPROTO** to **static**.
- Change the value of **DEVICE** to **bond1**.
- Change the value of **IPADDR** to the IP address to be allocated to bond1. If the IP address planned for the user-defined VLAN does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that BMSs communicating through the user-defined VLAN are in the same network segment as the user-defined VLAN. An example value is **10.10.10.3**.
- Set the value of **NETMASK** to the subnet mask of the IP address configured for bond1.

Retain values of other parameters.

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 9 Run the following command to enable port group bond1 of the user-defined VLAN:

ifup *bond1*

```
Determining if ip address 10.10.10.3 is already in use for device bond1…
```

Step 10 Perform the preceding operations to configure other BMSs.

Step 11 After all BMSs are configured, ping the IP addresses of other BMSs from each BMS.

```
bms-multinics-test-0001:/etc/sysconfig/network # tcpdump -i bond1 -nne host 10.10.10.4
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on bond1, link-type EN10MB (Ethernet), capture size 96 bytes
18:51:55.196928 40:7d:0f:f4:ff:5c > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806), length 60: arp who-has 10.10.10.3 tel
l 10.10.10.4
18:51:55.196951 f4:4c:7f:3f:da:07 > 40:7d:0f:f4:ff:5c, ethertype ARP (0x0806), length 42: arp reply 10.10.10.3 is-at
 f4:4c:7f:3f:da:07
18:51:55.197005 40:7d:0f:f4:ff:5c > f4:4c:7f:3f:da:07, ethertype IPv4 (0x0800), length 98: 10.10.10.4 > 10.10.10.3:
ICMP echo request, id 25888, seq 1, length 64
18:51:55.197031 f4:4c:7f:3f:da:07 > 40:7d:0f:f4:ff:5c, ethertype IPv4 (0x0800), length 98: 10.10.10.3 > 10.10.10.4:
ICMP echo reply, id 25888, seq 1, length 64
18:51:56.196847 40:7d:0f:f4:ff:5c > f4:4c:7f:3f:da:07, ethertype IPv4 (0x0800), length 98: 10.10.10.4 > 10.10.10.3:
ICMP echo request, id 25888, seq 2, length 64
18:51:56.196852 f4:4c:7f:3f:da:07 > 40:7d:0f:f4:ff:5c, ethertype IPv4 (0x0800), length 98: 10.10.10.3 > 10.10.10.4:
```

```
bms-multinics-test-0002:/etc/sysconfig/network # ping 10.10.10.3
PING 10.10.10.3 (10.10.10.3) 56(84) bytes of data.
64 bytes from 10.10.10.3: icmp_seq=1 ttl=64 time=0.546 ms
64 bytes from 10.10.10.3: icmp_seq=2 ttl=64 time=0.047 ms
64 bytes from 10.10.10.3: icmp_seq=3 ttl=64 time=0.040 ms
64 bytes from 10.10.10.3: icmp_seq=4 ttl=64 time=0.038 ms
64 bytes from 10.10.10.3: icmp_seq=5 ttl=64 time=0.036 ms
64 bytes from 10.10.10.3: icmp_seq=6 ttl=64 time=0.035 ms
64 bytes from 10.10.10.3: icmp_seq=7 ttl=64 time=0.038 ms
64 bytes from 10.10.10.3: icmp_seq=8 ttl=64 time=0.036 ms
^C
--- 10.10.10.3 ping statistics ---
8 packets transmitted, 8 received, 0% packet loss, time 7000ms
rtt min/avg/max/mdev = 0.035/0.102/0.546/0.167 ms
```

----End

6.5.5 Configuring a User-defined VLAN (Ubuntu)

This section uses Ubuntu 16.04 LTS (Xenial Xerus x86_64) as an example to describe how to configure a user-defined VLAN for BMSs.

> **NOTE**
>
> The configuration methods of other Ubuntu OSs are similar to that of Ubuntu 16.04 LTS (Xenial Xerus x86_64).

Step 1 Use a key or password to log in to the BMS as user **root**.

Step 2 On the BMS CLI, run the following command to check the NIC information:

ip link

Information similar to the following is displayed:

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group
default qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:1c:35:37 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:1c:35:37 brd ff:ff:ff:ff:ff:ff
4: enp129s0f0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default
qlen 1000
    link/ether f4:4c:7f:3f:da:07 brd ff:ff:ff:ff:ff:ff
5: enp129s0f1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT group default
qlen 1000
    link/ether f4:4c:7f:3f:da:08 brd ff:ff:ff:ff:ff:ff
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode
DEFAULT group default qlen 1000
    link/ether fa:16:3e:1c:35:37 brd ff:ff:ff:ff:ff:ff
```

> **NOTE**
>
> Among the devices, eth0 and eth1 bear the VPC, and enp129s0f0 and enp129s0f1 bear the user-defined VLAN. In the following steps, enp129s0f0 and enp129s0f1 are used to configure a user-defined VLAN.

Step 3 Run the following command to check whether the **/etc/udev/rules.d/** directory contains the **80-persistent-net.rules** file:

ll /etc/udev/rules.d/ | grep 80-persistent-net.rules

- If yes, and the file contains all NICs except **bond0** and **lo** obtained in step **Step 2** and their MAC addresses, go to step **Step 6**.
- If no, go to step **Step 4**.

Step 4 Run the following command to copy the **/etc/udev/rules.d/70-persistent-net.rules** file and name the copy as **/etc/udev/rules.d/80-persistent-net.rules**.

cp -p /etc/udev/rules.d/70-persistent-net.rules /etc/udev/rules.d/80-persistent-net.rules

Step 5 Configure the udev rules:

Add the NICs and their MAC addresses obtained in step **Step 2**, except **lo**, **eth0**, **eth1**, and **bond0**, to the **/etc/udev/rules.d/80-persistent-net.rules** file. This ensures that the names and sequence of NICs will not change after the BMS is restarted.

> **NOTE**
>
> Ensure that the NIC MAC address and names are lowercase letters.

vim /etc/udev/rules.d/80-persistent-net.rules

The modification result is as follows:

```
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="e8:4d:d0:c8:99:5b", NAME="eth0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="e8:4d:d0:c8:99:5c", NAME="eth1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:3f:da:07",
NAME="enp129s0f0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="f4:4c:7f:3f:da:08",
NAME="enp129s0f1"
```

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 6 Run the following command to copy the **/etc/network/interfaces.d/50-cloud-init.cfg** file to generate the **/etc/network/interfaces.d/60-cloud-init.cfg** file:

cp -p /etc/network/interfaces.d/50-cloud-init.cfg /etc/network/interfaces.d/60-cloud-init.cfg

> **NOTE**
>
> If the **/etc/network/interfaces.d/50-cloud-init.cfg** file does not exist, copy the **/etc/network/interfaces** file and run the following commands:
>
> **mkdir /etc/network/interfaces.d**
>
> **cp -p /etc/network/interfaces /etc/network/interfaces.d/60-cloud-init.cfg**

Step 7 Run the following command to edit the **/etc/network/interfaces.d/60-cloud-init.cfg** file of devices **enp129s0f0** and **enp129s0f1**:

vim /etc/network/interfaces.d/60-cloud-init.cfg

Edit the file as follows:

```
auto enp129s0f0
iface enp129s0f0 inet manual
bond_mode 1
bond-master bond1
bond_miimon 100
mtu 8888
auto enp129s0f1
iface enp129s0f1 inet manual
bond_mode 1
bond-master bond1
bond_miimon 100
mtu 8888
auto bond1
iface bond1 inet static
bond_miimon 100
bond-slaves none
bond_mode 1
address 10.10.10.3
netmask 255.255.255.0
hwaddress f4:4c:7f:3f:da:07
mtu 8888
```

Where,

- **enp129s0f0** and **enp129s0f1** are the NICs that bear the user-defined VLAN.

- **hwaddress** is the MAC address of enp129s0f0.

- Change the value of **address** to the IP address allocated to bond1. If the IP address planned for the user-defined VLAN does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that BMSs communicating through the user-defined VLAN are in the same network segment as the user-defined VLAN.

- Set the value of **netmask** to the subnet mask of the IP address configured for bond1.

Set values of other parameters. For example, set **mtu** to **8888**, **bond_miimon** to **100**, and **bond_mode** to **1**.

After the modification, press **Esc**, enter **:wq**, save the configuration, and exit.

Step 8 Run the following commands to restart the network:

ifup *enp129s0f0*

ifup *enp129s0f1*

◫ NOTE

enp129s0f0 and **enp129s0f1** are the NICs that bear the user-defined VLAN.

Step 9 Run the following commands to check the NIC device status and whether the **bond1** configuration file takes effect:

ip link

```
root@bms-af1d:~# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group default qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode DEFAULT group default qlen 1000
    link/ether fa:16:3e:1c:35:37 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond0 state UP mode DEFAULT group default qlen 1000
    link/ether fa:16:3e:1c:35:37 brd ff:ff:ff:ff:ff:ff
4: enp129s0f0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond1 state UP mode DEFAULT group default qlen 1000
    link/ether f4:4c:7f:3f:da:07 brd ff:ff:ff:ff:ff:ff
5: enp129s0f1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond1 state UP mode DEFAULT group default qlen 1000
    link/ether f4:4c:7f:3f:da:07 brd ff:ff:ff:ff:ff:ff
7: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode DEFAULT group default qlen 1000
    link/ether fa:16:3e:1c:35:37 brd ff:ff:ff:ff:ff:ff
8: bond1: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP mode DEFAULT group default qlen 1000
    link/ether f4:4c:7f:3f:da:07 brd ff:ff:ff:ff:ff:ff
root@bms-af1d:~#
```

ifconfig

```
root@bms-af1d:~# ifconfig
bond0     Link encap:Ethernet  HWaddr fa:16:3e:1c:35:37
          inet addr:192.168.20.195  Bcast:192.168.20.255  Mask:255.255.255.0
          inet6 addr: fe80::f816:3eff:fe1c:3537/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
          RX packets:77 errors:0 dropped:18 overruns:0 frame:0
          TX packets:74 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:6569 (6.5 KB)  TX bytes:12236 (12.2 KB)

bond1     Link encap:Ethernet  HWaddr f4:4c:7f:3f:da:07
          inet addr:10.10.10.3  Bcast:10.10.10.255  Mask:255.255.255.0
          inet6 addr: fe80::f64c:7fff:fe3f:da07/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:10 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 B)  TX bytes:776 (776.0 B)

enp129s0f0 Link encap:Ethernet  HWaddr f4:4c:7f:3f:da:07
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 B)  TX bytes:0 (0.0 B)

enp129s0f1 Link encap:Ethernet  HWaddr f4:4c:7f:3f:da:07
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:10 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 B)  TX bytes:776 (776.0 B)

eth0      Link encap:Ethernet  HWaddr fa:16:3e:1c:35:37
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
          RX packets:3236 errors:0 dropped:3177 overruns:0 frame:0
          TX packets:78 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:197273 (197.2 KB)  TX bytes:12847 (12.8 KB)

eth1      Link encap:Ethernet  HWaddr fa:16:3e:1c:35:37
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
          RX packets:6366 errors:0 dropped:18 overruns:0 frame:0
          TX packets:18224 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:444846 (444.8 KB)  TX bytes:1550404 (1.5 MB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
```

Step 10 Perform the preceding operations to configure other BMSs.

Step 11 After all BMSs are configured, ping the IP addresses of other BMSs from each BMS.

```
root@bms-7b5c:/etc/network/interfaces.d# ifconfig bond1
bond1     Link encap:Ethernet  HWaddr 40:7d:0f:f4:ff:5c
          inet addr:10.10.10.4  Bcast:10.10.10.255  Mask:255.255.255.0
          inet6 addr: fe80::427d:fff:fef4:ff5c/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
          RX packets:11 errors:0 dropped:7 overruns:0 frame:0
          TX packets:20 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:736 (736.0 B)  TX bytes:1308 (1.3 KB)

root@bms-7b5c:/etc/network/interfaces.d# ping 10.10.10.3
PING 10.10.10.3 (10.10.10.3) 56(84) bytes of data.
64 bytes from 10.10.10.3: icmp_seq=1 ttl=64 time=0.061 ms
64 bytes from 10.10.10.3: icmp_seq=2 ttl=64 time=0.053 ms
64 bytes from 10.10.10.3: icmp_seq=3 ttl=64 time=0.046 ms
64 bytes from 10.10.10.3: icmp_seq=4 ttl=64 time=0.038 ms
64 bytes from 10.10.10.3: icmp_seq=5 ttl=64 time=0.050 ms
64 bytes from 10.10.10.3: icmp_seq=6 ttl=64 time=0.035 ms
^C
--- 10.10.10.3 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 4997ms
rtt min/avg/max/mdev = 0.035/0.047/0.061/0.009 ms
root@bms-7b5c:/etc/network/interfaces.d#
root@bms-7b5c:/etc/network/interfaces.d#
```

```
root@bms-af1d:~# ifconfig bond1
bond1     Link encap:Ethernet  HWaddr f4:4c:7f:3f:da:07
          inet addr:10.10.10.3  Bcast:10.10.10.255  Mask:255.255.255.0
          inet6 addr: fe80::f64c:7fff:fe3f:da07/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:8888  Metric:1
          RX packets:5 errors:0 dropped:1 overruns:0 frame:0
          TX packets:14 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:376 (376.0 B)  TX bytes:1056 (1.0 KB)

root@bms-af1d:~# tcpdump -i bond1 -nne host 10.10.10.4
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on bond1, link-type EN10MB (Ethernet), capture size 262144 bytes
10:04:52.930343 40:7d:0f:f4:ff:5c > f4:4c:7f:3f:da:07, ethertype IPv4 (0x0800), length 98: 10.10.10.4 > 10
.10.10.3: ICMP echo request, id 19052, seq 1, length 64
10:04:52.930360 f4:4c:7f:3f:da:07 > 40:7d:0f:f4:ff:5c, ethertype IPv4 (0x0800), length 98: 10.10.10.3 > 10
.10.10.4: ICMP echo reply, id 19052, seq 1, length 64
10:04:53.929346 40:7d:0f:f4:ff:5c > f4:4c:7f:3f:da:07, ethertype IPv4 (0x0800), length 98: 10.10.10.4 > 10
.10.10.3: ICMP echo request, id 19052, seq 2, length 64
10:04:53.929354 f4:4c:7f:3f:da:07 > 40:7d:0f:f4:ff:5c, ethertype IPv4 (0x0800), length 98: 10.10.10.3 > 10
.10.10.4: ICMP echo reply, id 19052, seq 2, length 64
```

----End

6.5.6 Configuring a User-defined VLAN (Windows Server)

This section uses Windows Server 2012 R2 Standard as an example to describe how to configure a user-defined VLAN for BMSs.

📖 **NOTE**

The configuration methods of other Windows Server OSs are similar to that of Windows Server 2012 R2 Standard.

Step 1 Log in to a Windows BMS.

Step 2 On the Windows PowerShell CLI of the BMS, run the following command to check the NIC information:

Get-NetAdapter

Information similar to the following is displayed.

```
PS C:\Users\Administrator> Get-NetAdapter

Name                    InterfaceDescription               ifIndex Status    MacAddress            LinkSpeed
----                    --------------------               ------- ------    ----------            ---------
eth3                    Intel(R) 82599 10 Gigabit ???????       18 Up        F4-4C-7F-3F-DA-08      10 Gbps
eth2                    Intel(R) 82599 10 Gigabit ???????       16 Up        F4-4C-7F-3F-DA-07      10 Gbps
eth1                    Intel(R) 82599 10 Gigabit ???????       15 Up        E8-4D-D0-C8-99-5C      10 Gbps
eth0                    Intel(R) 82599 10 Gigabit ???????       17 Up        E8-4D-D0-C8-99-5B      10 Gbps
Team1                   Microsoft Network Adapter Multiplexo... 23 Up        FA-16-3E-C8-C4-73      10 Gbps

PS C:\Users\Administrator> _
```

□ NOTE

Among the devices, eth0 and eth1 bear the VPC, and eth2 and eth3 bear the user-defined VLAN. The following steps use eth2 and eth3 to configure a user-defined VLAN.

Step 3 To improve the outbound traffic on the OS, perform the operations in **Method 1**. If there is no special requirement on traffic, perform the operations in **Method 2**.

- **Method 1: Use the switch independent mode for the team in the OS. The outbound traffic is distributed across all active NICs, and the inbound traffic is received through one of the NICs in the team.**

  1. Run the following command to create a port group for the user-defined VLAN:

     New-NetLbfoTeam -Name *qinq* **-TeamMembers "***eth2***","***eth3***" - TeamingMode SwitchIndependent -LoadBalancingAlgorithm Dynamic - Confirm:$false**

     ```
     PS C:\Users\Administrator> New-NetLbfoTeam -Name qinq -TeamMembers "eth2","eth3" -TeamingMode SwitchIndependent
     -LoadBalancingAlgorithm Dynamic -Confirm:$false

     Name                 : qinq
     Members              : {eth2, eth3}
     TeamNics             : qinq
     TeamingMode          : SwitchIndependent
     LoadBalancingAlgorithm : Dynamic
     Status               : Degraded
     ```

     □ NOTE

     In the command, *qinq* is the name of the port group planned for the user-defined VLAN, and *eth2* and *eth3* are the network devices that bear the user-defined VLAN obtained in step **Step 2**.

  2. Run the following command to query the network adapters:

     Get-NetLbfoTeamMember

     ```
     PS C:\Users\Administrator> Get-NetLbfoTeamMember

     Name                    : eth0_d7a1277d-7cd9-4fd4-a1ff-a7c4d8009361
     InterfaceDescription    : Intel(R) Ethernet Connection X722 for 10GbE SFP+
     Team                    : Team1
     AdministrativeMode      : Standby
     OperationalStatus       : Standby
     TransmitLinkSpeed(Gbps) : 10
     ReceiveLinkSpeed(Gbps)  : 10
     FailureReason           : AdministrativeDecision

     Name                    : eth1_d7a1277d-7cd9-4fd4-a1ff-a7c4d8009361
     InterfaceDescription    : Intel(R) Ethernet Connection X722 for 10GbE SFP+ #2
     Team                    : Team1
     AdministrativeMode      : Active
     OperationalStatus       : Active
     TransmitLinkSpeed(Gbps) : 10
     ReceiveLinkSpeed(Gbps)  : 10
     FailureReason           : NoFailure

     Name                    : eth2
     InterfaceDescription    : Intel(R) 82599 10 Gigabit ???????
     Team                    : qinq
     AdministrativeMode      : Active
     OperationalStatus       : Active
     TransmitLinkSpeed(Gbps) : 10
     ReceiveLinkSpeed(Gbps)  : 10
     FailureReason           : NoFailure

     Name                    : eth3
     InterfaceDescription    : Intel(R) 82599 10 Gigabit ???????
     Team                    : qinq
     AdministrativeMode      : Active
     OperationalStatus       : Active
     TransmitLinkSpeed(Gbps) : 10
     ReceiveLinkSpeed(Gbps)  : 10
     FailureReason           : NoFailure
     ```

     Get-NetAdapter

```
PS C:\Users\Administrator> Get-NetAdapter

Name                          InterfaceDescription               ifIndex Status       MacAddress
----                          --------------------               ------- ------       ----------
qinq                          Microsoft Network Adapter Multiple...#2   33 Up          DC-99-14-93-DE-C2
eth1_d7a1277d-7...8009361     Intel(R) Ethernet Connection X722 ...#2   19 Up          2C-97-B1-D2-B4-87
LOM4                          Intel(R) Ethernet Connection X722 fo...   17 Disconnected 2C-97-B1-D2-B4-89
Team1                         Microsoft Network Adapter Multiplexo...   24 Up          FA-16-3E-35-C9-F3
eth0_d7a1277d-7...8009361     Intel(R) Ethernet Connection X722 fo...   15 Up          2C-97-B1-D2-B4-86
LOM3                          Intel(R) Ethernet Connection X722 ...#2   18 Disconnected 2C-97-B1-D2-B4-88
eth3                          Intel(R) 82599 10 Gigabit ???????         14 Up          DC-99-14-93-DE-C3
eth2                          Intel(R) 82599 10 Gigabit ???????         16 Up          DC-99-14-93-DE-C2
```

- **Method 2: Use the active-active mode for the team in the OS.**

1. Run the following command to create a port group for the user-defined VLAN:

   New-NetLbfoTeam -Name *Team2* **-TeamMembers "***eth2***","***eth3***" -TeamingMode SwitchIndependent -LoadBalancingAlgorithm IPAddresses -Confirm:$false**

```
PS C:\Users\Administrator> New-NetLbfoTeam -Name Team2 -TeamMembers "eth2","eth3"  -TeamingMode SwitchIndependent
 -LoadBalancingAlgorithm IPAddresses -Confirm:$false


Name                  : Team2
Members               : {eth2, eth3}
TeamNics              : Team2
TeamingMode           : SwitchIndependent
LoadBalancingAlgorithm : IPAddresses
Status                : Down


PS C:\Users\Administrator>
```

   ☐ **NOTE**

   In the command, *Team2* is the name of the port group planned for the user-defined VLAN, and *eth2* and *eth3* are the network devices that bear the user-defined VLAN obtained in step **Step 2**.

2. Run the following command to set a network port of port group Team2 created in **Step 3.1** to the standby port:

   Set-NetLbfoTeamMember -Name "*eth2***" -AdministrativeMode Standby -Confirm:$false**

   ☐ **NOTE**

   The port group configured for the user-defined VLAN supports only the active/standby mode. *eth2* is one of the ports of the port group. You can determine which port is configured as the standby port based on your planning.

   get-NetLbfoTeamMember

```
PS C:\Users\Administrator> get-NetLbfoTeamMember

Name                    : eth2
InterfaceDescription    : Intel(R) 82599 10 Gigabit ???????    #2
Team                    : Team2
AdministrativeMode      : Standby
OperationalStatus       : Standby
TransmitLinkSpeed(Gbps) : 10
ReceiveLinkSpeed(Gbps)  : 10
FailureReason           : AdministrativeDecision

Name                    : eth3
InterfaceDescription    : Intel(R) 82599 10 Gigabit ???????    #4
Team                    : Team2
AdministrativeMode      : Active
OperationalStatus       : Active
TransmitLinkSpeed(Gbps) : 10
ReceiveLinkSpeed(Gbps)  : 10
FailureReason           : NoFailure

Name                    : eth0
InterfaceDescription    : Intel(R) 82599 10 Gigabit ???????    #3
Team                    : Team1
AdministrativeMode      : Standby
OperationalStatus       : Standby
TransmitLinkSpeed(Gbps) : 10
ReceiveLinkSpeed(Gbps)  : 10
FailureReason           : AdministrativeDecision

Name                    : eth1
InterfaceDescription    : Intel(R) 82599 10 Gigabit ???????
Team                    : Team1
AdministrativeMode      : Active
OperationalStatus       : Active
TransmitLinkSpeed(Gbps) : 10
ReceiveLinkSpeed(Gbps)  : 10
FailureReason           : NoFailure
```

Get-NetAdapter

```
PS C:\Users\Administrator> Get-NetAdapter

Name          InterfaceDescription                  ifIndex Status    MacAddress          LinkSpeed
----          --------------------                  ------- ------    ----------          ---------
Team2         Microsoft Network Adapter Multiple...#2    33 Up        F4-4C-7F-3F-DA-08     10 Gbps
eth3          Intel(R) 82599 10 Gigabit ???????         18 Up        F4-4C-7F-3F-DA-08     10 Gbps
eth2          Intel(R) 82599 10 Gigabit ???????         16 Up        F4-4C-7F-3F-DA-07     10 Gbps
eth1          Intel(R) 82599 10 Gigabit ???????         15 Up        E8-4D-D0-C8-99-5C     10 Gbps
eth0          Intel(R) 82599 10 Gigabit ???????         17 Up        E8-4D-D0-C8-99-5B     10 Gbps
Team1         Microsoft Network Adapter Multiplexo...    23 Up        FA-16-3E-C8-CA-73     10 Gbps
```

Step 4 Run the following command to enter the **Network Connections** page:

ncpa.cpl

Then enter the following page.
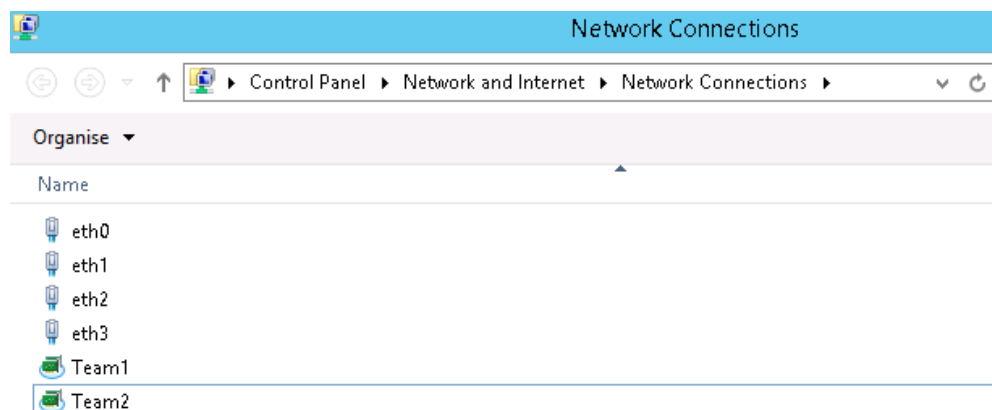
Step 5 Configure a user-defined VLAN.

1. On the **Network Connections** page, double-click port group **Team2** created in **Step 3** to switch to the **Team2 Status** page.

2. Click **Next** to switch to the **Team2 Properties** page.

3. On the **Networking** tab page, double-click **Internet Protocol Version 4 (TCP/IPv4)** to switch to the **Internet Protocol Version 4 (TCP/IPv4) Properties** page.

4. Select **Use the following IP address**, configure the IP address and subnet mask, and click **OK**.

☐ NOTE

If the IP address planned for the user-defined VLAN does not conflict with the VPC network segment, you can plan the IP address as needed, only to ensure that BMSs communicating through the user-defined VLAN are in the same network segment as the user-defined VLAN.

Step 6 Perform the preceding operations to configure other BMSs.

Step 7 After all BMSs are configured, ping the IP addresses of other BMSs from each BMS.

Step 8 If you want to configure VLAN sub-interfaces to isolate network planes, perform the following operations:

Run the following command to create a VLAN sub-interface based on the existing Team2:

Add-NetLbfoTeamNIC -Team "Team2" -VlanID *XXX* **-Confirm:$false**

In the preceding command, **Team2** indicates the bond name, and *XXX* indicates the VLAN ID.

After the VLAN sub-interface is created, configure the IP address and subnet mask of network port Team2-VLAN 500 by referring to **Step 4** and **Step 5**.

----End

6.6 IB Network

6.6.1 Overview

IB Network

The IB network features low latency and high bandwidth and is used in a number of High Performance Computing (HPC) projects. It uses the 100 Gbit/s Mellanox IB NIC, dedicated IB switch, and controller software UFM to ensure network communication and management, and uses the Partition Key to isolate IB networks of different tenants (similar to VLANs in the Ethernet). The IB network supports two communication modes, RDMA and IPoIB.

To create an IB network, you must select a flavor that supports the IB network during BMS creation. After an IB network is provisioned, BMSs can communicate with each other in RDMA mode. In the IPoIB communication mode, you need to configure IP addresses on the IB network port. You can use static IP addresses or

IP addresses dynamically assigned by DHCP. Examples of static IP addresses are as follows:

```
#/etc/sysconfig/network/ifcfg-ib0
DEVICE=ib0
TYPE=InfiniBand
ONBOOT=yes
HWADDR=80:00:00:4c:fe:80:00:00:00:00:00:00:f4:52:14:03:00:7b:cb:a1
BOOTPROTO=none
IPADDR=172.31.0.254
PREFIX=24
NETWORK=172.31.0.0
BROADCAST=172.31.0.255
IPV4_FAILURE_FATAL=yes
IPV6INIT=no
MTU=65520
CONNECTED_MODE=yes
NAME=ib0
```

⚠ CAUTION

In the IB network, an IP address is assigned to a new BMS in DHCP mode by default. You can manually specify a static IP address not in use to the BMS.

For more information about the IPoIB communication mode, see **https://www.kernel.org/doc/Documentation/infiniband/ipoib.txt**.

View IB Networks

IB networks are presented to you through the BMS specifications shown in **Figure 6-10**. You need to configure and plan the VLANs and IP addresses.

Figure 6-10 BMS extended configuration

	Flavor name	CPU	Memory	Local Disk	Extended Configuration
○	physical.d2.xmedium	36 core Intel Xe…	384 GB DDR4	2*600G SAS RAID 1+ 24*1.8T …	2 x 2*10GE
○	physical.h2.large	36 core Intel(R)…	12*16 GB DDR4	1*1.6TB NVMe SSD Disk	1*100G IB + 2*10GE
Sold ○	physical.m2.medium	96 core 4*24Co…	32*64 GB DIMM	2*600GB SAS System Disk RA…	2x2*10GE
○	physical.s3.large	20 core Intel Xe…	128 GB DDR4	2*600G SAS System Disk RAI…	2 x 2*10GE
Sold ○	physical.s4.3xlarge	44 core Intel Xe…	384 GB DDR4	NA	2 x 2*10GE
Sold ○	physical.s4.large	20 core Intel Xe…	192 GB DDR4	NA	2 x 2*10GE

6.7 User-defined Network

6.7.1 Overview

User-defined Network

BMSs support secondary virtualization. The user-defined network enables communication between the VM network and the VPC, and provides functions similar to the VPC network for VMs. Application scenarios and benefits include:

- Private clouds such as XenServer, VMware, and Hyper-V can be provisioned using BMSs.

- VMs or containers provisioned in the private cloud on BMSs can communicate with the Internet.

- VMs or containers provisioned in the private cloud on BMSs can communicate with each other.

- Network ACLs can be configured for the subnets of the VMs or containers provisioned in the private cloud on BMSs.

- VMs or containers provisioned in the private cloud on BMSs can communicate with ECSs in VPCs.

- The private cloud system based on BMSs and Direct Connect can communicate with the private cloud in the on-premise data center.

The best practice Virtualization on BMS is provided for you to use user-defined networks. For details, see **Virtualization on BMS**.

Notes

- If you have multiple VPCs, you can only create user-defined networks in one VPC.

- If each of your BMSs has only one NIC, you cannot create user-defined networks.

- You must create a BMS with a user-defined VLAN before creating a user-defined network.

- A maximum of two user-defined networks can be created.

- Security groups cannot be configured for user-defined subnets.

6.7.2 Creating and Managing a User-defined Network

Create a User-defined Network

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Click the **User-defined Networks** tab and click **Create User-defined Network**.

4. Set **Name**, **VPC**, **AZ**, VLAN, and **User-defined Subnet**.

Figure 6-11 Creating a user-defined network

NOTE

- **User-defined Subnet**: The subnet mask ranges from 16 to 29. The gateway IP address of the user-defined subnet cannot be one used by any other tenant. In addition, the CIDR of the user-defined subnet cannot conflict with that of the VPC subnet.

- The VLAN ranges from 1 to 4094. User-defined subnets of the same user-defined network must have different VLANs.

5. Click **OK**.

    After the user-defined network is created, you can view its status, VPC, and subnets in the list of user-defined networks.

    Figure 6-12 User-defined networks

    NOTE

    If you have not created a multi-NIC BMS, message "Insufficient multi-NIC BMSs." is displayed and the user-defined network fails to be created.

Change the Name of a User-defined Network

1. On the BMS console, click the **User-defined Networks** tab. Locate the row that contains the target user-defined network and click **Modify** in the **Operation** column.

2. Enter a new name and click **OK**.

View Details About a User-defined Network

1. On the BMS console, click the **User-defined Networks** tab.

2. In the list of user-defined networks, click the name of the target user-defined network.

   You can view the name, AZ, number of user-defined subnets, and VPC information of the user-defined network.

   You can also change the name of the user-defined network by clicking ✎ next to its name, entering a new name, and then clicking ✔ .

Manage a User-defined Subnet

- Add a user-defined subnet.

  On the page showing details of the user-defined network, click **Add User-defined Subnet**. In the displayed dialog box, enter an IP address, set the VLAN, and click **OK**.

- View details about a user-defined subnet.

  In the list of user-defined subnets, click the name of the target user-defined subnet to go to its details page.

  On the **IP Addresses** page, you can view the IP address of the user-defined subnet and the EIP bound to the subnet.

  On the **BMSs** page, you can view information about all the BMSs with multiple NICs.

  Figure 6-13 BMSs tab

IP Address	BMSs	
Name	ID	NIC Name
bms-dsy-hyper-v-0002	2b6fdbaf-89b4-4893-b083-beb53a1445c2	enhanced-nic2
bms-dsy-hyper-v-0001	061d6fdc-e18d-4964-9af1-e49c38c83e37	enhanced-nic2

- Modify a user-defined subnet.

  You cannot directly modify a user-defined subnet. You must delete the user-defined network and then create a new user-defined network and a user-defined subnet again.

- Delete a user-defined subnet.

  In the list of user-defined subnets, locate the target user-defined subnet and click **Delete** in the **Operation** column.

  If a user-defined network has only one subnet, you are not allowed to delete the subnet.

Delete a User-defined Network

1. On the BMS console, click the **User-defined Networks** tab.

2. Locate the target user-defined network, click **Delete** in the **Operation** column, and click **OK** in the displayed dialog box.

After the last multi-NIC BMS using a user-defined network is deleted, the user-defined network will be unavailable and needs to be deleted manually. To use user-defined networks later, you must create new ones.

6.8 User-defined Network ACL

6.8.1 Overview

User-defined Network ACL

A user-defined network ACL is an optional layer of security for your subnets. You can associate one or more user-defined subnets with a user-defined network ACL for controlling traffic in and out of the subnets.

Figure 6-14 User-defined Network ACL
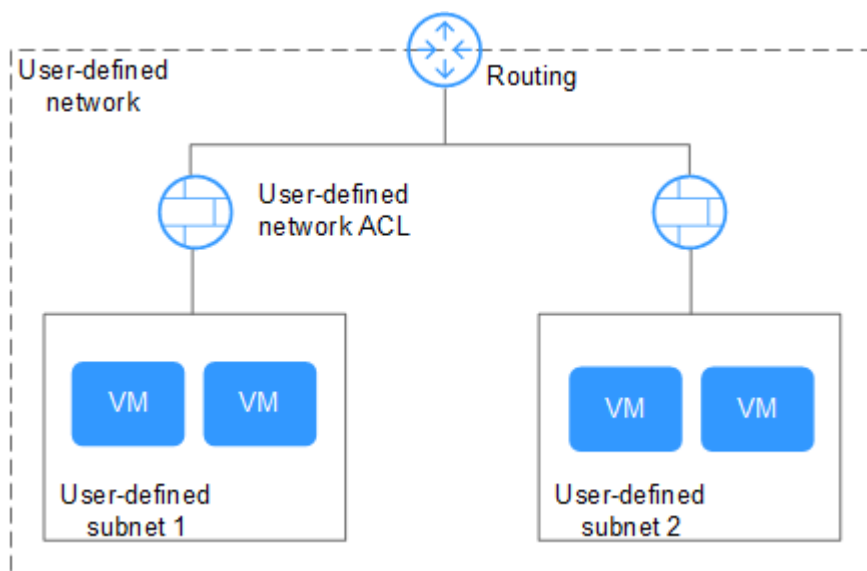
Similar to security groups, user-defined network ACLs provide access control functions and add an additional layer of defense to your VPC. Security groups have only the allow rules, while user-defined network ACLs have both the allow rules and deny rules to implement fine-grained and comprehensive access control. Currently, user-defined networks do not support the security group function. You can use user-defined network ACLs to meet your network configuration requirements.

Notes

- User-defined network ACLs are similar to the network ACLs in the VPC service. For more information, see **Network ACL**.
- By default, your user-defined network has no ACL. You can create a user-defined network ACL and associate it with a subnet. By default, a user-defined network ACL denies all inbound traffic to and outbound traffic from the associated subnet until you add rules.

● A user-defined network ACL can be associated with multiple user-defined subnets. However, a user-defined subnet can be associated with only one user-defined network ACL.

User-defined network ACL Configuration Process

Figure 6-15 Configuration process
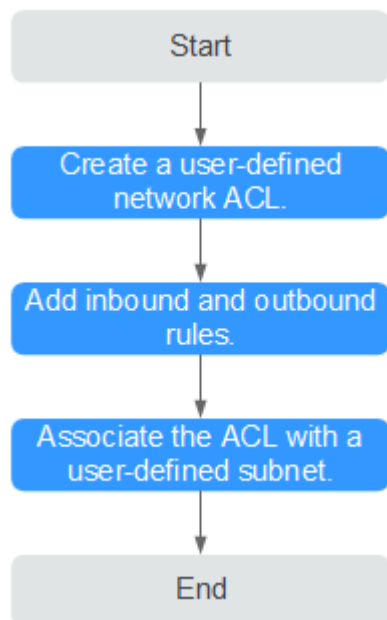
For details, see **Creating and Managing a User-defined Network ACL**.

6.8.2 Configuration Examples

Example 1: Denying Access from Specific Ports

In this example, you need to prevent WannaCry ransomware attacks and deny access from the port that can be exploited by WannaCry, for example TCP 445. You can add a user-defined network ACL rule to deny all incoming traffic from TCP port 445.

Table 6-1 lists the inbound rule required.

Table 6-1 User-defined network ACL rule

Direc tion	Polic y	Proto col	Source	Source Port Range	Destina tion	Destina tion Port Range	Description
Inbo und	Deny	TCP	0.0.0.0/ 0	1-6553 5	0.0.0.0/0	445	Denies inbound traffic from any IP address through TCP port 445.

Example 2: Allowing Access from Specific Ports and Protocols

In this example, a BMS in a user-defined subnet is used as the web server, and you must allow inbound traffic from HTTP port 80 and HTTPS port 443 and allow all outbound traffic.

Table 6-2 lists the inbound and outbound user-defined network ACL rules required.

Table 6-2 User-defined network ACL rules

Direc tion	Polic y	Proto col	Source	Source Port Range	Destina tion	Destina tion Port Range	Description
Inbo und	Perm it	TCP	0.0.0.0/ 0	1-6553 5	0.0.0.0/0	80	Allows inbound HTTP traffic from any IP address to BMSs in the user- defined subnet through port 80.
Inbo und	Perm it	TCP	0.0.0.0/ 0	1-6553 5	0.0.0.0/0	443	Allows inbound HTTP traffic from any IP address to BMSs in the user- defined subnet through port 443.

Direction	Policy	Protocol	Source	Source Port Range	Destination	Destination Port Range	Description
Outbound	Permit	All	0.0.0.0/0	0	0.0.0.0/0	0	Allow all outbound traffic from the user-defined subnet.

6.8.3 Creating and Managing a User-defined Network ACL

Create a User-defined Network ACL

Each user can create a maximum of 200 user-defined network ACLs by default.

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Click the **User-defined Networks** tab and click **Create User-defined Network ACL**.

4. Enter a name and description and click **OK**.

5. After the user-defined network ACL is created successfully, it is displayed in the list of user-defined network ACLs.

Add an Inbound/Outbound Rule

Add an inbound or outbound rule based on your network security requirements.

1. On the BMS console, click the **User-defined Networks** tab and select **User-defined Network ACLs**.

2. Click the name of the target user-defined network ACL.

3. On the **Inbound Rules** or **Outbound Rules** tab, click **Add Rule** to add an inbound or outbound rule.

   You can click **+** to add more rules.

Table 6-3 Parameter description

Parameter	Description	Example Value
Policy	Specifies the policy of the user-defined network ACL. This parameter is mandatory. You can select a value from the drop-down list.	Permit

Parameter	Description	Example Value
Protocol	Specifies the protocol supported by the user-defined network ACL. This parameter is mandatory. You can select a value from the drop-down list. Currently, only **TCP**, **UDP**, **ICMP**, and **All protocols** are available. If you select **ICMP** or **All protocols**, port information cannot be specified.	TCP
Source	Specifies the source IP address from which the traffic is permitted. The default value is **0.0.0.0/0**, which indicates that traffic from all IP addresses is permitted. For example: xxx.xxx.xxx.xxx/32 (IP address) xxx.xxx.xxx.0/24 (CIDR block) 0.0.0.0/0 (any IP address)	0.0.0.0/0
Source Port Range	Specifies the source port number or port number range. The value ranges from 1 to 65535. For a port number range, enter two port numbers connected by a hyphen (-), such as **1-100**. You must specify this parameter if **TCP** or **UDP** is selected for **Protocol**.	22 or 22-30
Destination	Specifies the destination IP address to which the traffic is permitted. The default value is **0.0.0.0/0**, which indicates that traffic from all IP addresses is permitted. For example: xxx.xxx.xxx.xxx/32 (IP address) xxx.xxx.xxx.0/24 (CIDR block) 0.0.0.0/0 (any IP address)	0.0.0.0/0
Destination Port Range	Specifies the destination port number or port number range. The value ranges from 1 to 65535. For a port number range, enter two port numbers connected by a hyphen (-), such as **1-100**. You must specify this parameter if **TCP** or **UDP** is selected for **Protocol**.	22 or 22-30

Parameter	Description	Example Value
Description	Specifies description of the user-defined network ACL rule. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (<) or (>).	-

4.   Click **OK**.

Associate a User-defined Network ACL with a User-defined Subnet

When you need to associate a user-defined network ACL with a user-defined subnet, go to the user-defined network ACL details page and click the **User-defined Subnets** tab. By default, a user-defined network ACL denies all inbound traffic to and outbound traffic from the associated subnet until you add rules.

1.   On the BMS console, click the **User-defined Networks** tab and select **User-defined Network ACLs**.

2.   Click the name of the target user-defined network ACL.

3.   On the **Associate Subnet** page, click **Associate**.

Figure 6-16 Associate Subnet
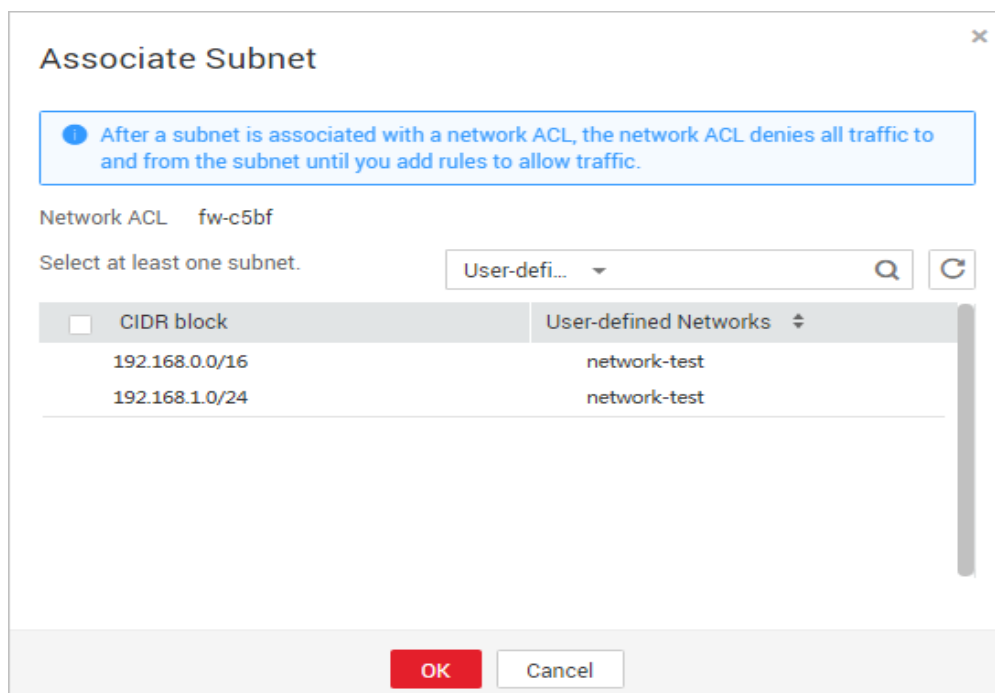
4.   Select the user-defined subnets you want to associate with the user-defined network ACL and click **OK**.

 NOTE

User-defined subnets that have already been associated with user-defined network ACLs will not be displayed on the page for you to select. One-click user-defined subnet association and disassociation are not supported. If you want to associate a user-defined subnet that has already been associated with another user-defined network ACL, you must first disassociate the subnet from the original network ACL.

7 Security

7.1 Security Group

7.1.1 Adding Security Group Rules

Scenarios

The default security group rule allows all outgoing data packets. BMSs in a security group can access each other without the need to add access rules. After a security group is created, you can create different access rules for the security group to protect the BMSs that are added to this security group.

📖 **NOTE**

You can add only one security group when creating a BMS. After the BMS is created, you can modify the security group of each NIC on the BMS details page.

Suggestions

- When adding a security group rule for a BMS, grant the minimum permissions possible:
  - Enable specific ports rather than a port range, for example, port 80.
  - Be cautious to authorize source address 0.0.0.0/0 (entire network segment).
- You are not advised to use one security group to manage all applications because isolation requirements for different layers vary.
- Configuring a security group for each BMS is unnecessary. Instead, you can add BMSs with the same security protection requirements to the same security group.
- Simple security group rules are recommended. For example, if you add a BMS to multiple security groups, the BMS may comply with hundreds of security group rules, and a change to any rule may cause network disconnection for the BMS.

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. In the BMS list, click the name of the BMS whose security group rules you want to modify.

   The page showing details of the BMS is displayed.

4. Click the **Security Groups** tab and then ⌄ to view security group rules.

5. Click the security group ID.

   The system automatically switches to the **Security Group** page.

6. Click **Manage Rule** in the **Operation** column. On the security group details page, add a rule.

   Value **Inbound** indicates that traffic enters the security group, and value **Outbound** indicates that traffic leaves the security group.

Table 7-1 Parameter description

Parameter	Description	Example Value
Protocol	Specifies the network protocol for which the security group rule takes effect. The value can be **TCP**, **UDP**, **ICMP**, **HTTP**, or others.	TCP
Port	Specifies the port or port range for which the security group rule takes effect. The value ranges from **0** to **65535**.	22 or 22-30
Source	Specifies the traffic source (inbound rule). This parameter is required for an inbound rule. The value can be an IP address or a security group.	0.0.0.0/0 default
Destination	Specifies the traffic destination (outbound rule). This parameter is required for an outbound rule. The value can be an IP address or a security group.	0.0.0.0/0 default

📖 **NOTE**

The default source IP address **0.0.0.0/0** indicates that all IP addresses can access BMSs in the security group.

7.1.2 Security Group Configuration Examples

Case 1: BMSs in Different Security Groups Need to Communicate with Each Other Through an Internal Network

- Scenario

  Resources on a BMS in a security group need to be copied to a BMS in another security group. The two BMSs are in the same VPC. Then, you can enable internal network communication between the two BMSs and copy resources.

- Security group configuration

  In the same VPC, BMSs associated with the same security group can communicate with one another by default, and no additional configuration is required. However, BMSs in different security groups cannot communicate with each other by default. You must add security group rules to enable the BMSs to communicate with each other through an internal network.

  However, BMSs in different security groups cannot communicate with each other by default. You must add security group rules to enable the BMSs to communicate with each other through an internal network.

Protocol	Direction	Port Range/ ICMP Protocol Type	Source
Protocol to be used for internal network communication. Supported values are **TCP**, **UDP**, **ICMP**, and **All**.	Inbound	Port number range or ICMP protocol type	IPv4 address, IPv4 CIDR block, or another security group ID

Case 2: Only Specified IP Addresses Can Remotely Access BMSs in a Security Group

- Scenario

  To prevent BMSs from being attacked, you can change the port number for remote login and configure security group rules that allow only specified IP addresses to remotely access the BMSs.

- Security group configuration

  To allow IP address **192.168.20.2** to remotely access Linux BMSs in a security group over the SSH protocol and port 22, you can configure the following security group rule.

Protocol	Direction	Port Range	Source
SSH (22)	Inbound	22	IPv4 address, IPv4 CIDR block, or another security group ID For example, 192.168.20.2

Case 3: Remotely Connecting to a Linux BMS Through SSH

- Scenario

  To remotely connect to a Linux BMS through SSH, you need to add a security group rule.

  NOTE

  The default security group comes with this rule. If you use the default security group, you do not need to configure the rule again.

- Security group configuration

Protocol	Direction	Port Range	Source
SSH (22)	Inbound	22	0.0.0.0/0

Case 3: Remotely Connecting to a Windows BMS Through RDP

- Scenario

  To remotely connect to a Windows BMS through RDP, you need to add a security group rule.

  NOTE

  The default security group comes with this rule. If you use the default security group, you do not need to configure the rule again.

- Security group configuration

Protocol	Direction	Port Range	Source
RDP (3389)	Inbound	3389	0.0.0.0/0

Case 5: Pinging a BMS from the Internet

- Scenario

  To ping BMSs from each other to check connectivity, you need to add a security group rule.

- Security group configuration

Protocol	Direction	Port Range	Source
ICMP	Inbound	All	0.0.0.0/0

7.1.3 Changing a Security Group

Scenarios

This section describes how to change the security group of the BMS NIC or associate multiple security groups with the BMS.

📖 **NOTE**

When multiple security groups are associated with the BMS, all the security group rules take effect.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. Click the name of the target BMS.

   The page showing details of the BMS is displayed.
4. Click the **Security Groups** tab. Then, click **Change Security Group**.
5. In the displayed **Change Security Group** dialog box, select the target security group and click **OK**.

   Figure 7-1 Changing a security group
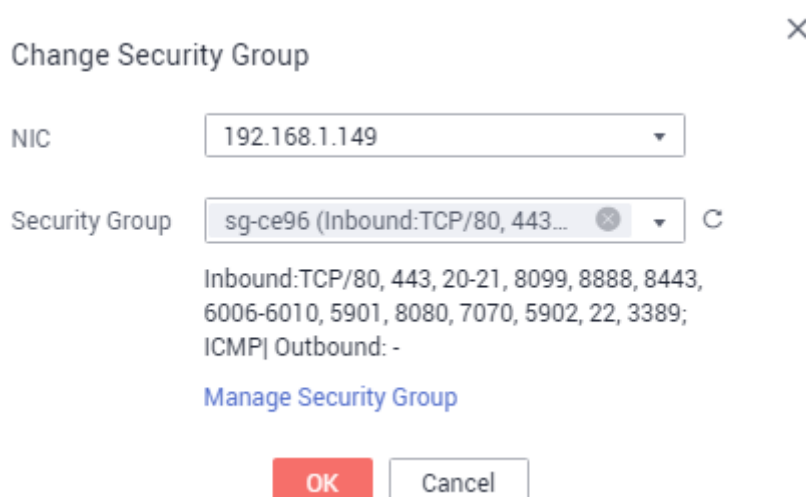
   To associate multiple security groups with the BMS, select the groups.

Result

On the BMS details page, click the **Security Groups** tab. The security group has been changed, or new security groups are contained in the list.

8 Resources and Tags

8.1 Tag

8.1.1 Overview

To facilitate your management of instances, disks, images, and other cloud resources, you can add a tag to each resource to allocate your own metadata to the resource. Tag Management Service (TMS) is a visualized service for fast and unified cross-region tagging and categorization of cloud services.

Basics of Tags

Tags are used to identify cloud resources. When you have many cloud resources of the same type, you can use tags to classify cloud resources by dimension (for example, use, owner, or environment).

Figure 8-1 Example tags

Figure 8-1 shows how tags work. In this example, you assign two tags to each cloud resource. Each tag contains a key and a value that you define. The key of one tag is **Owner**, and the key of another tag is **Use**. Each tag has a value.

You can quickly search for and filter specific cloud resources based on the tags added to them. For example, you can define a set of tags for cloud resources in an account to track the owner and usage of each cloud resource, making resource management easier.

Tag Usage

- BMS-related services that support tags include ECS, IMS, and EVS.

- Each tag consists of a key and a value.

- A BMS can have a maximum of nine tags.

- For each resource, each tag key must be unique and can have only one tag value.

- **Table 8-1** provides the tag key and value requirements.

Table 8-1 Tag key and value requirements

Parameter	Requirement	Example Value
Tag key	- Cannot be left blank. - Can only contain letters, digits, underscores (_), and hyphens (-). - Contains a maximum of 36 characters.	Organization
Tag value	- Cannot be left blank. - Can only contain letters, digits, underscores (_), periods (.), and hyphens (-). - Contains a maximum of 43 characters.	Apache

8.1.2 Adding Tags

Tags are used to identify cloud resources, such as instances, images, and disks. If you have multiple types of cloud resources which are associated with each other, you can add tags to the resources to classify and manage them easily. For more information, see **Overview**.

You can add tags to a BMS in either of the following ways:

- **Add Tags During BMS Creation**

- **Add Tags on the BMS Details Page**

Add Tags During BMS Creation

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Click **Buy BMS**.

4. Configure the BMS parameters.

   Select **Configure now** for **Advanced Settings** and add a tag key and tag value. For the tag key and tag value requirements, see **Table 8-1**.

   📖 **NOTE**

   For details about other parameters, see **Creating a BMS**.

Add Tags on the BMS Details Page

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.
3. In the BMS list, click the name of the target BMS.

   The page showing details of the BMS is displayed.
4. Click the **Tags** tab and then **Add Tag**. In the displayed dialog box, enter the tag key and tag value. For the tag key and tag value requirements, see **Table 8-1**.

   You can change the tag value after the tag is added.

Tips

If you want to add the same tag to multiple BMSs or other resources, you can create a predefined tag on the TMS console and then select the tag for the BMSs or resources. This free you from having to repeatedly enter tag keys and values. The procedure is as follows:

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. In the navigation pane on the left, choose **Predefined Tags**. In the right pane, click **Create Tag**, and enter a key (for example **project**) and a value (for example **A**) in the displayed dialog box.
4. Choose **Service List** > **Computing** > **Bare Metal Server**. On the BMS console, click the name of the BMS to which you want to add the predefined tag.

   The page showing details of the BMS is displayed.
5. Click the **Tags** tab and then **Add Tag**. In the displayed dialog box, select the predefined tag you have created.

8.1.3 Searching for Resources by Tag

After tags are added to resources, you can search for resources by tag using either of the following methods.

Filter Resources in the Resource List

On the BMS console, query BMSs by tag key and value.

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.

   The BMS console is displayed.

3. Click **Search by Tag** above the upper right corner of the BMS list to expand the search area.

4. Enter the tag of the target BMS.

    Both the tag key and value are mandatory. If the tag key or value is matched, the system automatically displays the target BMSs.

5. Click [+] to add multiple tags.

    You can add multiple tags. The system will display BMSs that have all the tags.

6. Click **Search**.

    The system searches for BMSs based on the specified tag keys and values.

Filter Resources on the TMS Console

1. Log in to the management console.

2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.

3. On the **Resource Tags** page, set the search criteria, including **Region**, **Resource Type**, and **Resource Tag**.

4. Click **Search**.

    All the resources that meet the search criteria will be displayed in the **Search Result** area.

8.1.4 Deleting Tags

If you no longer need a tag, delete it in either of the following ways:

Procedure

1. Log in to the management console.

2. Under **Computing**, click **Bare Metal Server**.

    The BMS console is displayed.

3. In the BMS list, click the name of the target BMS.

    The page showing details of the BMS is displayed.

4. Click the **Tags** tab. Locate the row containing the tag to be deleted and click **Delete** in the **Operation** column. In the **Delete Tag** dialog box, click **Yes**.

8.2 Resource Location

Some resources are available in all regions around the globe, while others are only available in specified regions or AZs.

Resource	Type	Description
HUAWEI CLOUD account	Global	You can use the same HUAWEI CLOUD account in all regions.

Resource	Type	Description
Predefined tags	Global	You can use the same predefined tag in all regions.
Key pair	Global or regional	A key pair you create on the management console is associated with the region where it is created. You can create your own RSA key pair and import it into the region where you want to use it. Therefore, you can upload a key pair to each region to use it globally. For details about key pairs, see **Using an SSH Key Pair**.
Resource identifier	Regional	Each resource identifier (such as instance ID, EVS disk ID, and VPC ID) is associated with a region and can be used only in the region where the resource is created.
User-defined resource name	Regional	Each resource name (such as the security name and key pair name) is associated with a region and can be used only in the region where the resource is created. Although you can create resources with the same name in different regions, the resources are not associated with each other.
VPC	Regional	A VPC is associated with a region and can only be associated with instances in the same region.
EIP	Regional	An EIP is associated with a region and can only be associated with instances in the same region.
Security group	Regional	A security group is associated with a region and can only be allocated to instances in the same region. The security group rule cannot be used to enable communication between instances in different regions.
Image	Regional	An image is associated with a region and can only be associated with instances in the same region. The image can be a public, private, or shared image.
Instance	AZ	An instance is associated with an AZ, but the instance ID is associated with a region.
Disk	AZ	A disk is associated with an AZ and can only be attached to instances in the same AZ.
Subnet	AZ	A subnet is associated with an AZ and can only be associated with instances in the same AZ.

8.3 Adjusting Resource Quotas

What Is Quota?

Quotas are enforced for service resources on the platform to prevent unforeseen spikes in resource usage. Quotas can limit the number or amount of resources available to users, such as the maximum number of BMSs or EVS disks that can be created.

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

📖 **NOTE**

The BMS service has no independent quota. It shares the number of instances, vCPUs, and memory with the ECS service. You can view BMS quota in the **Elastic Cloud Server** row.

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.

   Figure 8-2 My Quotas
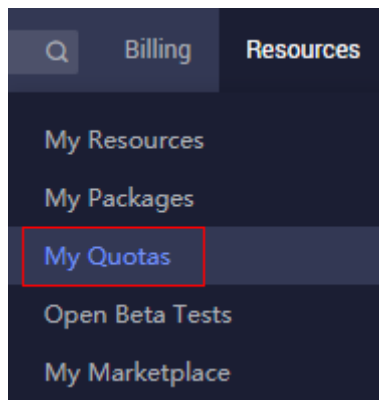
4. View the used and total quota of each type of resources on the displayed page.

   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.

Figure 8-3 My Quotas
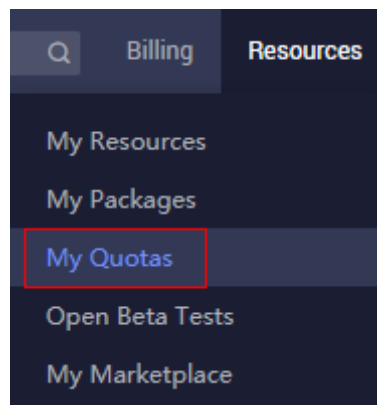
3. Click **Increase Quota**.

4. On the **Create Service Ticket** page, configure parameters as required.

   In **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 Tenant Authorization Letter and Privacy Statement** and click **Submit**.

9 Server Monitoring

9.1 Overview

Server Monitoring

Server monitoring provided by Cloud Eye includes basic monitoring and OS monitoring. Basic monitoring refers to monitoring of ECS metrics automatically reported (BMS does not support basic monitoring). OS monitoring provides system-wide, active, and fine-grained (10-second interval) monitoring for BMSs, on which Agent is installed. Agent occupies less than 1.5% CPU and less than 50 MB memory of the OS.

To meet the basic monitoring and O&M requirements for servers, **Server Monitoring** monitors more than 40 metrics, such as CPU, memory, disk, and network metrics.

Constraints

- Agent can only be installed on BMSs running a 64-bit Linux OS.
- Private images do not support this function.

  Table 9-1 lists the Linux images that support server monitoring.

Table 9-1 Linux images that support server monitoring

OS Type (64-bit)	Version
SUSE	Enterprise11 SP4 and Enterprise12 SP1
CentOS	6.9, 7.2, and 7.3
EulerOS	2.2
Debian	8.6

Installation Methods

After a BMS is created, you need to manually install Agent to monitor the BMS. For details about the installation, see **Installing and Configuring the Agent on a Linux ECS or BMS**.

9.2 Supported Monitoring Metrics (with Agent Installed)

Description

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

📖 **NOTE**

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

Namespace

SERVICE.BMS

Metrics

Supported BMS **OS Monitoring** metrics include CPU metrics listed in **Table 9-2**, CPU load metrics listed in **Table 9-3**, memory metrics listed in **Table 9-4**, disk metrics listed in **Table 9-5**, disk I/O metrics listed in **Table 9-6**, file system metrics listed in **Table 9-7**, NIC metrics listed in **Table 9-8**, software RAID metrics listed in **Table 9-9**, and process metrics in **Table 9-10**.

📖 **NOTE**

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

Currently, BMSs running the Windows OS cannot be monitored.

Table 9-2 CPU metrics

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

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

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

Table 9-3 CPU load metrics

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

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

Table 9-4 Memory metrics

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

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

Table 9-5 Disk metrics

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

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

Table 9-6 Disk I/O metrics

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

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

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

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

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

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

Table 9-7 File system metrics

Metri c ID	Metric	Description	Valu e Rang e	Monit ored Objec t	Monitorin g Interval (Raw Data)
disk_f s_rwst ate	(Agent) File System Read/Write Status	Read and write status of the mounted file system of the monitored object Possible values are **0** (read and write) and **1** (read only). Check file system information in the fourth column in the **/proc/ mounts** file.	0 and 1	BMS	1 minute
disk_i nodes Total	(Agent) Disk inode Total	Total number of index nodes on the disk Run the **df -i** command to check information in the **Inodes** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	BMS	1 minute
disk_i nodes Used	(Agent) Total inode Used	Number of used index nodes on the disk Run the **df -i** command to check data in the **IUsed** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	BMS	1 minute

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

Table 9-8 NIC metrics

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

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

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

Table 9-9 Software RAID metrics

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
md1_status _device:1	(Agent) Status	Software RAID status of the monitored object. Its value is **0** if the RAID is abnormal. Run the plug-in script **/usr/local/ telescope/plugins/raid-monitor.sh** in a collection period. Obtain its value by checking data changes in the **/proc/mdstat** file and run **mdadm -D/dev/md0** (**md0** indicates the RAID name).	0 and 1	BMS	1 minute

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

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

Table 9-10 Process metrics

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

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

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

10 Troubleshooting

10.1 What Do I Do If I Cannot Log In to My BMS or the BMS EVS Disk Is Lost After the BMS Is Started or Restarted?

Issue Description

After a BMS is started or restarted, the user cannot log in to the BMS or the BMS EVS disk is lost.

Possible Causes

The BMS cannot obtain the IP address or the EVS disk cannot be attached to the BMS because packet loss caused by network congestion occurs.

Solution

Restart the BMS. If the fault still exists after you restart the BMS for several times, contact the customer service.

10.2 What Do I Do If a Key Pair Created Using PuTTYgen Cannot Be Imported to the Management Console?

Symptom

When a key pair created using PuTTYgen was imported to the management console, the system displayed a message indicating that importing the public key failed.

Possible Causes

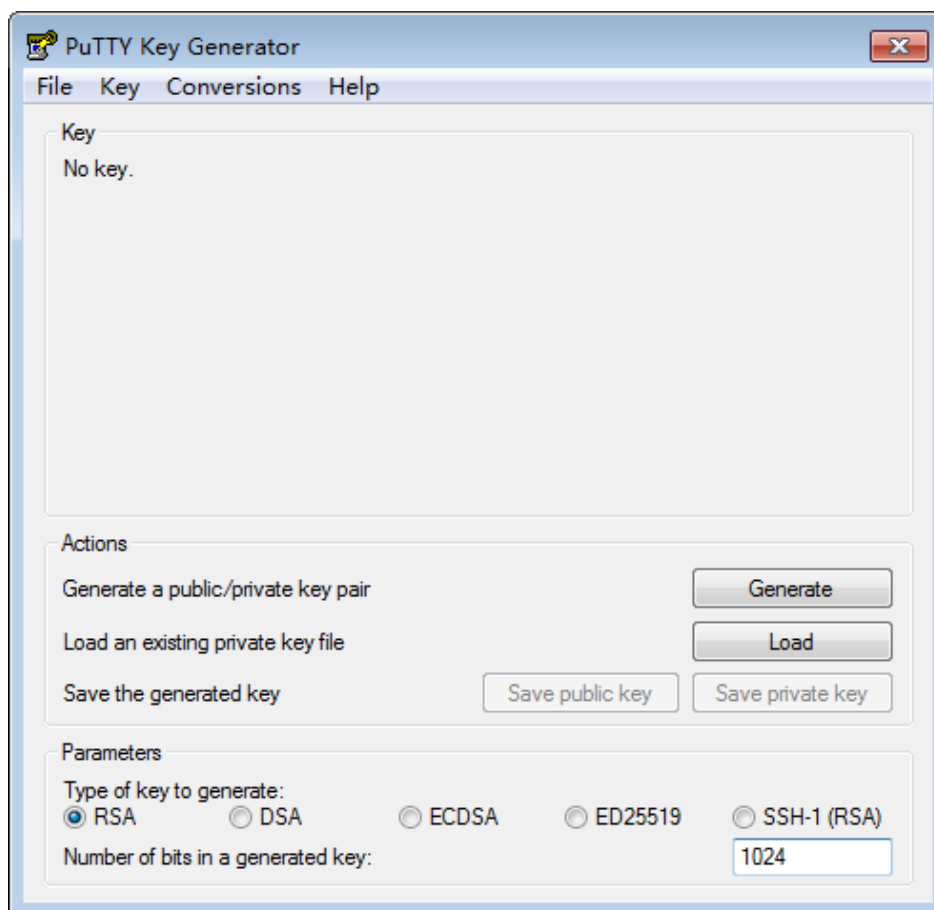The format of the public key content does not meet system requirements.

Storing a public key by clicking **Save public key** of PuTTYgen will change the format of the public key content. Such a key cannot be imported to the management console.

Solution

Use the locally stored private key and **PuTTY Key Generator** to restore the format of the public key content. Then, import the public key to the management console.

1. Double-click **puttygen.exe**. The **PuTTY Key Generator** window is displayed.

   Figure 10-1 PuTTY Key Generator

2. Click **Load** and select the private key.

   The system automatically loads the private key and restores the format of the public key content in **PuTTY Key Generator**. The content in the red box in **Figure 10-2** is the public key with the format meeting system requirements.

Figure 10-2 Restoring the format of the public key content
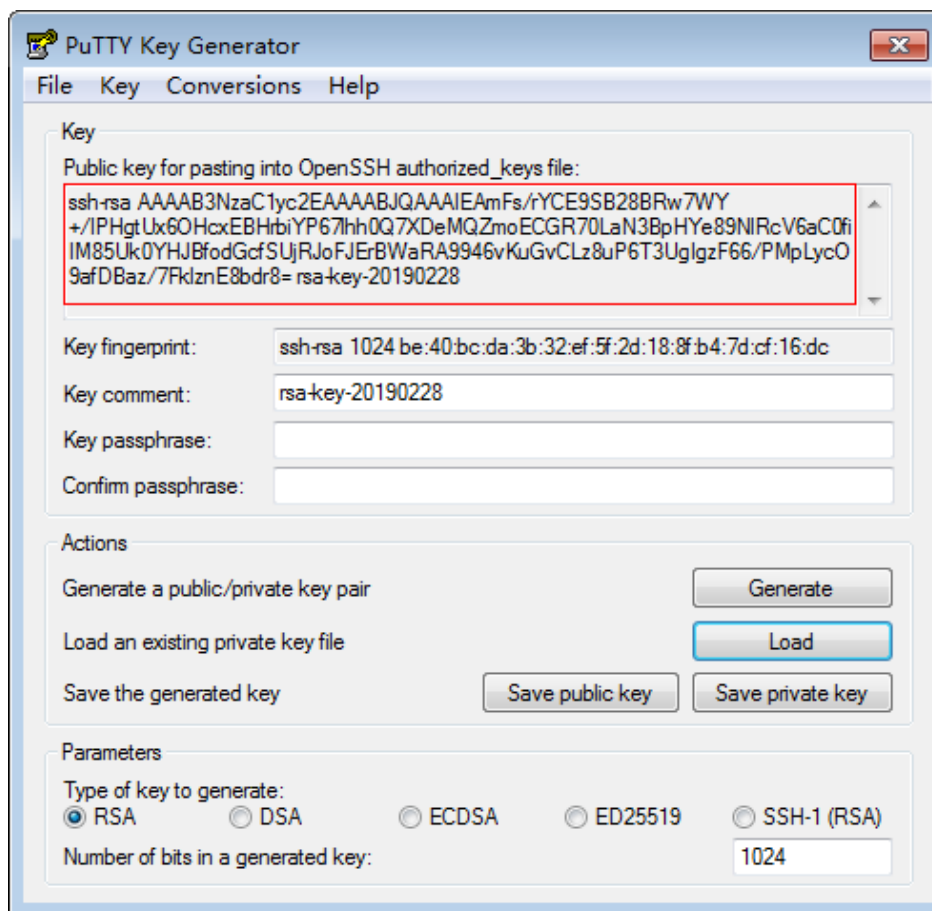
3. Copy the public key content to a .txt file and save the file in a local directory.
4. Import the public key to the management console.

    a. Log in to the management console.

    b. Under **Computing**, click **Bare Metal Server**.

      The BMS console is displayed.

    c. In the navigation tree, choose **Key Pair**.

    d. On the right side of the page, click **Import Key Pair**.

    e. Copy the public key content in the .txt file to **Public Key Content** and click **OK**.

10.3 What Do I Do If Disks Cannot Be Attached to a BMS That Restarts Abnormally?

Symptom

After a BMS provisioned using a local disk with data volumes restarts abnormally, no volume information exists in the BMS OS, and disks cannot be attached to the BMS on the management console.

Abnormal restart indicates that a BMS is powered off and then powered on abnormally, which is not caused by the tenant's operation on the management console.

Solution

Locate the row that contains the BMS, click **More** in the **Operation** column, and select **Restart**. Disks are attached to the BMS automatically after the BMS restarts.

If disks still cannot be attached to the BMS after it is restarted, contact the customer service.

10.4 What Do I Do If an EVS Disk Attached to a Windows BMS Is in Offline State?

Symptom

After an EVS disk is attached to a Windows BMS, start **Control Panel**, choose **System and Security** > **Administrative Tools**, and double-click **Computer Management**. On the **Computer Management** page, choose **Storage** > **Disk Management**. The EVS disk attached to the BMS is in **Offline** state.

Solution

1. Log in to the Windows BMS.

2. Click **Start**, enter **cmd** in **Search programs and files**, and press **Enter** to open the command-line interface (CLI).

3. Type **diskpart**.
    C:\Users\Administrator>**diskpart**

4. Type **san**.
    DISKPART> **san**
    SAN Policy: Online All

5. Type **san policy=onlineall**.
    DISKPART> **san policy=onlineall**
    DiskPart successfully changed the SAN policy for the current operating system

6. Type **list disk** to display all disks of the BMS.
    DISKPART> **list disk**
    Disk ### Status   Size   Free   Dyn   Gpt
    Disk 0   Online   838 GB   0B
    Disk 1   Offline 838 GB   838 GB
    Disk 2   Offline 838 GB   838 GB
    Disk 3   Offline 838 GB   838 GB
    ...

7. Type **select disk** *num*. *num* indicates the disk number. Replace it with the specific disk number.
    DISKPART> **select disk** *4*

8. Type **attributes disk clear readonly**.
    DISKPART> **attributes disk clear readonly**
    DiskPart succeed to clear disk attributes.

9. Type **online disk**.
    DISKPART> **online disk**
    DiskPart succeed to make the selected disk online.

10. After the modification, format the EVS disk.

A Change History

Released On	Description
2020-09-10	This issue is the eighth official release. Added the precautions for using the IB network in **Overview**.
2020-01-20	This issue is the seventh official release. Added **Tag**.
2019-07-30	This issue is the sixth official release. Added the following content: ● **Creating a Dedicated BMS** ● **Changing the Name of a BMS** ● **Private Image Overview** ● **Overview** Modified the following content: ● Adjusted the outline of **Instance**. ● Added a table displaying the requirements and differences between login methods in **Linux BMS Login Methods**. ● Added suggestions for using security groups in **Adding Security Group Rules**.
2019-05-30	This issue is the fifth official release. ● Added descriptions about virtualization in **Common Operations**. ● Added **User-defined Network ACL**.

Released On	Description
2019-04-25	This issue is the fourth official release. ● Added the guide to changing the SID of a Windows Server 2012 BMS in **Creating a BMS**. ● Added verification operations in **Resetting the BMS Password**. ● Optimized descriptions in **Metadata**.
2019-03-18	This issue is the third official release. Optimized the whole document, including adjusting the outline, optimizing feature descriptions, and adding scenario descriptions.
2018-10-31	This issue is the second official release. ● Added **Deleting the Password of a Windows BMS**. ● Added the method of automatically updating the BMS host name in **How Do I Configure the Static Host Name of a BMS?**
2018-06-30	This issue is the first official release.