



**Bare Metal Server**

# **User Guide**

**Date**      **2023-03-30**

---

# Contents

---

<b>1 Overview.....</b>	<b>1</b>
1.1 What Is BMS?.....	1
1.2 BMS Advantages.....	3
1.3 Application Scenarios.....	4
1.4 Image.....	4
1.4.1 Image.....	5
1.5 EVS Disk.....	7
1.6 Network.....	8
1.7 Security.....	10
1.7.1 License Type.....	10
1.7.2 Cloud-Init.....	11
1.7.3 Identity Authentication and Access Control.....	12
1.8 Billing.....	14
1.8.1 Billing.....	15
1.9 Region and AZ.....	18
1.10 Related Services.....	19
1.11 Features and Constraints.....	19
<b>2 Getting Started.....</b>	<b>21</b>
2.1 Quick Start.....	21
2.2 Making Preparations.....	22
2.3 Step 1: Create a BMS.....	22
2.4 Step 2: Log In to the BMS.....	24
2.5 Step 3: Deploy an Application.....	24
2.6 Step 4: Release the BMS.....	25
<b>3 Instance.....</b>	<b>27</b>
3.1 Creating a BMS.....	27
3.1.1 Introduction.....	27
3.1.2 Creating a Common BMS.....	27
3.1.3 Creating a BMS Supporting Quick Provisioning.....	34
3.1.4 Creating a BMS from a Private Image.....	34
3.2 Viewing BMS Information.....	35
3.2.1 Viewing BMS Creation Statuses.....	35

3.2.2 Viewing BMS Details.....	36
3.3 Logging In to a Linux BMS.....	36
3.3.1 Linux BMS Login Methods.....	36
3.3.2 Remotely Logging In to a BMS.....	37
3.3.3 Logging In to a BMS Using an SSH Key Pair.....	38
3.3.4 Logging In to a BMS Using an SSH Password.....	41
3.4 Logging In to a Windows BMS.....	42
3.4.1 Windows BMS Login Methods.....	42
3.4.2 Logging In to a BMS Remotely Using MSTSC.....	42
3.5 Managing BMSs.....	43
3.5.1 Changing the Name of a BMS.....	43
3.5.2 Stopping a BMS.....	44
3.5.3 Restarting a BMS.....	44
3.5.4 Reinstalling the OS.....	45
3.5.5 Rebuilding a BMS.....	47
3.5.6 Backing Up a BMS.....	48
3.5.7 Releasing a BMS.....	49
3.6 User Data and Metadata.....	50
3.6.1 Injecting User Data.....	50
3.6.2 Retrieving Metadata.....	57
<b>4 Image.....</b>	<b>65</b>
4.1 Private Image Overview.....	65
4.2 Creating a Private Image from a BMS.....	66
4.3 Creating a Private Image from an External Image File.....	67
<b>5 Disk.....</b>	<b>70</b>
5.1 Disk Types.....	70
5.2 Attaching Data Disks.....	72
5.3 Initializing Data Disks.....	73
5.3.1 Introduction to Data Disk Initialization Scenarios and Partition Styles.....	73
5.3.2 Initializing a Windows Data Disk (Windows Server 2016).....	75
5.3.3 Initializing a Linux Data Disk (fdisk).....	86
5.3.4 Initializing a Linux Data Disk (parted).....	91
5.3.5 Initializing a Windows Data Disk Greater Than 2 TB (Windows Server 2012).....	95
5.3.6 Initializing a Linux Data Disk Greater Than 2 TB (parted).....	103
5.4 Detaching a Disk.....	108
5.5 Expanding Disk Capacity.....	108
<b>6 Key Pair and Password.....</b>	<b>110</b>
6.1 Using an SSH Key Pair.....	110
6.2 Obtaining the Password of a Windows BMS.....	114
6.3 Deleting the Password of a Windows BMS.....	115
<b>7 Network.....</b>	<b>116</b>

7.1 EIP.....	116
7.1.1 Overview.....	116
7.1.2 Binding an EIP to a BMS.....	117
7.1.3 Unbinding an EIP from a BMS.....	118
7.2 VPC.....	118
7.2.1 Overview.....	118
7.2.2 Binding a Virtual IP Address to a BMS.....	120
7.2.3 Setting the Source/Destination Check for a NIC.....	121
7.3 High-Speed Network.....	122
7.3.1 Overview.....	122
7.3.2 Managing High-Speed Networks.....	124
7.4 Enhanced High-Speed Network.....	125
7.4.1 Overview.....	126
7.4.2 Adding an Enhanced High-Speed NIC.....	128
7.4.3 Deleting an Enhanced High-Speed NIC.....	129
7.4.4 Configuring an Enhanced High-Speed NIC (SUSE Linux Enterprise Server 12).....	130
7.4.5 Configuring an Enhanced High-Speed NIC (SUSE Linux Enterprise Server 11).....	133
7.4.6 Configuring an Enhanced High-Speed NIC (Red Hat, CentOS, Oracle Linux, and EulerOS).....	136
7.4.7 Configuring an Enhanced High-Speed NIC (Ubuntu).....	143
7.4.8 Configuring an Enhanced High-Speed NIC (Windows Server).....	149
7.5 User-defined VLAN.....	154
7.5.1 Overview.....	154
7.5.2 Configuring a User-defined VLAN (SUSE Linux Enterprise Server 12).....	155
7.5.3 Configuring a User-defined VLAN (SUSE Linux Enterprise Server 11).....	161
7.5.4 Configuring a User-defined VLAN (Red Hat, CentOS, Oracle Linux, and EulerOS).....	166
7.5.5 Configuring a User-defined VLAN (Ubuntu).....	169
7.5.6 Configuring a User-defined VLAN (Windows Server).....	173
7.6 IB Network.....	178
7.6.1 Overview.....	178
<b>8 Security.....</b>	<b>180</b>
8.1 Security Group.....	180
8.1.1 Adding Security Group Rules.....	180
8.1.2 Security Group Configuration Examples.....	181
8.1.3 Changing a Security Group.....	183
<b>9 Resources and Tags.....</b>	<b>185</b>
9.1 Tag.....	185
9.1.1 Overview.....	185
9.1.2 Adding Tags.....	186
9.1.3 Searching for Resources by Tag.....	187
9.1.4 Deleting Tags.....	187
9.2 Resource Location.....	188
9.3 Adjusting Resource Quotas.....	189

<b>10 Server Monitoring</b>	<b>191</b>
10.1 Overview	191
10.2 Installing and Configuring Agent	192
10.3 Monitored Metrics (with Agent Installed)	193
<b>11 Troubleshooting</b>	<b>198</b>
11.1 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?	198
11.2 What Do I Do If a Key Pair Created Using PuTTYgen Cannot Be Imported to the Management Console?	198
11.3 What Do I Do If Disks Cannot Be Attached to a BMS That Restarts Abnormally?	200
11.4 What Do I Do If an EVS Disk Attached to a Windows BMS Is in Offline State?	201
<b>12 FAQs</b>	<b>203</b>
12.1 General FAQ	203
12.1.1 What Are the Restrictions on Using BMSs?	203
12.1.2 How Are BMSs Different from ECSs?	204
12.1.3 What Are the Differences Between BMSs and Traditional Physical Servers?	204
12.1.4 What Are the Differences Between BMS and Dedicated Host (DeH)?	204
12.1.5 How Do BMSs Ensure Data Security?	205
12.1.6 Can I Use APIs to Access BMSs?	205
12.1.7 Will Services Be Affected If Hyper-Threading Is Configured for BMS?	205
12.1.8 How Do I View and Increase the BMS Quota?	205
12.2 Instance FAQ	206
12.2.1 How Long Does It Take to Create a BMS?	206
12.2.2 What Do I Do If I Cannot Find My BMS on the Management Console?	206
12.2.3 How Can I Obtain board_type of a BMS Flavor?	207
12.2.4 Why Is the BMS Creation Task Displayed as Failed But the BMS List Shows the BMS?	208
12.2.5 How Do I Create a BMS That Can be Quickly Provisioned?	208
12.2.6 What Are the Advanced Features of BMSs Using EVS Disks?	208
12.2.7 Can I Transfer a BMS to Another Account?	209
12.2.8 Is the BMS Host Name with Suffix novalocal Normal?	209
12.2.9 How Do I Monitor BMS Metrics?	210
12.2.10 How Can I Check the BMS Monitoring Status?	210
12.2.11 How Do I Create an Agency for Server Monitoring of the BMS?	210
12.3 Key Pair and Password FAQ	210
12.3.1 How Do I Change the Password of a BMS in Its OS?	211
12.4 Login FAQ	213
12.4.1 What Need to Be Prepared for Logging In to a BMS?	213
12.4.2 What Do I Do If I Have Selected the Key Pair Authentication When I Created a BMS But Want to Log In to the BMS Using a Password?	214
12.4.3 What Do I Do If I Cannot Log In to a Windows BMS?	214
12.4.4 What Do I Do If I Cannot Log In to a Linux BMS?	215
12.4.5 What Browser Versions Can Be Used to Remotely Log In to a BMS?	218

12.4.6 What Do I Do If the Login Page Does Not Respond?.....	218
12.4.7 What Do I Do If the BMS Console Is Displayed Improperly After I Remotely Log In to a BMS?.....	220
12.4.8 What Do I Do If the Numeric Keypad Does Not Work During Remote Login?.....	221
12.4.9 What Do I Do If the SSH Login or Data Transmission Is Slow?.....	222
12.5 Network and Security FAQ.....	222
12.5.1 Can BMSs of Different Accounts Communicate with Each Other over an Internal Network?.....	223
12.5.2 How Do Two BMSs in the Same Region But Different AZs Communicate with Each Other?.....	223
12.5.3 Are My BMSs in the Same Subnet?.....	223
12.5.4 Can I Associate a BMS with Multiple Security Groups?.....	223
12.5.5 Can BMSs Communicate with ECSs in the Same VPC?.....	223
12.5.6 What Are the Differences Between the Primary and Extension NICs of BMSs?.....	223
12.5.7 Can I Bind Multiple EIPs to a BMS?.....	224
12.5.8 Can I Configure the EIP?.....	224
12.5.9 Will I Obtain an EIP That Has Been Released?.....	224
12.5.10 What Are the Differences Between EIPs, Private IP Addresses, and Virtual IP Addresses?.....	224
12.5.11 How Can I Modify the Network Configuration or Restart the Network If I Can Log In to a BMS Using Only SSH?.....	224
12.5.12 What Do I Do If the Communication Between the Primary NIC and Extension NIC of the BMS is Abnormal?.....	224
12.5.13 How Can I Configure a Static IP Address for a BMS?.....	225
12.5.14 How Do I Configure the DNS Server?.....	227
12.5.15 How Do I Change the MTU Value of a Linux BMS NIC?.....	230
12.5.16 How Do I Change the MTU Value of a Windows BMS NIC?.....	232
12.6 Disk FAQ.....	235
12.6.1 Can EVS Disks Be Attached to BMSs? .....	235
12.6.2 What Are the Restrictions for Attaching a Disk to a BMS?.....	236
12.6.3 How Do I Know Whether EVS Disks Are Available in a Flavor?.....	236
12.6.4 How Do I Change the Disk Identifier in the fstab file to UUID?.....	236
12.6.5 How Do I Obtain the Drive Letter of an EVS Disk?.....	237
12.6.6 Are the EVS Disk Device Names on the Console and the Device Names in BMS OSs Consistent?.....	237
12.6.7 Why Is the EVS Disk Size Not Updated in the BMS OS After the EVS Disk Capacity Has Been Expanded?.....	241
12.6.8 How Can I Restore System Disk Data Using the Snapshot?.....	241
12.6.9 What Do I Do to Prevent Risks of Attaching or Detaching the System Disk?.....	241
12.6.10 How Should I Select Storage?.....	242
12.6.11 Why Is the Disk Capacity Displayed in the BMS OS Less Than That Displayed on the Official Website?.....	242
12.7 OS FAQ.....	242
12.7.1 Can I Install or Upgrade BMS OSs by Myself?.....	242
12.7.2 Can the BMS OS Be Replaced?.....	242
12.7.3 Is a GUI Provided for BMS OSs?.....	242
12.7.4 Is an Upload Tool Delivered with BMS OSs?.....	242
12.7.5 How Do I Configure the Static Host Name of a BMS?.....	243

---

12.7.6 How Do I Set the Password Validity Period?.....	245
12.7.7 How Do I Set SSH Configuration Items?.....	246
12.7.8 How Can I Handle the Eight-Hour Difference Between the Windows BMS and Local Time.....	247
12.7.9 How Can I Activate a Windows BMS?.....	248
12.7.10 How Do I Change the SID of a Windows Server 2012 BMS?.....	249
12.7.11 How Do I Change the Kernel Version of CentOS 7 BMSs?.....	251
12.7.12 How Do I Reserve Log Space If the Root Partition Automatically Expands Disks?.....	252
12.7.13 How Do I Roll Back the Kernel Version If I Mistakenly Upgrade the Kernel?.....	255
12.7.14 How Do I Increase the Swap Partition Size?.....	256
12.7.15 How Do I Increase the Size of the Root Partition of a BMS Which Is Quickly Provisioned?.....	257
12.7.16 Common Linux Commands.....	259
12.7.17 How Do I Update the Disk Metadata After the LVM Volume Is Remounted?.....	260
12.7.18 How Do I Handle a Network Failure After Services Are Switched from a Windows BMS Booted from an EVS Disk to an HA BMS?.....	260
<b>A Change History.....</b>	<b>263</b>

# 1 Overview

---

## 1.1 What Is BMS?

### Overview

A Bare Metal Server (BMS) features both the scalability of Elastic Cloud Servers (ECSs) and high performance of physical servers. It provides dedicated servers on the cloud, delivering the performance and security required by core databases, critical applications, high-performance computing (HPC), and Big Data.

The BMS self-service feature allows you to apply for and use a BMS on demand. To apply for a BMS, you need to specify the server type, image, required network, and other configurations. You can obtain the BMS you require within 30 minutes.

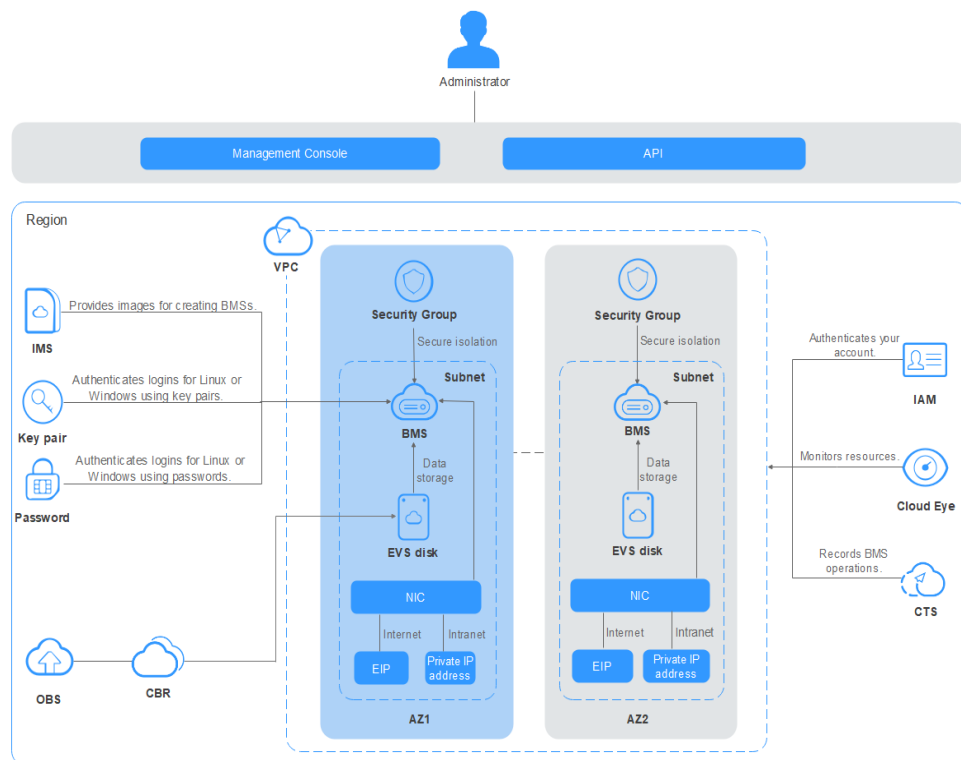
### System Architecture

BMS works with other cloud services to provide computing, storage, network, and image functions.

- BMSs are deployed in multiple availability zones (AZs) connected with each other through an internal network. If an AZ becomes faulty, other AZs in the same region will not be affected.
- Virtual Private Cloud (VPC) allows you to create a dedicated network for BMSs and configure subnets and security groups. BMSs in a VPC can communicate with the external network through EIPs (bandwidth support required).
- Image Management Service (IMS) allows you to install OSs on BMSs or create BMSs using private images for rapid service deployment.
- Elastic Volume Service (EVS) provides storage and Volume Backup Service (VBS) provides data backup and restoration.
- Cloud Eye is a key measure to monitor BMS performance, reliability, and availability. Using Cloud Eye, you can monitor BMS in real time.
- Cloud Backup and Recovery (CBR) backs up data for EVS disks and BMSs, and uses snapshot backups to restore the EVS disks and BMSs.



**Figure 1-1** System architecture



## BMSs, Physical Servers, and ECSs

**Table 1-1** compares BMSs, physical servers, and ECSs. Y indicates supported and N indicates unsupported.

### NOTE

BMSs have all the features and advantages of physical servers. Your applications can access the physical CPU and memory without any virtualization overhead.

**Table 1-1** Comparison between BMSs, physical servers, and ECSs

Category	Function	BMS	Physical Server	ECS
Provisioning	Automatic provisioning	Y	N	Y
Compute	No feature loss	Y	Y	N
	No performance loss	Y	Y	N
	Exclusive resources	Y	Y	N
Storage	Local storage	Y	Y	N

Category	Function	BMS	Physical Server	ECS
	Booting from an EVS disk (system disk)	Y	N	Y
	Using an image (free from OS installation)	Y	N	Y
Network	VPC	Y	N	Y
	Communication between physical servers and VMs through a VPC	Y	N	Y
Management and control	Consistent remote login experience as VMs	Y	N	Y
	Monitoring and auditing of key operations	Y	N	Y

## Access Methods

The public cloud provides a web-based service management system (management console). You can access BMS through the management console or HTTPS APIs. The two access methods differ as follows:

- API  
If you want to integrate BMS into a third-party system for secondary development, use APIs to access the BMS service.
- Management console  
For all other purposes, use the management console.

## 1.2 BMS Advantages

### High Security and Reliability

BMS allows you to use dedicated computing resources, add servers to VPCs and security groups for network isolation, and integrate related components for server security. BMS can interconnect with dedicated storage to ensure the data security and reliability required by enterprise services.

## High Performance

BMS has no virtualization overhead, allowing dedicated computing resources for service running. BMS can use high-bandwidth, low-latency storage and networks on the cloud, meeting the deployment density and performance requirements of critical services such as enterprise databases, big data, containers, HPC, and AI.

## Quick Provisioning and Unified O&M

The required BMSs can be provisioned within minutes after you submit an order. You can manage your BMSs through their lifecycle from the management console or using open APIs with SDKs.

## Quick Integration of Cloud Services and Solutions

Based on the unified VPC model, cloud services and cloud solutions (such as database, big data, container, HPC, and AI solutions) can be quickly integrated to run on BMSs. This accelerates cloud transformation.

# 1.3 Application Scenarios

## Database

Mission-critical database services of governments and financial institutions must be deployed on physical servers with dedicated resources, isolated networks, and guaranteed performance. The BMS service properly meets these database service requirements by providing high-performance servers dedicated for individual users.

## Big Data

For Internet services involving big data storage and analysis, the BMS service provides both local storage and compute-storage decoupling backed by OBS.

## Container

Containers enable elastic load balancing for Internet services. BMSs provide more agile container deployment with higher density and lower resource overhead than VMs. Cloud native technologies reduce the cost of cloud transformation.

## HPC/AI

In high-performance computing (HPC) such as supercomputing, DNA sequencing, and AI, a large amount of data needs to be processed. The BMS service meets the requirements of HPC services for high computing performance, high stability, and high real-time performance of servers.

# 1.4 Image

## 1.4.1 Image

### What Is an Image?

An image is a template of the BMS running environment. It contains an OS and runtime environment, and some pre-installed applications. An image file is equivalent to a copy file that contains all data in the system disk.

### Image Types

Images can be classified into public images, private images, and shared images.

**Table 1-2** Image types

Image Type	Description
Public image	A public image is provided by the cloud platform and is available to all users. It contains an OS and preinstalled public applications.
Private image	A private image is created by a user and is available only to the user who created it. It contains an OS, pre-installed public applications, and the user's private applications. Using a private image to create BMSs frees you from repeatedly configuring BMSs.
Shared image	A shared image is a private image other users share with you.

### Public Images

Public images are provided by the system. These images are available to all users, compatible with BMSs and most mainstream OSs, and are pre-installed with necessary plug-ins. Public images available to you vary depending on the BMS flavor you selected.

#### Characteristics

- OS types: Linux and Windows OSs that are updated and maintained periodically
- Pre-installed software: plug-ins that BMS storage, networks, and basic functions depend on



These plug-ins are necessary for BMSs to run properly. Do not delete or modify any of them. Otherwise, basic BMS functions will be affected.

---

**Table 1-3** Pre-installed software

Software	Description
Cloud-Init	Cloud-Init is an open-source cloud initialization program, which initializes specific configurations, such as the host name, key, and user data, of a newly created BMS.
bms-network-config	This plug-in is used to automatically configure BMS networks during BMS provisioning and restore the BMS network when the network is interrupted due to faults.
SDI iNIC frontend driver plug-in	This plug-in is installed in the image so that EVS disks can be attached to BMSs. In this way, BMSs can be booted from EVS disks, facilitating quick BMS provisioning.

- Compatibility: compatible with server hardware
- Security: highly stable and licensed
- Restrictions: no restrictions on usage

## Private Images

A private image contains an OS, preinstalled public applications, and a user's private applications. You can use a private image to create BMSs without having to repeatedly configure them.

### Characteristics

- Compatibility: Private images can be used to deploy servers that are of the same model as the source BMS and may fail to deploy servers of other models.
- Functions: You can create and delete private images, as well as create BMSs and reinstall the BMS OS using private images.
- Restrictions: You can create a maximum of 50 private images.

## Shared Images

A shared image is a private image other users share with you.

## Application Scenarios

- Deploying software environments in a batch  
Prepare a BMS with an OS, the partition arrangement you prefer, and software installed to create a private image. You can use the image to create batch clones of your custom BMS.
- Backing up a BMS  
Create an image from a BMS to back up the BMS. If the software of the BMS becomes faulty, you can use the image to restore the BMS.

## 1.5 EVS Disk

### What Is Elastic Volume Service (EVS)?

EVS offers scalable block storage for BMSs. EVS disks feature high reliability, high performance, and rich specifications, and are ideal for distributed file systems, development and test environments, data warehouse applications, and high-performance computing (HPC) scenarios.

Unlike traditional servers that can only use local disks, BMSs can use EVS disks that are not constrained by capacity. Shared EVS disks allow concurrent reads and writes by multiple BMSs, enabling you to deploy core applications in clusters.

### EVS Disk Types

BMSs support the following types of EVS disks:

- **Common I/O:** This EVS disk type delivers a maximum of 2200 IOPS. It is ideal for application scenarios that require large capacity, medium read/write speed, and fewer transactions, such as enterprise applications and small-scale testing.
- **High I/O:** This EVS disk type delivers a maximum of 5,000 IOPS and a minimum of 1 ms read/write latency. It is designed to meet the needs of mainstream high-performance, high-reliability application scenarios, such as enterprise applications, large-scale development and testing, and web server logs.
- **Ultra-high I/O:** This EVS disk type delivers a maximum of 33,000 IOPS and a minimum of 1 ms read/write latency. It is excellent for ultra-high I/O, ultra-high bandwidth, and read/write-intensive application scenarios, such as distributed file systems in HPC or NoSQL/RDS in I/O-intensive scenarios.
- **Extreme SSD:** This EVS disk type delivers up to 128,000 IOPS and sub-millisecond read latencies. With RDMA integrated with low-latency congestion control algorithms, this disk type is suitable for application scenarios that require ultra-high bandwidth and ultra-low latency.

### EVS Disk Performance

The key indicators of EVS disk performance include read/write latency, IOPS, and throughput.

- **IOPS:** number of read/write operations performed by an EVS disk per second
- **Throughput:** amount of data successfully transmitted by an EVS disk per second, that is, the amount of data read from and written into an EVS disk
- **Read/write latency:** minimum interval between two consecutive read/write operations of an EVS disk

For more details, see *Elastic Volume Service User Guide*.

### EVS Disk Device Types

BMS supports only Small Computer System Interface (SCSI) EVS disks.

On the management console, you can create EVS disks with **Device Type** set to **SCSI**. The EVS disks support transparent SCSI command transmission, allowing BMS OSs to directly access underlying storage media. The EVS disks support basic read/write SCSI commands and advanced SCSI commands.

**NOTE**

BMS public image OSs are preinstalled with the driver required to use SCSI disks, so you do not need to install the driver. To know how to install the driver, see "Installing the SDI Card Driver" in *Bare Metal Server Private Image Creation Guide*.

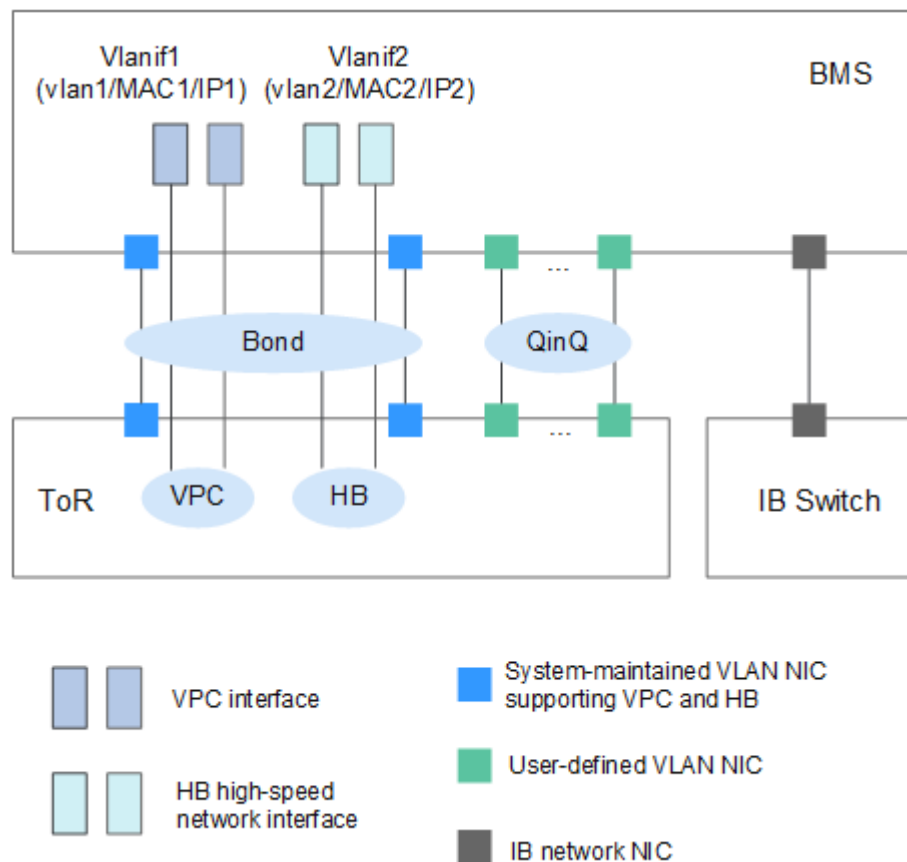
## 1.6 Network

BMS provides five types of networks: VPC, high-speed network, enhanced high-speed network, user-defined VLAN, and IB network. They are isolated from each other.

VPC and high-speed network interfaces are VLAN sub-interfaces created after system maintenance VLAN NICs are bonded. You can manage and configure NICs of user-defined VLANs and IB networks.

The QingTian architecture improves the network performance and VPC NIC bonding provides high reliability.

**Figure 1-2** BMS networks



 **NOTE**

- In the preceding figure, ToR indicates the cabling mode in the server cabinet. The access switch is placed on top of the rack and the server is placed beneath it. HB indicates a high-speed network. QinQ indicates an 802.1Q tunnel.
- VPC and high-speed network interfaces are generated by the system and cannot be changed. They are configured in the same NIC bond.
- BMSs can communicate with ECSs through VPCs or IB networks (if any).
- Only VPC supports security groups, EIPs, and ELB.
- For a high-speed network and user-defined VLAN, BMSs in the same network communicate with each other only through layer-2 connections.

## VPC

A VPC is a logically isolated, configurable, and manageable virtual network. It helps improve the security of cloud resources and simplifies network deployment. You can create security groups and VPNs, configure IP address ranges, and specify bandwidth sizes in your VPC. With a VPC, you can easily manage and configure internal networks and change network configurations. You can also customize access rules to control BMS access within a security group and across different security groups to enhance BMS security.

For more information, see *Virtual Private Cloud User Guide*.

## High-Speed Network

A high-speed network is an internal network between BMSs. It provides high bandwidth for connecting BMSs in the same AZ. If you want to deploy services that require high throughput and low latency, you can create high-speed networks. Currently, the BMS service supports high-speed networks with a maximum bandwidth of 10 Gbit/s.

Enhanced high-speed networks use upgraded hardware and software and provide performance superior to high-speed networks.

Enhanced high-speed networks have the following advantages:

- The bandwidth is 10 Gbit/s or higher.
- The number of network planes can be customized and a maximum of 4000 subnets are supported.

## User-defined VLAN

You can use the 10GE Ethernet NICs that are not being 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. User-defined VLAN NICs are in pairs. You can configure NIC bonding to achieve high availability. User-defined VLANs in different AZs cannot communicate with each other.



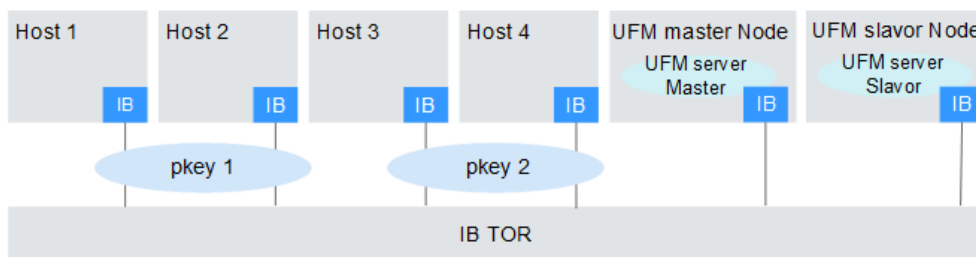
 **NOTE**

QinQ is a layer 2 tunnel protocol based on IEEE 802.1Q encapsulation. It adds a public VLAN tag to a frame with a private VLAN tag and allows the frame with double VLAN tags to be transmitted over the service provider's backbone network based on the public VLAN tag. This provides a layer 2 VPN tunnel for customers.

## IB Network

An 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 an Ethernet).

**Figure 1-3** IB network isolation



 **NOTE**

Unified Fabric Manager (UFM) is the IB switch controller of an IB network based on OpenSM software and provides northbound interfaces. It is deployed in active/standby mode.

## 1.7 Security

### 1.7.1 License Type

#### Use License from the System

You can use OS licenses provided by the cloud platform. You need to pay for the licenses which are billed on a pay-as-you-go basis. The licenses are managed by the cloud platform.

#### Bring Your Own License (BYOL)

##### What Is BYOL?

Bring Your Own License (BYOL) allows you to use your own OS licenses. You do not need to pay for the licenses but need to manage them by yourself.

##### How Can I Use BYOL?

If you choose BYOL, you need to manage licenses by yourself. The cloud platform provides functions you need for maintaining license compliance during the lifecycle of your license.

## Application Scenarios

You can choose BYOL when you create a BMS.

The system will not allow you to change the license type after you create the BMS or when you reinstall its OS.

## 1.7.2 Cloud-Init

### What Is Cloud-Init?

Cloud-Init is an open-source cloud initialization program, which initializes specific configurations, such as the host name, key, and user data, of a newly created BMS.

By default, Cloud-Init has been installed for all public images.

### Impact on IMS

To ensure that BMSs you create using private images support customized configurations, you must install Cloud-Init or Cloudbase-Init when you create private images.

- For Windows OSs, download and install Cloudbase-Init.
- For Linux OSs, download and install Cloud-Init.

After Cloud-Init or Cloudbase-Init is installed in an image, Cloud-Init or Cloudbase-Init will automatically initialize the BMS that created from the image. For details about how to install Cloud-Init and Cloudbase-Init, see *Bare Metal Server Private Image Creation Guide*.

### Impact on BMS

- When you create a BMS, if Cloud-Init has been installed in the image you select, you can initialize the BMS by injecting customized configurations (such as the BMS login password) into it. For details, see [Injecting User Data](#).
- For a BMS with Cloud-Init installed, you can view the BMS metadata and configure and manage the BMS. For more information, see [Retrieving Metadata](#).

### Notes

- If Cloud-Init has been installed, enable DHCP in the VPC to which the BMS belongs.
- If Cloud-Init has been installed, ensure that security group rules in the outbound direction meet the following requirements so that you can access the metadata service:
  - Protocol: TCP
  - Port Range: 80
  - Destination: 169.254.0.0/16

 **NOTE**

If you use the default security group rule in the outbound direction, the preceding requirements have been met. The default outbound security group rule is as follows:

- Protocol: ANY
- Port Range: ANY
- Destination: 0.0.0.0/16

## 1.7.3 Identity Authentication and Access Control

Identity and Access Management (IAM) provides functions such as user identity authentication, permission assignment, and access control. You can use IAM to securely control user access to your BMSs. IAM permissions define which actions on your cloud resources are allowed or denied. After creating an IAM user, add it to a user group and grant the permissions required by BMS to the user group. Then, all users in this group will be automatically granted with these permissions.

### Account Security

If you are an enterprise administrator, you can use IAM to create a user and grant permissions to the user. Enterprise employees can use the user account to access the system, and you do not need to share your account password or key pair with them. This helps you manage resources efficiently. You can also configure account security policies to protect these user accounts and reduce security risks for your enterprise information.

### Fine-Grained Authorization

You can grant refined permissions to employee accounts to ensure that cloud services are properly used.

### Security Group

A security group is a virtual firewall that detects status and filters data packets. It is an important network isolation method used for access control of ECSs, BMSs, load balancers, and databases.

You can configure security group rules to allow instances in a security group to access the public or private network.

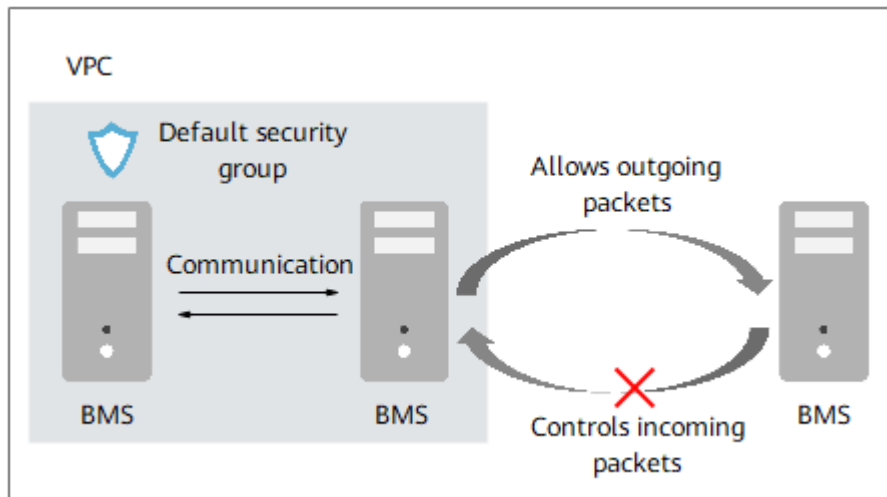
- A security group is a logical group. You can add BMSs that have the same security protection requirements within a region to the same security group.
- By default, BMSs in the same security group can communicate with each other through an internal network, whereas BMSs in different security groups cannot.
- You can modify a security group rule at any time, and the modification takes effect immediately.

### Default Security Group

When you create a BMS in a region, the system will create a default security group if there is no security group in the region.

The default security group rule allows all outgoing data packets and blocks incoming data packets. BMSs in this security group can access each other already. You do not need to add additional rules.

**Figure 1-4** Default security group



**Table 1-4** describes the rules of the default security group.

**Table 1-4** Rules in the default security group

Direction	Protocol	Port/Range	Source/Destination	Description
Outbound	All	All	Destination: 0.0.0.0/0	Allows all outbound traffic.
Inbound	All	All	Source: the current security group (for example, sg-xxxxx)	Allows communications among BMSs within the security group and denies all inbound traffic (incoming data packets).
Inbound	TCP	22	Source: 0.0.0.0/0	Allows all IP addresses to access Linux BMSs over SSH.
Inbound	TCP	3389	Source: 0.0.0.0/0	Allows all IP addresses to access Windows BMSs over RDP.

For more information, see *Virtual Private Cloud User Guide*.

## Key Pair Authentication

- What is a key pair?

A key pair, or SSH key pair, is an authentication method used when you remotely log in to Linux instances. A key pair is generated using an encryption algorithm. It

contains a public key, and a private key reserved for you. The public key is used to encrypt data (for example, a password), and the private key is used to decrypt the data.

The cloud platform stores the public key, and you need to store the private key. Do not share your private key with anyone. Keep your private key secure.

- Advantages

A key pair is more secure and easier to use than username/password in authentication.

**Table 1-5** Comparison between the key pair and username/password

Item	Key pair	Username and Password
Security	<ul style="list-style-type: none"><li>• More secure than username/password and free from brute-force attacks</li><li>• Cannot be derived from the public key</li></ul>	Less secure
Easy to use	Simultaneous login to a large number of Linux instances, simplifying management	Login to only one Linux instance at one time, giving no chance of batch maintenance

- Constraints

- Only Linux instances support the key pair method.
- Only RSA key pairs are supported. RSA keys are typically 1024, 2048, or 4096 bits long.
- A Linux instance can have only one key pair. If a key pair has been bound to your BMS and you bind a new key pair to the BMS, the new key pair will replace the original one.

- Generation

- Create a key pair on the management console.

 **NOTE**

When a key pair is generated, download and properly keep it.

- Use PuTTYgen to create a key pair and import the key pair into the cloud platform.

## Helpful Links

[Using an SSH Key Pair](#)

## 1.8 Billing

## 1.8.1 Billing

### Billing Items

The billing items include BMS, EVS disk (optional), and EIP (optional). For details, see [Table 1-6](#).

**Table 1-6** BMS billing

Billing Item	Description
BMS	Pricing for the BMS is based on the specifications you choose, including CPU, memory, local disks, and extended configurations.
(Optional) EVS disk	EVS disks that you create when you create a BMS are billed in the same manner as the BMS.
(Optional) EIP	You can bind an EIP to a BMS and pay for the EIP by bandwidth or traffic.

### Billing Modes

BMSs are billed on a pay-per-use basis or by reserved instance (RI).

- **Pay-per-use:** It is a pay-after-use mode. Billing starts when a BMS is provisioned and stops when the BMS is deleted. You can stop paying for a BMS when you no longer need it. There is no upfront payment for excess capacity.
- **RI:** This billing mode provides a larger discount than pay-per-use and is recommended for long-term users. An RI is billed based on the purchased duration specified in the order.

### Billing Involved in Specifications Changes

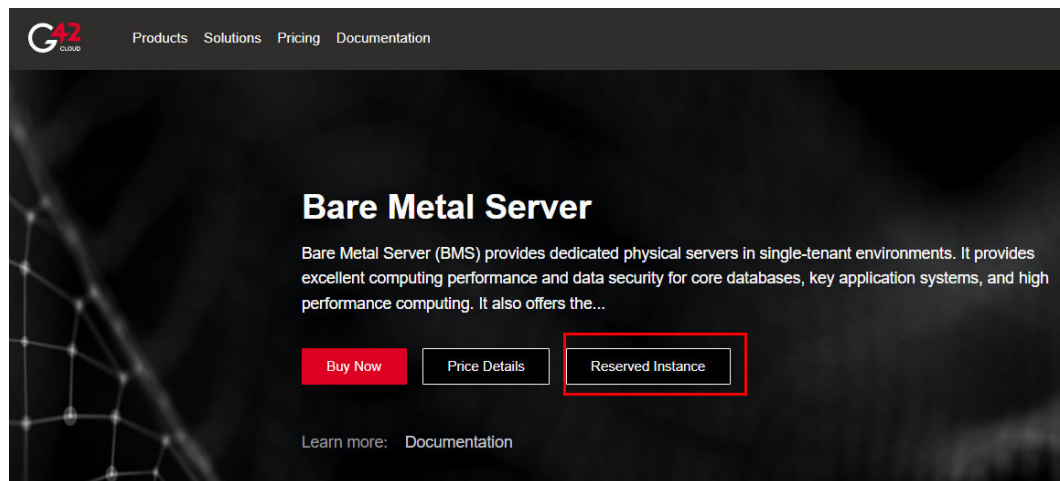
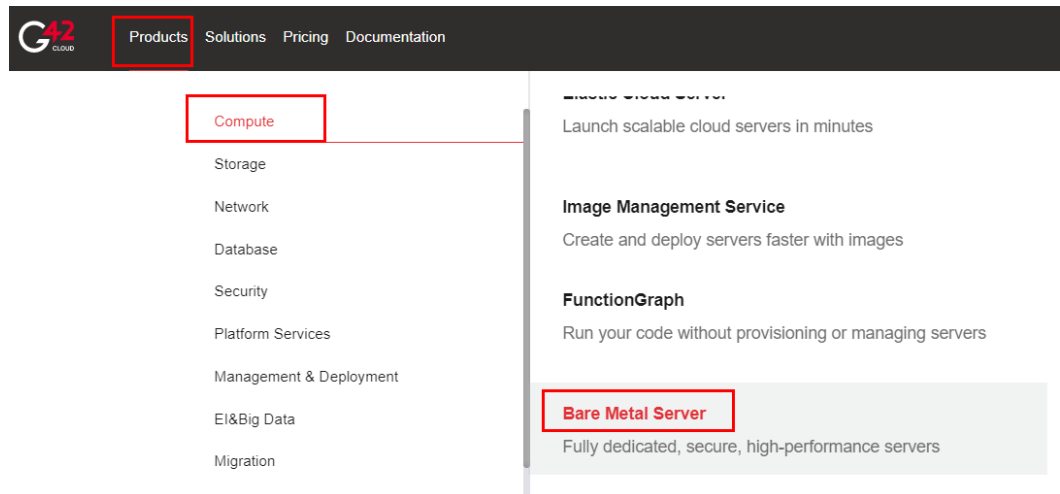
If you require additional storage space, you can either expand the capacity of EVS disks that are attached to a BMS or attach more EVS disks to the BMS. The additional storage space will be billed depending on the billing modes you specified for the EVS disks.

### What Is an RI?

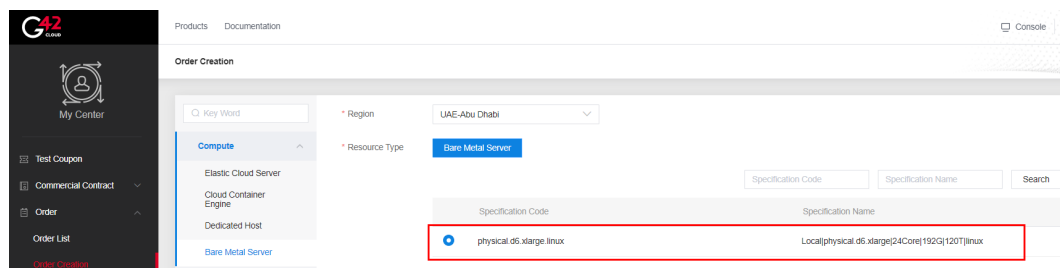
Reserved instances (RIs) help you save money by committing to one-year or three-year plans for BMSs. Committing allows you to get a discount on the BMSs you use. RIs provide a billing discount and don't affect the runtime state of your BMSs. After you purchase an RI, the discount automatically applies to matched BMSs. You can pay for the entire reserved instance term with one upfront payment or on a monthly basis, and you will get more discounts if you choose all upfront payment.

## Purchasing an RI

**Step 1** Go to the BMS portal and click **Reserved Instance**.



**Step 2** Select the specification of the pay-per-use BMSs for which you need a discount.



**Step 3** Select the duration and payment mode.

Subscription Name	Duration	Upfront (AED)	Recurring Charge (AED/Month)	Effective Hourly (AED)	Savings over On-Demand	On-Demand Hourly (AED)
<input checked="" type="radio"/> physical.d5.xlarge.linux.1 year without upfront	12 months	--	5,890.8131	8.0696	41%	13.6773
<input type="radio"/> physical.d5.xlarge.linux.1 year with all upfront	12 months	65,897.2314	--	7.5225	45%	
<input type="radio"/> physical.d5.xlarge.linux.3 year without upfront	36 months	--	4,492.9931	6.1547	55%	
<input type="radio"/> physical.d5.xlarge.linux.3 year with all upfront	36 months	143,775.7776	--	5.4709	60%	

\* Quantity: 1      BMS Price: 5,890.8131 AED / Month  
This price is an estimate and may differ from the final price.

**Step 4** Confirm and submit the purchase information.

**Subscription Features**

**Specification and Quantity**

Region	UAE-Abu Dhabi	Cloud Service	Bare Metal Server
Resource Type	Bare Metal Server	Specification Code	physical.d5.xlarge.linux
Specification Name	Localphysical.d5.xlarge(24Core)(192Gi)(120T)linux	Duration	36 months
Begin Date	Mar 01 2023 00:00:00	End Date	Feb 28 2026 23:59:59
Billing Mode	Without Upfront	Quantity	1

**Payment**

Recurring Charge	4,492.9931 AED	Discount	0% OFF	Recurring Charge After Discount	4,492.9931 AED
Total Amount	161,747.7516 AED	Total Payment Amount	161,747.7516 AED		

----End

**RI Billing**

Before ordering RIs, you can first understand the matching rules and billing rules.

**Matching dimensions:**

RIs and pay-per-use BMSs are matched based on their regions, account names, cloud services, resource types, and specifications.

**Matching rules:**

Starting from the next month of an RI order, the system automatically checks for pay-per-use BMSs that have the same specifications as the RIs in the order at 4 am on the first day of each month. If the number of RIs in the order is no less than that of the pay-per-use BMSs, all the pay-per-use BMSs with the specifications will enjoy a discount. Otherwise, the bill consumption amount of each BMS is sorted from largest to smallest, and the largest amount is matched first. The remaining unmatched BMSs are billed on a pay-per-use basis.

After the RIs are successfully matched, they will be associated with the pay-per-use BMSs and the usage of the RIs in the order will be updated. When the number of unused RIs is 0, this order will no longer participate in the match.

**Billing rules:**



At the beginning of each month, you can view RI billing records of the last month from the user center. After an RI is purchased successfully, regardless of whether it can match a pay-per-use BMS, you need to pay for it within the validity period.

If an RI matches a pay-per-use BMS, you will not need to pay the pay-per-use fee for it. Otherwise, you need to pay for the BMS on the original pay-per-use basis.

## Unsubscribing from an RI

In principle, RIs cannot be unsubscribed and will automatically expire after the validity period expires.

# 1.9 Region and AZ

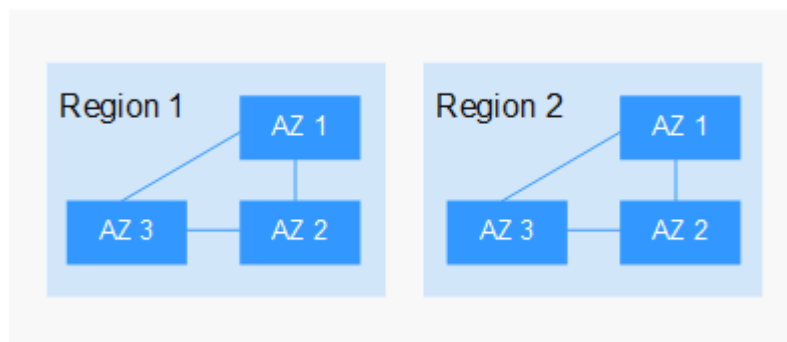
## Concept

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

- A region is a physical data center, which is completely isolated to improve fault tolerance and stability. The region that is selected during resource creation cannot be changed after the resource is created.
- An AZ is a physical location where resources use independent power supplies and networks. A region contains one or more AZs that are physically isolated but interconnected through internal networks. Because AZs are isolated from each other, any fault that occurs in one AZ will not affect others.

**Figure 1-5** shows the relationship between regions and AZs.

**Figure 1-5** Regions and AZs



## Selecting a Region

Select a region closest to your target users for lower network latency and quick access.

## Selecting an AZ

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

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

## Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see [Regions and Endpoints](#).

# 1.10 Related Services

## Relationships Between BMS and Other Services

- **Image Management Service (IMS)**  
You can quickly create BMSs using images. You can also create private images using BMSs.
- **Virtual Private Cloud (VPC)**  
You can configure a logically isolated network for your BMSs and configure security groups, VPN, IP address segments, and bandwidth. With a VPC, you can easily manage and configure internal networks and change network configurations. You can also customize access rules to control BMS access within a security group and across different security groups to enhance BMS security.
- **Elastic Volume Service (EVS)**  
You can attach EVS disks to a BMS and expand their capacity at any time.
- **Cloud Eye**  
After you obtain a BMS and install and configure Agent on the BMS, you can view the monitoring data of the BMS in Cloud Eye.

# 1.11 Features and Constraints

## Features

### BMS

- Automatic BMS provisioning and remote login to BMSs through the management console
- Managing the lifecycle of a BMS, including querying, starting, stopping, restarting, and deleting a BMS
- Rebuilding a BMS if the BMS hardware or SDI iNIC is damaged
- Injecting scripts to simplify BMS configuration and initialization
- Using APIs to manage BMSs
- Server monitoring, with which you can obtain the CPU, memory, and disk I/O metrics of your BMSs
- Tagging BMSs to make them easier to identify and search

### Disk

- Attaching EVS disks to or detaching EVS disks from Linux or Windows BMSs
- Shared EVS disks
- Dynamic capacity expansion of EVS disks

### Image

- Using a public, private, or shared image to create BMSs
- Creating a private image from a BMS
- Creating a private image from an external image file
- Sharing images and exporting images to an OBS bucket

### Network

- VPC
- Creating a security group and defining rules to protect BMS security
- Binding an EIP to a BMS to enable the BMS to access the Internet

## Constraints

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

# 2 Getting Started

---

## 2.1 Quick Start

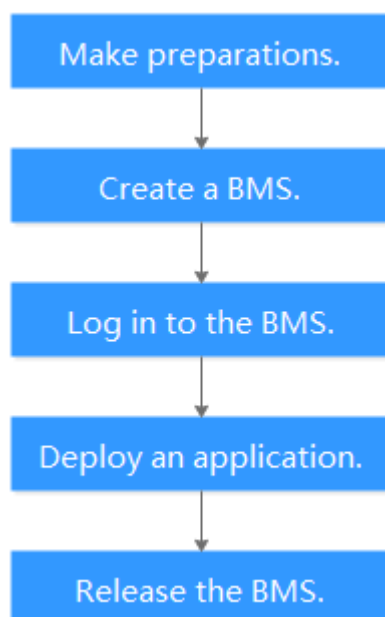
This section uses a web application server as an example to describe how you can create and use BMSs. This helps you choose an appropriate BMS, log in to it, and deploy Nginx on it.

 **NOTE**

This section is applicable only to the management console. If you use APIs, see *Bare Metal Server API*.

The following figure shows how to use BMSs.

**Figure 2-1** Process of using BMSs



## 2.2 Making Preparations

### (Optional) Create a Key Pair

The cloud platform uses public key cryptography to protect the login information of your BMS. You need to specify the key pair name and provide the private key when logging in to the BMS using SSH if you choose the key pair login mode. If you choose the password login mode, skip this section.

If you do not have a key pair, create one on the management console.

#### NOTE

If you want to create BMSs in multiple regions, you need to create a key pair in each region. For more information about regions, see [Region and AZ](#).

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 populated 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. The file name is the specified key pair name with a suffix of `.pem`. Store the private key file securely. In the displayed dialog box, click **OK**.

---

#### CAUTION

You can save the private key file only once. When you create a BMS, provide the key pair name. Each time you log in to the BMS using SSH, you need to provide the private key.

---

## 2.3 Step 1: Create a BMS

### Scenarios

This section helps you quickly create a BMS that will be used as a web server. For details about all the parameters used for creating a BMS, see [Region and AZ](#).

### Procedure

1. Log in to the Cloud Server Console.
2. In the navigation pane, choose **Bare Metal Server**.

3. In the upper right corner, click **Apply for BMS**.
4. Configure parameters.

- Specify **Region** and **AZ**.

 **NOTE**

After the BMS is created, you cannot change its region or AZ.

- Set **Flavor**.

Available flavors vary depending on the region and AZ you select. Web servers are mainly used for web page access and do not require strong computing capabilities. In addition, only a small amount of storage is required for recording logs. Therefore, select **physical.s3.large**.

- Set **Image**.

Select **Public image** and then **CentOS 7.4 64bit for BareMetal**.

- Specify **Disk**.

An EVS disk can be attached to a BMS. However, whether an EVS disk can be attached is determined by the flavor and image you select.

- Set **VPC** and **NIC**.

Retain the default values. When you use cloud services for the first time, the system automatically creates a VPC **default-vpc** and a subnet **default-subnet** for you. You can also create VPCs and subnets.

 **NOTE**

The system creates a security group for you by default. The default security group rule allows all outgoing data packets and blocks incoming data packets. In this way, the default security group rule ensures the security of basic BMS communications.

- Set **EIP**.

BMSs without an EIP cannot be connected to the Internet and are only used for deploying services in a private network or used in a cluster. Select **Not required**.

- Set **Login Mode**.

Select **Password** and set a password for user **root**.

- Configure **Advanced Settings**.

Select **Do not configure**.

- Set **BMS Name**.

The BMS name is in the format **bms-four random digits**. To easily identify it, you can add the function to its name, for example, **bms-7676-nginx**.

- Set **Quantity**.

Set the value to **1**.

5. Click **Apply Now**. Confirm the specifications and click **Submit**.

## Result

The BMS creation process requires about 5 to 30 minutes to complete. Refresh the BMS list. After the BMS status changes from **Creating** to **Running**, the BMS is created successfully.

## Follow-up Operations

A BMS that functions as a web server must allow ICMP traffic on ports 80 and 443. These rules are not configured for the default security group. You need to add the rules after you create the BMS. For details, see *Virtual Private Cloud User Guide*.


Protocol	Direction	Port Range	Source
TCP	Inbound	80	0.0.0.0/0
TCP	Inbound	443	0.0.0.0/0
ICMP	Inbound	All	0.0.0.0/0

## 2.4 Step 2: Log In to the BMS

### Scenarios

After you create a BMS, you can log in to it using multiple methods. This section describes the procedure to log in to a BMS from the management console. For more login modes, see [Linux BMS Login Methods](#) or [Windows BMS Login Methods](#).

### Procedure

1. Log in to the Cloud Server Console.
2. In the navigation pane, choose **Bare Metal Server**.
3. In the upper left corner, click  and select a region.
4. In the BMS list, locate the instance **bms-7676-nginx** and click **Remote Login** in the **Operation** column.
5. Wait for about one minute till the login page is displayed. Press **Enter** and enter username **root** and the password set in [Step 1: Create a BMS](#). Press **Enter**.

The login is successful if the following information is displayed:

```
[root@bms-7676-nginx ~]#
```

#### NOTE

If you have forgotten the login password, you can reset the password.

## 2.5 Step 3: Deploy an Application

This section describes how to deploy an application on a BMS.

### Install and Start Nginx

1. Run the **yum install nginx** command to install Nginx and enter **y** as prompted.

If the information shown in the following figure is displayed, Nginx is installed successfully.

```
Installed:
nginx.x86_64 1:1.12.2-3.el7

Dependency Installed:
dejavu-fonts-common.noarch 0:2.33-6.el7
fontconfig.x86_64 0:2.13.0-4.3.el7
gd.x86_64 0:2.0.35-26.el7
libX11.x86_64 0:1.6.5-2.el7
libXau.x86_64 0:1.0.8-2.1.el7
libjpeg-turbo.x86_64 0:1.2.90-6.el7
libxslt.x86_64 0:1.1.28-5.el7
nginxfilesystem.noarch 1:1.12.2-3.el7
nginx-mod-http-image-filter.x86_64 1:1.12.2-3.el7
nginx-mod-http-xslt-filter.x86_64 1:1.12.2-3.el7
nginx-mod-stream.x86_64 1:1.12.2-3.el7
dejavu-sans-fonts.noarch 0:2.33-6.el7
fontpackages-filesystem.noarch 0:1.44-8.el7
gperftools-libs.x86_64 0:2.6.1-1.el7
libX11-common.noarch 0:1.6.5-2.el7
libXpm.x86_64 0:3.5.12-1.el7
libxcb.x86_64 0:1.13-1.el7
nginx-all-modules.noarch 1:1.12.2-3.el7
nginx-mod-http-geoip.x86_64 1:1.12.2-3.el7
nginx-mod-http-perl.x86_64 1:1.12.2-3.el7
nginx-mod-mail.x86_64 1:1.12.2-3.el7

Complete!
```

2. Enter `systemctl start nginx.service` to start Nginx.

#### NOTE

This command applies to CentOS 7.4 64-bit, which is used as an example.

3. Enter `wget http://127.0.0.1` to test Nginx.

```
root@bms-~# wget http://127.0.0.1
--2019-07-04 11:06:32-- http://127.0.0.1/
Connecting to 127.0.0.1:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 3700 (3.6K) [text/html]
Saving to: 'index.html'

100%[=====>] 3,700 --.-K/s in 0s
2019-07-04 11:06:32 (532 MB/s) - 'index.html' saved [3700/3700]
```

## Access the Default Web Page

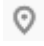
Open a browser and enter `http://BMS EIP` in the address box. If the Nginx welcome page is displayed, Nginx is installed successfully.

## 2.6 Step 4: Release the BMS

### Scenarios

If you no longer require the BMS, you can release it to avoid consuming the BMS quota.

### Procedure

1. Log in to the Cloud Server Console.
2. In the navigation pane, choose **Bare Metal Server**.
3. In the upper left corner, click  and select a region.
4. In the BMS list, locate **bms-7676-nginx**. Click **More** in the **Operation** column and select **Delete** from the drop-down list.
5. In the displayed dialog box, confirm the information and click **OK**.

If the BMS has associated resources, such as EVS disks and EIP, you can choose whether to delete these resources.



## Result

The deleted BMS will no longer be displayed in the BMS list.

# 3 Instance

---

## 3.1 Creating a BMS

### 3.1.1 Introduction

You can:

- Create a common BMS meeting your basic requirements as instructed in [Creating a Common BMS](#).
- Create a BMS that can be quickly provisioned. For details, see [Creating a BMS Supporting Quick Provisioning](#).
- Create a private image containing a required OS and applications and use it to create a BMS. For details, see [Creating a BMS from a Private Image](#).

### 3.1.2 Creating a Common BMS

#### Scenarios

This section describes how to create a BMS on the management console. When creating a BMS, you need to configure basic information such as the specifications, image, storage, network, and security groups for the BMS. You can also configure additional information for personalized deployment and management.

#### Prerequisites

- You have completed [Preparations](#).
- To inject user data, you have prepared [user data scripts](#).

#### Step 1: Configure Basic Settings

1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.  
The BMS console is displayed.
3. Click **Apply for BMS**.

4. Select a region.


BMSs in different regions cannot communicate with each other over an intranet. You are advised to select the region closest to your services to lower the network delay and improve the access speed. Note that after a BMS is created, its region cannot be changed.
5. 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.
6. Select a flavor.

**Flavor** contains 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 the BMS has two 10GE NICs, each with two ports. One NIC is used for the BMS to connect to a VPC and the other is used for the BMS to communicate with other BMSs in a high-speed network.

 **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 will be installed on the EVS disk attached to the BMS.
7. Set **Image**.
  - **Public Image**

A public image is a standard OS image provided by the system and 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). If you need other applications or software, configure them on the new BMSs.
  - **Private Image**

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

A shared image is a private image shared by another user.
8. Set **Disk**.

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.

9. Click **Next: Configure Network**.


## Step 2: Configure Network

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.

If you are not using VPC for the first time, perform the following operations to configure the network:

1. Configure the BMS network. That is, configure the primary and extension NICs.

Select a VPC and then a subnet in the VPC, and set a private IP address for the primary NIC by selecting **Automatically-assigned IP address** or **Manually-specified IP address**.

You can add extension NICs for the BMS as needed. Click  next to **Extension NIC**. Then, select a subnet and set a private IP address for the NIC.

---

 **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.

---

2. (Optional) Configure high-speed NICs.

High-speed NICs provide high-speed network ports for communication between BMSs. They provide high bandwidth.

Each high-speed NIC of a BMS must be in a different high-speed network.

 **NOTE**

If high-speed networks have been upgraded to enhanced high-speed networks with higher performance, you need to configure enhanced high-speed NICs.

3. Configure enhanced high-speed NICs.

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 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.

#### 4. Configure a security group.

Similar to firewall, a security group is a logical group used to control network access. 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.

#### 5. Set **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:

- **Not required**: The BMS cannot communicate with the Internet and can only be used to deploy services or clusters in a private network.
- **Automatically assign**: The system automatically assigns an EIP with a dedicated bandwidth to the BMS.
- **Use existing**: An existing EIP is assigned to the BMS.

 NOTE

If you select **Use existing**, you can create only one BMS at a time.

6. (Optional) Set **Bandwidth** for the EIP.  
This parameter is mandatory when **EIP** is set to **Automatically assign** in 5.
7. Click **Next: Configure Advanced Settings**.

### Step 3: Configure Advanced Settings

1. Set **BMS Name**.

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

If you create multiple BMSs at a time, suffixes will be added to the BMSs in sequence. For example, if you enter **bms**, the BMS names will be **bms-0001**, **bms-0002**, ... If you create 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.

2. Configure **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 3-1](#) 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 one.

- 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 3-1](#).

**Table 3-1** Password requirements

Parameter	Requirements	Example Value
Password	<ul style="list-style-type: none"> <li>• Consists of 8 to 26 characters.</li> <li>• Must contain at least three of the following character types:                             <ul style="list-style-type: none"> <li>- Uppercase letters</li> <li>- Lowercase letters</li> <li>- Digits</li> <li>- Special characters !@\$%^&amp;*_+=[]{}:./?</li> </ul> </li> <li>• Cannot contain the username or the username spelled backwards.</li> <li>• Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows BMSs.)</li> </ul>	Test12\$@

3. (Optional) Enable 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 6.

- 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 have selected **Enable auto backup**, the system will use the default backup policy.

**Figure 3-1** Default backup policy



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

4. (Optional) Configure **Advanced Options**.

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

- **User Data Injection** enables the BMS to automatically inject user data when the BMS starts for the first time. After this function is enabled, the BMS automatically injects user data upon its first startup.

This parameter is available only when **Key pair** is selected for **Login Mode**. For detailed operations, see [Injecting User Data](#).


– **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. If you have no agency, click **Create Agency** to create one. Currently, agencies are mainly used for server monitoring.

5. Click **Next: Confirm**.

## Step 4: Confirm

1. Check the configurations. If any configuration is incorrect or needs to be modified, click  next to **Configure Basic Settings**, **Configure Network**, or **Configure Advanced Settings** to return to the corresponding page for modification.
2. Set **Quantity**.  
A maximum of 24 BMSs can be created at a time.  
You can create only one BMS at a time if you:
  - Manually specify an IP address for a NIC or high-speed NIC.
  - Use an existing EIP.
3. Click **Apply Now** to return to the BMS list page.
4. Wait for about 30 minutes until the BMS status changes to **Running**. 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?](#)
- 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.



## 3.1.3 Creating a BMS Supporting Quick Provisioning

### Scenarios

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

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

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

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.

### Procedure

You can create a BMS supporting quick provisioning by following the instructions in [Creating a Common BMS](#).

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

- **Flavor:** Select a flavor that supports quick provisioning.
- **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.

## 3.1.4 Creating a BMS from 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 Common BMS](#).

Note for setting the parameters:

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

## 3.2 Viewing BMS Information

### 3.2.1 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**:

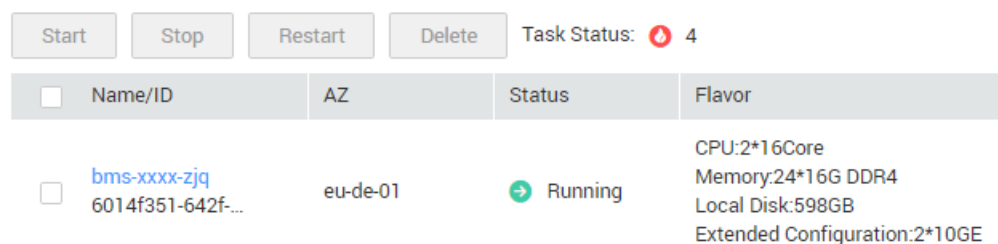
- **Processing:** The system is processing the task.
- **Failed:** The system has failed to process the task. The system rolls back the failed task and displays an error code, for example, **(BMS.3033) Failed to create system disk**.

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 create a BMS, the **Task Status** area will show the task processing status.

**Figure 3-2** BMS application status



<input type="checkbox"/>	Name/ID	AZ	Status	Flavor
<input type="checkbox"/>	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 the BMS Creation Task Displayed as Failed But the BMS List Shows the BMS?](#)

## 3.2.2 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 status, name, BMS ID, flavor, and private IP address.
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**, 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 *Cloud Eye User Guide*.

## 3.3 Logging In to a Linux BMS

### 3.3.1 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 3-2** Linux BMS login methods

Access to the Internet	On-premise OS	Login Method
Yes/No	Windows or Linux	<a href="#">Remotely Logging In to a BMS</a>
Yes	Windows	Use a remote login tool, such as PuTTY. <ul style="list-style-type: none"><li>• For how to log in to a BMS using an SSH key pair, see <a href="#">Logging In to a BMS Using an SSH Key Pair</a>.</li><li>• For how to log in to a BMS using an SSH password, see <a href="#">Logging In to a BMS Using an SSH Password</a>.</li></ul>
Yes	Linux	Run commands. <ul style="list-style-type: none"><li>• For how to log in to a BMS using an SSH key pair, see <a href="#">Logging In to a BMS Using an SSH Key Pair</a>.</li><li>• For how to log in to a BMS using an SSH password, see <a href="#">Logging In to a BMS Using an SSH Password</a>.</li></ul>

## 3.3.2 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.
- 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.
- 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 3-3** Setting a password for user **root**

```
[root@serverc28ef36e-08ef-4d94-8921-155fa4d4332b ~]# passwd
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
[root@serverc28ef36e-08ef-4d94-8921-155fa4d4332b ~]#
```

## 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?](#)

## 3.3.3 Logging In to a BMS Using an SSH Key Pair

### Scenarios

This section describes how to log in to a Linux BMS using an SSH key pair from a Windows or Linux PC.

### 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 [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 PC

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.

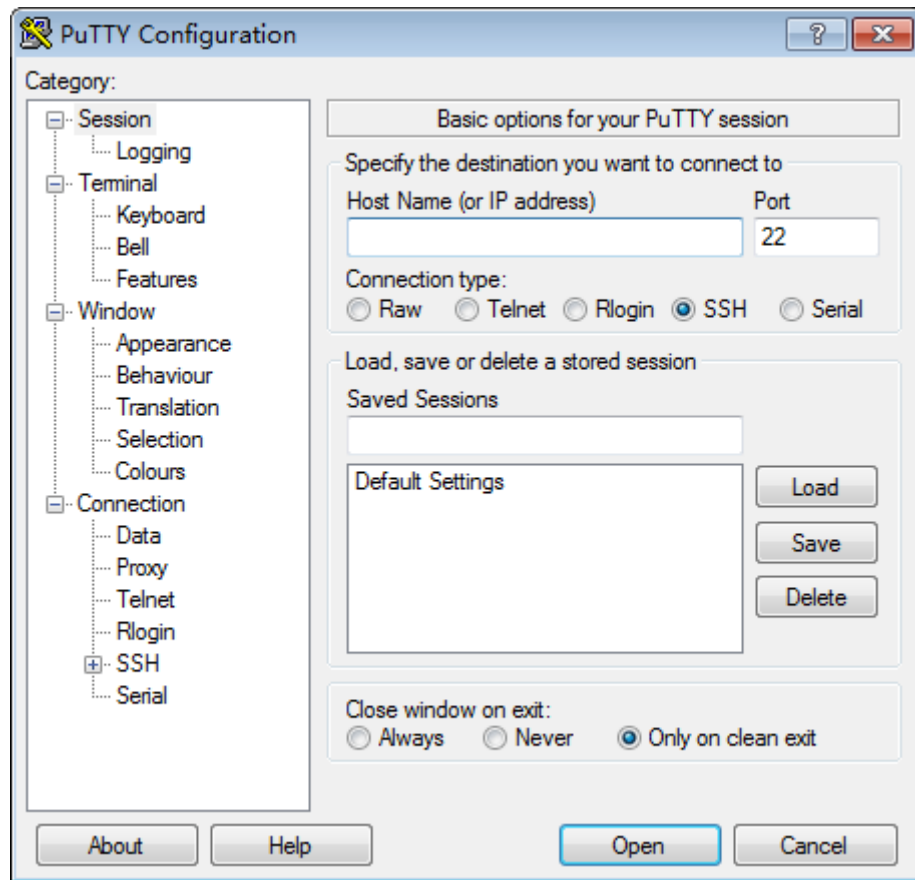
1. Check whether the private key file has been converted to **.ppk** format.
  - If yes, go to step **7**.
  - If no, go to step **2**.
2. Visit the following website and download PuTTY and PuTTYgen:  
<https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>

#### 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 3-4 PuTTY Configuration



8. Choose **Connection > Data**. Enter the image username in **Auto-login username**.

**NOTE**

Contact the operation administrator to obtain the image 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 the 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**.
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 PC

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 xxx@EIP of the BMS
```

### NOTE

- In the preceding command, *path* refers to the path under which the key file is stored.
- *xxx* indicates the username of the BMS image. Contact the operation administrator to obtain the username.

## 3.3.4 Logging In to a BMS Using an SSH Password

### Scenarios

This section describes how to log in to a Linux BMS using an SSH password from a Windows or Linux PC.

### 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 to 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
```

## 3.4 Logging In to a Windows BMS

### 3.4.1 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.

### 3.4.2 Logging In to a BMS Remotely Using MSTSC

#### Scenarios

This section describes how to log in to a Windows BMS using MSTSC (a remote login tool) from your local PC.

#### 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.

## 3.5 Managing BMSs

### 3.5.1 Changing the Name of a BMS



#### Scenarios

To make it easy for you to identify and manage each BMS, the cloud platform 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:

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

Check that automatic synchronization is enabled.

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

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

[METADATA]
enable_preserve_hostname = False
auto_synchronize_hostname = True

[IB]
enable_ib = True

[ROCE]
enable_roce = True
```

 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.

## 3.5.2 Stopping a BMS

### Scenarios

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

 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.
- You can stop a BMS only on the management console and cannot run `shutdown` to stop it. It is because that the `shutdown` and other commands attempting to stop a BMS will be regarded as unexpected operations and will not take effect.

### 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](#)

## 3.5.3 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**.

## 3.5.4 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.
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 3-5](#).
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 is configured for the BMS, 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.

## 3.5.5 Rebuilding a BMS

### Scenarios

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

#### NOTE

A BMS cannot be rebuilt automatically. You need to contact the operation administrator to rebuild it.

### 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.

### Constraints

- 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 operation 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
```

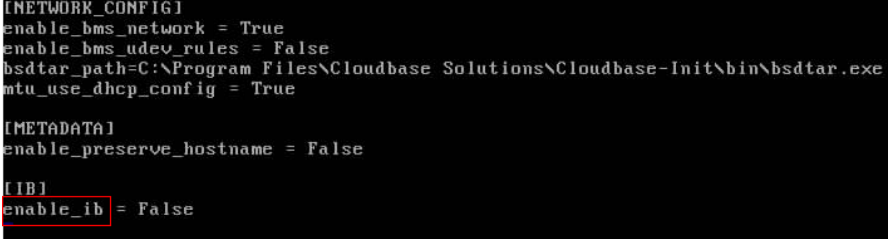
- iii. Change the value of `enable_ib` in the `bms-network-config.conf` file to **False**.

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

Check that the value has been changed.

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

**Figure 3-5** Checking the value of `enable_ib`



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

[METADATA]
enable_preserve_hostname = False

[IB]
enable_ib = False
```

- 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
```

## 3.5.6 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.

## Constraints

- 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 User Guide*.

## 3.5.7 Releasing a BMS

### Scenarios

You can delete BMSs you no longer need.

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.

### 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 **Delete** from the drop-down list. To delete multiple BMSs, select them and click **Delete** at the top of the BMS list.



4. In the displayed dialog box, click **Yes**.  
If the BMS has associated resources, such as EVS disks and EIP, you can choose whether to delete these resources.

## 3.6 User Data and Metadata

### 3.6.1 Injecting User Data

#### Application Scenarios

You can inject user data to configure BMSs.

- Use scripts to simplify BMS configuration.
- Use scripts to initialize BMS OSs.
- Upload scripts to BMSs at creation time.
- Use scripts for other purposes.


#### Constraints

- 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 BMSs must be a public image, a private image created based on a public image, or a private image with Cloud-Init installed.
  - The script format must comply with user data script specifications for Linux BMSs.
  - DHCP must be enabled for the VPC, and port 80 must be enabled for the security group in the outbound direction.
  - If password login is used, user data injection will be unavailable.
- 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 BMSs 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.

## Procedure

1. Create a user data script. The format must comply with user data script specifications. For details, see [Helpful 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** text box or upload the user data file.

**Figure 3-6** 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

User data scripts of Linux BMSs are customized by using the open-source Cloud-Init architecture. This architecture uses BMS metadata as the data source for automatically configuring the BMSs. The 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 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: user-data scripts and Cloud-Config data scripts

**Table 3-3** 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.

-	User-Data Script	Cloud-Config Data
Format	A script must be started with <b>#!</b> , for example, <b>#!/bin/bash</b> and <b>#!/usr/bin/env python</b> . When the BMS is started for the first time, the script will be executed at the rc.local-like level, indicating a low priority in the boot sequence.	The first line must be <b>#cloud-config</b> , 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 depending on the applications installed on the BMS.

- How can I view the user data injected into a Linux BMS?
  - a. Log in to the BMS.
  - b. Run the following command to view the user data as user **root**:  
**curl http://169.254.169.254/openstack/latest/user\_data**

- 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** to **As Text** and enter the user data script content.

```
#!/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** to **As Text** and enter the user data script content.

```
#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 3-7** Viewing the execution result

```
localhost localhost.localdomain localhost4 localhost4.localdomain4
:1 localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.1.130 us.archive.ubuntu.com
```

## User Data Scripts of Windows BMSs

User data scripts of Windows BMSs are customized by using the open-source Cloudbase-Init architecture. This architecture uses BMS metadata as the data source for initializing and automatically configuring the BMSs. The 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: batch-processing program scripts and PowerShell scripts

**Table 3-4** Windows BMS script types

-	Batch-Processing Program Script	PowerShell Script
Format	The script must be started with <b>rem cmd</b> , 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 <b>#ps1</b> , 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 user data injected into a Windows BMS?
  - a. Log in to the BMS.
  - b. Enter the following URL in the address box of a browser and view the injected user data:

**[http://169.254.169.254/openstack/latest/user\\_data](http://169.254.169.254/openstack/latest/user_data)**

- Examples

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

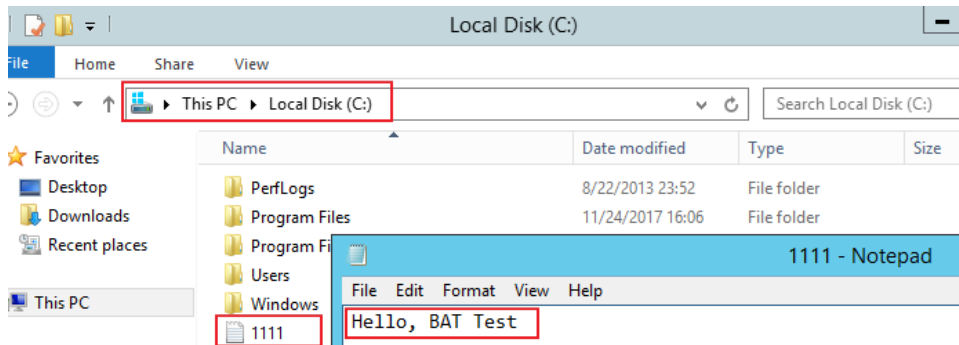
### Example 1: Inject a batch-processing program script.

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

```
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 3-8** Text file 1111.txt



To view the user data injected into the Windows BMS, log in at [http://169.254.169.254/openstack/latest/user\\_data](http://169.254.169.254/openstack/latest/user_data).

**Figure 3-9** Viewing user data in 1111.txt




**Example 2: Inject a PowerShell script.**

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

```
#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 3-10** Text file aaaa.txt



To view the user data injected into the Windows BMS, log in at [http://169.254.169.254/openstack/latest/user\\_data](http://169.254.169.254/openstack/latest/user_data).

**Figure 3-11** Viewing user data in aaaa.txt



## Case 1

This case illustrates how to inject user data so as 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 when you are creating BMSs in a batch.

The script 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 inject user data so as to reset the password for logging in to a Linux BMS.

In this example, the password of user **root** will be reset to "\*\*\*\*\*".

### NOTE

The new password must meet the password complexity requirements listed in [Table 3-5](#).

**Table 3-5** Password requirements

Parameter	Requirements	Example Value
Password	<ul style="list-style-type: none"> <li>Consists of 8 to 26 characters.</li> <li>Must contain at least three of the following character types:               <ul style="list-style-type: none"> <li>Uppercase letters</li> <li>Lowercase letters</li> <li>Digits</li> <li>Special characters <code>!@%&amp;^_+=[]{};./?</code></li> </ul> </li> <li>Cannot contain the username or the username spelled backwards.</li> <li>Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows BMSs.)</li> </ul>	Test12\$@

The script 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 new 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 inject user data so as to create a user on a Windows BMS and set a 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 3-6](#).

**Table 3-6** Password requirements

Parameter	Requirements	Example Value
Password	<ul style="list-style-type: none"> <li>Consists of 8 to 26 characters.</li> <li>Must contain at least three of the following character types:               <ul style="list-style-type: none"> <li>Uppercase letters</li> <li>Lowercase letters</li> <li>Digits</li> <li>Special characters !@\$%^_+=[]{};./?</li> </ul> </li> <li>Cannot contain the username or the username spelled backwards.</li> <li>Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows BMSs.)</li> </ul>	Test12\$@

The script is as follows:

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

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

### Case 4

This case illustrates how to inject user data so as 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 script is as follows:

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

## Case 5

This case illustrates how to inject user data so as to assign the 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 script 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
```

## Helpful 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/>

## 3.6.2 Retrieving 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 retrieved using compatible OpenStack and EC2 APIs listed in [Table 3-7](#).

**Table 3-7** 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 <a href="#">Table 3-8</a> .
	/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.



Metadata Type	Metadata Item	Description
	/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 <a href="#">Injecting User Data</a> .  For password-authenticated Linux BMSs, save the password injection script.
	/network_data.json	This interface is used to query network information of a BMS.
	/securitykey	This interface is used to obtain temporary security credentials: Access Key ID (AK) and Secret Access Key (SK).  Before obtaining temporary AK/SK on a BMS, you need to create an agency for BMS on IAM and assign required resource permissions to BMS.
EC2 type	/meta-data/hostname	This interface is used to query the host name of a BMS.  To remove the suffix <b>.novalocal</b> from a BMS, see: <a href="#">Is the BMS Host Name with Suffix novalocal Normal?</a>
	/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.

Metadata Type	Metadata Item	Description
	/meta-data/ security-groups	This interface is used to query the name of the security group of the BMS.

**Table 3-8** Metadata key fields

Parameter	Type	Description
uuid	String	Specifies the BMS ID.
availability_zone	String	Specifies the AZ where the BMS is located.
meta	Dict	Specifies the metadata information, including the image name, image ID, and VPC ID.
hostname	String	Specifies the hostname of the BMS. To remove the suffix <b>.novalocal</b> from a BMS, see: <a href="#">Is the BMS Host Name with Suffix novalocal Normal?</a>
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+dNwlydGlxJHUp+00poeUsAdBFkbPbYQTmpNwpoEb43k9z+96TyrekNKS
+iLYDdRny4kKGoNPEVBCc05Hg1TcDbIAPfJwgJS1okqEtlcofUhKmL3K0fto
+5KXEDU3GNuGwyZXjdVb9HQWU+E1jztAJjqsahnU+g/tawABTVySLBKlAT8fMGax1mTGGArucn/
WzDcy19DGioKPE7F8lLtSQ4Ww3VClK5VYB/h0x+4r7IVHrPmYX/
bi1Yhm3Dc4rRYNaTjdOV5gUOsbo3oAeQkmKwQ/
NO0N8qw5Ya4l8ZUW4tMav4mOsRySOOB35v0bvaJc6p
+50DTbWNeX5A2MLiEhTP3vsPrmvk4LRF7CLz2J2TGIM14OoVBw7LARwmv9cz532zHki/c8tlhRzLmOTXh/
wL36zFW10DeuReUGmXth7IGNmRMQKV6+mil78jm/KMPpgAdK3vwYF/
GcelOFJD2HghMUUCeMbwYnvijLTejuBpwhJMNIHA/NvlEsxDxqBCoss/Jfe+yCmUFyxovJ
+L8oNkTzkmCNzw3Ra0hiKchGhqK3BleToV/kVx5DdF081xrEA
+qyoM6CVyftEoz1zIRRYoo9bJ65Eg6Jd8dj1UCVsDqRY1pljgzE/
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**

```
ICAgICAgDQoiQSBjbG91ZCBkb2VlZlG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbiBqdXN0IHN1Y2ggYSBkaXJlY
3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkkl0IGZlZWxzIGFuIGltcHVsc2lubi4uLnRoaXMgaXMgdGhlIH
```

```
BsYWNIIHRvIGdviG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhIIHJlYXNvbnMgYW5kIHRobzZSBwYXR0ZXJucyBiZWpibmQgYWxslGNsb3VkcycwYV5kIHlvdSB3aWxslGtub3csiHRvbywgd2hIbiB5b3UgbGlmdCB5b3Vyc2VsZiBoaWdolGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiiNCg0KLVJpY2hpcmQgQmFjaA=
```

### NOTE

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

**Figure 3-12** 404 Not Found

```
[root@pythonsdktempest--server-1519783681 ~]# curl http://169.254.169.254/openstack/latest/user_data
<html>
  <head>
    <title>404 Not Found</title>
  </head>
  <body>
    <h1>404 Not Found</h1>
    The resource could not be found.<br /><br />
  </body>
</html>
```

## 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 security credentials: Access Key ID (AK) and Secret Access Key (SK).

 NOTE

- To obtain temporary AK/SK on a BMS, you need to create an agency for BMS on IAM and assign required resource permissions to BMS.
- The temporary AK/SK pair expires an hour later but is updated 10 minutes ahead of the expiration time. During the 10 minutes, both the new and old temporary AK/SK pairs can be used.
- When using temporary AK/SK, 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**

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbjBqdXN0IHN1Y2ggYSBkaXJlY3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLk0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhliH  
BsYWNIHRvIGdvlG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhliHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJl  
cyBiZWVhpbmQgYWxslGNsb3VkcycgYW5kIHlvdSB3aWxslGtub3csIHRvbywgZ2h1biB5b3UgbGlmdCB5b3  
Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2h1cmQgQmFjaA=  
=
```

## 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 AAAAB3NzaC1yc2EAAAADAQABAAQADl5Fw5k8Fgzajn1zJwLoV3+wMP+6CyvsSilc/  
hioggSnYu/AD0Yqm8vVO0kWlun1rFbdO+QUZKyVr/OPUjQSw4SRh4qsTKf/+eFoWTjplFvd1WCBZzS/  
WRenxlwR00KkcZHSJro763+wYcwKieb4eKRxaQoQvoFgVjLBULXajH4eKoKTVNtMXAvPP9aMy2SLgsJNt  
Mb9ArfziAiblQynq7UifLnN3VclzPeiWrqtzjyOp6CPUXnL0lVPTvbLe8sUteBsJZwLL6K4i  
+Y0lf3ryqnmQgC21yW4Dzu+kwk8FVT2MgWkCwiZd8gQ/+uJzrJFyMfUOBiklOBfuUENIJUuHAB  
Generated-by-Nova
```

# 4 Image

---

## 4.1 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 User Guide.

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. The following figure shows how to use private images.



Figure 4-1 Using private images



## 4.2 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" in *Bare Metal Server Private Image Creation Guide*.

- 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 **Create Image**.  
The page for creating an image is displayed.
5. Enter the image name, set a tag, and enter description as needed.  
Click **Apply Now**.
6. On the displayed **Details** page, confirm the configuration and click **Submit Application**.
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.

## 4.3 Creating a Private Image from an External Image File

### Scenarios

You can create and register a private image using an external image file. [Figure 4-2](#) shows the procedure.

**Figure 4-2** 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 *Object Storage Service User Guide*.

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

## 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.
    - **Description:** (Optional) Enter description of the image.
5. Click **Apply Now**.  
On the displayed **Details** page, confirm the configuration and click **Submit Application**.
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](#).

# 5 Disk

---

## 5.1 Disk Types

The cloud platform provides various storage products for your BMSs, including block storage based on the distributed storage architecture, dedicated storage based on 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 5-1** Comparison of storage products

Storage Product	Storage Type	Typical Application Scenarios	Process
Block storage	Shared storage pools	<ul style="list-style-type: none"> <li>• Enterprise daily work</li> <li>• Development and testing</li> <li>• Enterprise applications, including SAP, Microsoft Exchange, and Microsoft SharePoint</li> <li>• Distributed file systems</li> <li>• Various databases, including MongoDB, Oracle, SQL Server, MySQL, and PostgreSQL</li> </ul>	Create a disk and then attach the disk to the BMS.
DSS	Physically isolated storage pools and dedicated resources	<ul style="list-style-type: none"> <li>• Hybrid load: DSS supports hybrid deployment of HPC, database, email, OA, and web applications.</li> <li>• High-performance computing</li> <li>• OLAP applications</li> </ul>	<p>DSS can be used with BMSs in DeCs or those not in DeCs.</p> <ul style="list-style-type: none"> <li>• DeC scenario: Enable DeC, apply for a storage pool, create a disk in the storage pool, and attach the disk to the BMS.</li> <li>• Non-DeC scenario: Apply for a storage pool, create a disk in the storage pool, and attach the disk to the BMS.</li> </ul>
Local disks	Local disks of servers	<ul style="list-style-type: none"> <li>• Big data</li> <li>• Distributed cache</li> </ul>	Create a BMS and use its local disks directly.

## 5.2 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 "Creating an EVS Disk" in *Elastic Volume Service User Guide*.

### 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?](#)

## 5.3 Initializing Data Disks

### 5.3.1 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 5-2](#) lists the common disk partition styles. For Linux OSs, different disk partition styles require different partitioning tools.



**Table 5-2** Disk partition styles

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Linux Partitioning Tool
Master Boot Record (MBR)	2 TB	<ul style="list-style-type: none"> <li>4 primary partitions</li> <li>3 primary partitions and 1 extended partition</li> </ul> <p>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:</p> <ul style="list-style-type: none"> <li>3 primary partitions and 1 extended partition, with the extended partition containing 3 logical partitions</li> <li>1 primary partition and 1 extended partition, with the extended partition containing 5 logical partitions</li> </ul>	<ul style="list-style-type: none"> <li>fdisk</li> <li>parted</li> </ul>
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 \(parted\)](#)

## 5.3.2 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.

### 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 5-1** Server Manager



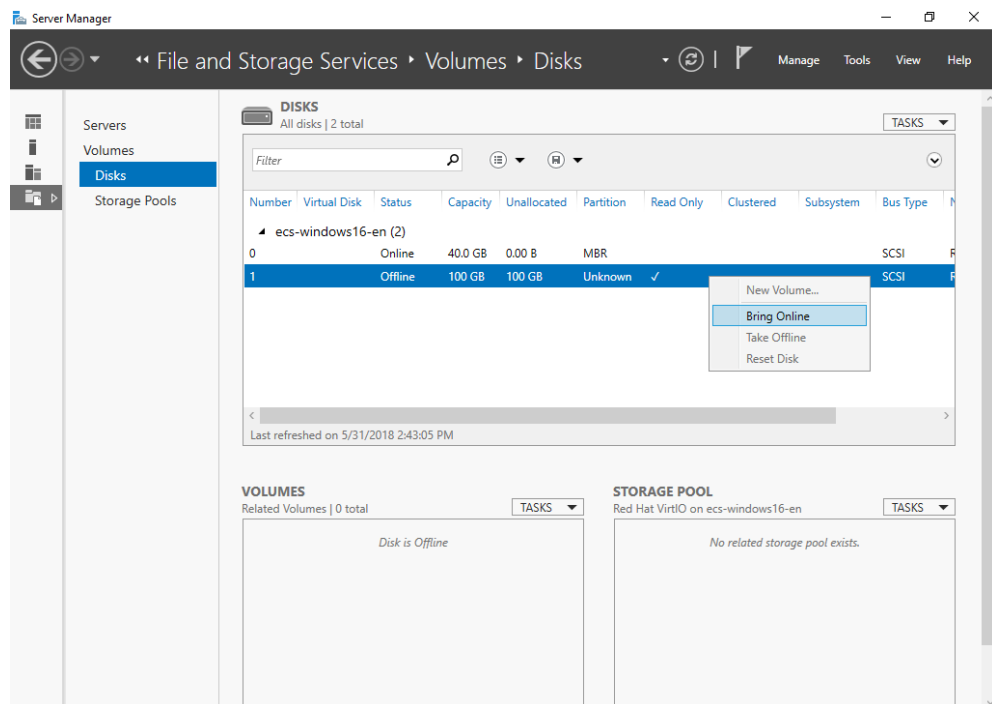
**Step 3** In the navigation tree on the left, choose **File and Storage Services**.  
The **Servers** page is displayed.

**Figure 5-2** Servers



**Step 4** In the navigation pane, choose **Disks**.  
The **Disks** page is displayed.


Figure 5-3 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 5-4 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 5-5 Bring online succeeded




**Step 6** After the disk has been brought online, initialize the disk.

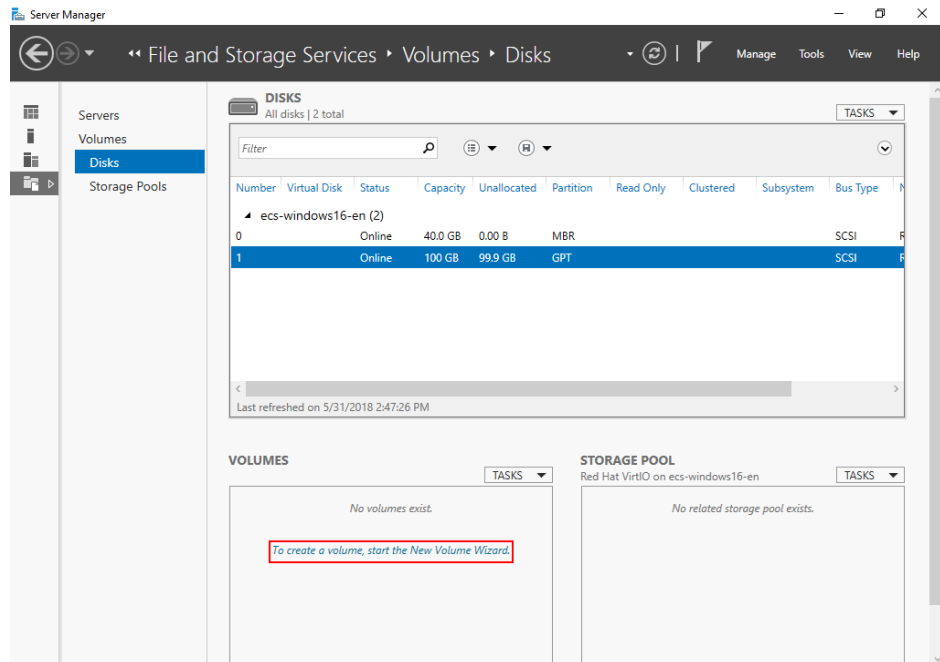
1. Right-click the new disk and choose **Initialize** from the shortcut menu. The **Initialize Disk** dialog box is displayed.

Figure 5-6 Initialize Disk (Windows 2016)



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

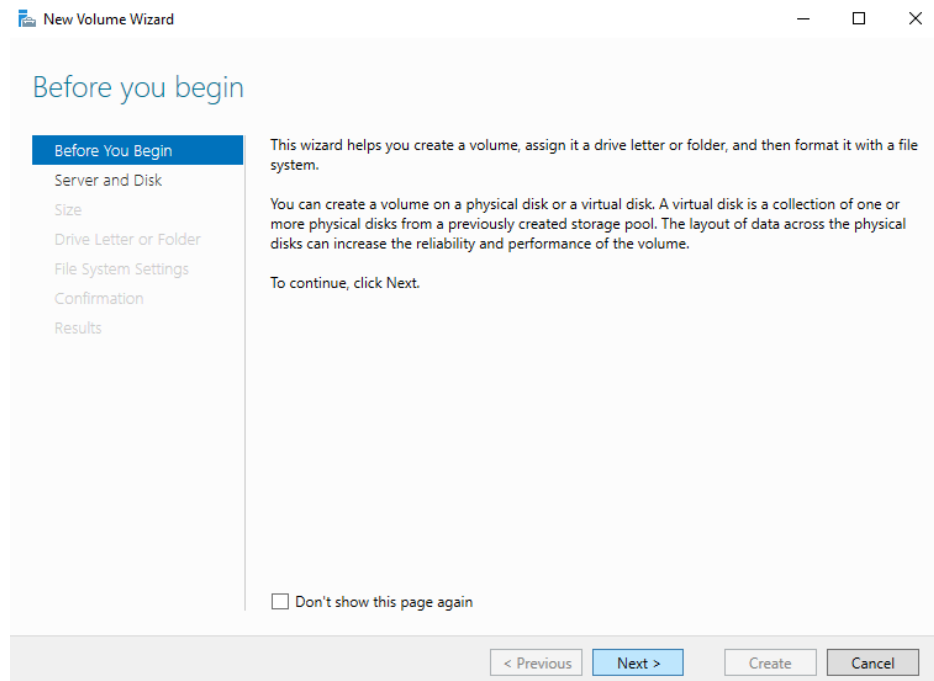
**Figure 5-7** Completing initialization



**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 5-8** New Volume Wizard



**Step 8** Follow the prompts and click **Next**.  
The **Select the server and disk** page is displayed.

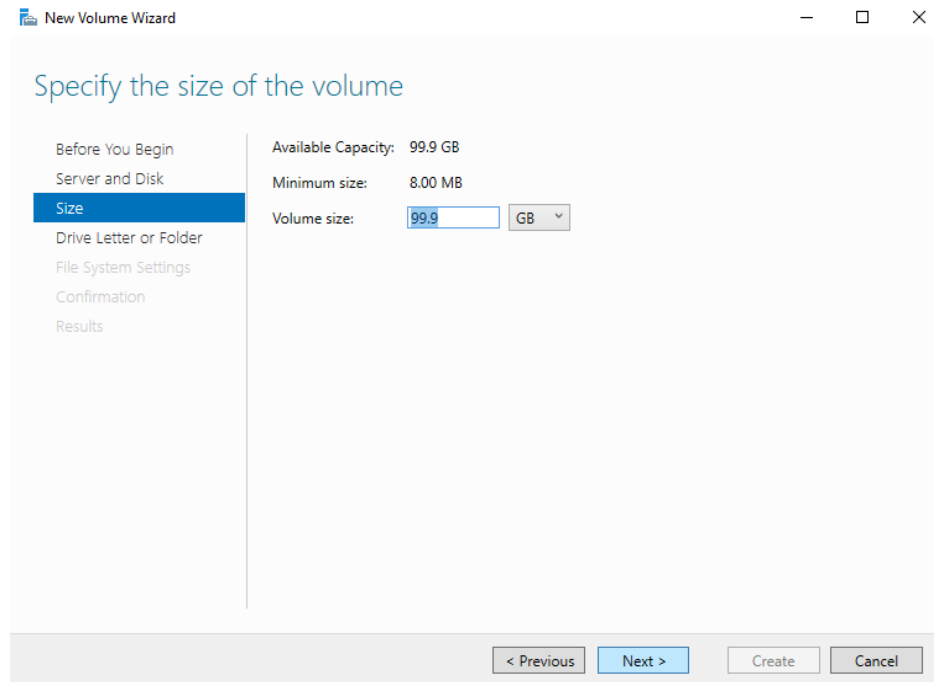
**Figure 5-9** Select the server and disk



**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.

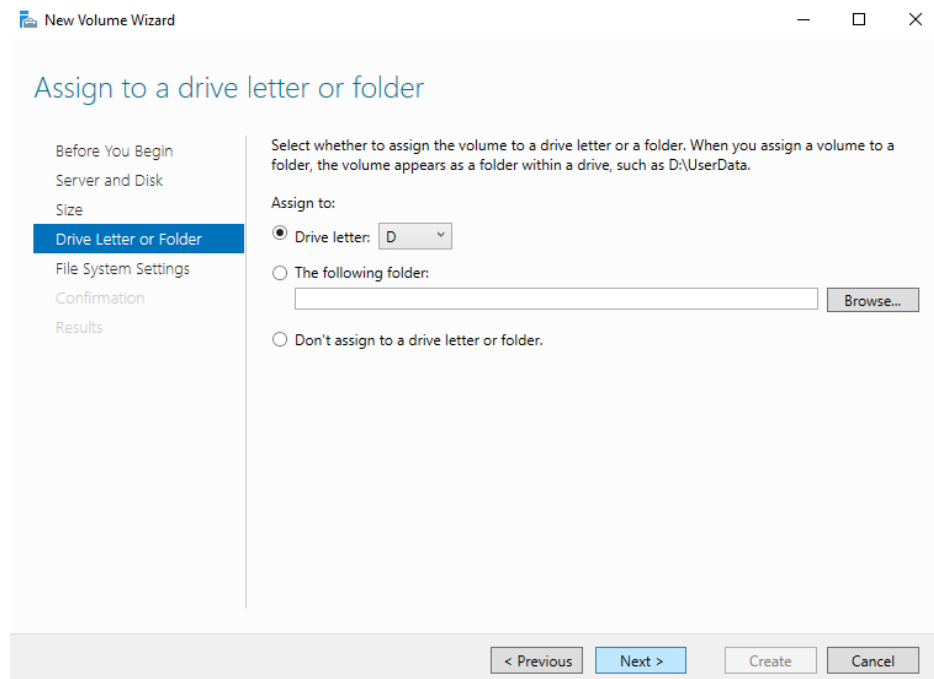
**Figure 5-10** Specify Volume Size (Windows 2016)



**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 5-11** 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 5-12** 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 5-13** 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 5-14](#) is displayed, the volume is successfully created.


**Figure 5-14** 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 5-15** 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.

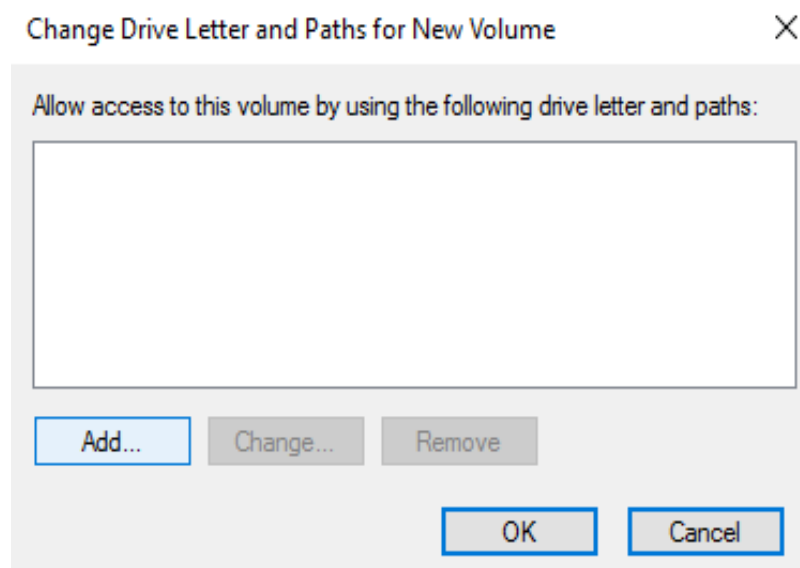
**Figure 5-16** Disk Management (Windows 2016)



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

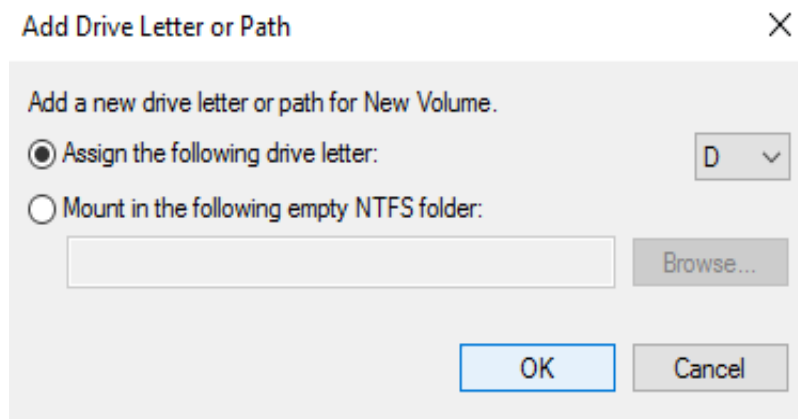
The **Change Drive Letter and Paths for New Volume** dialog box is displayed.

**Figure 5-17** Change Drive Letter and Paths for New Volume



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

**Figure 5-18** Add Drive Letter or Path



- 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

## 5.3.3 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.

### 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 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

## 5.3.4 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.

### 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 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

## 5.3.5 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.

## Prerequisites

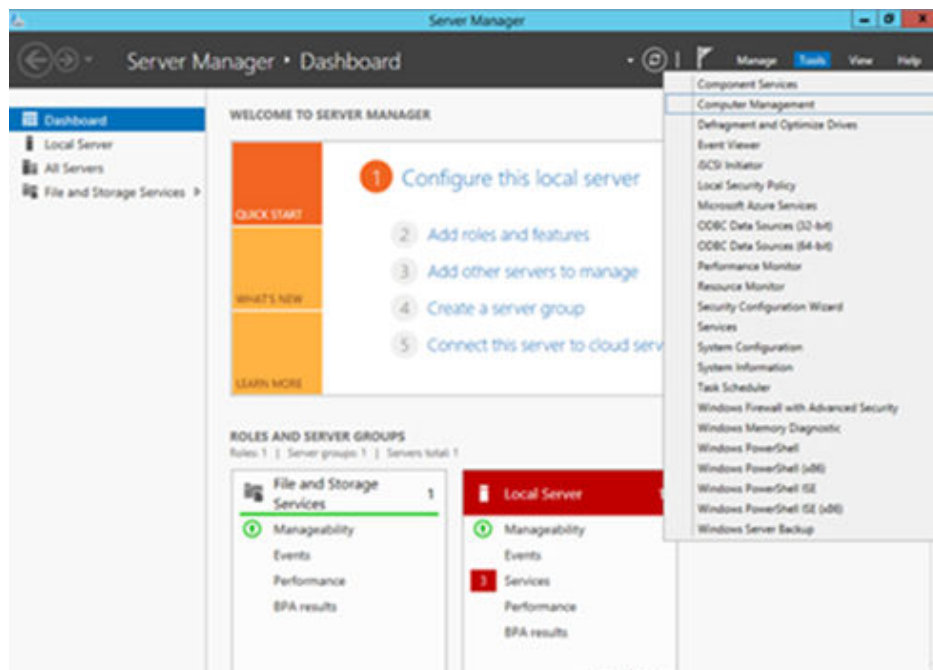
- You have logged in to the BMS.
- A data disk has been attached to the BMS and has not been initialized.

## Procedure

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

The **Server Manager** window is displayed.

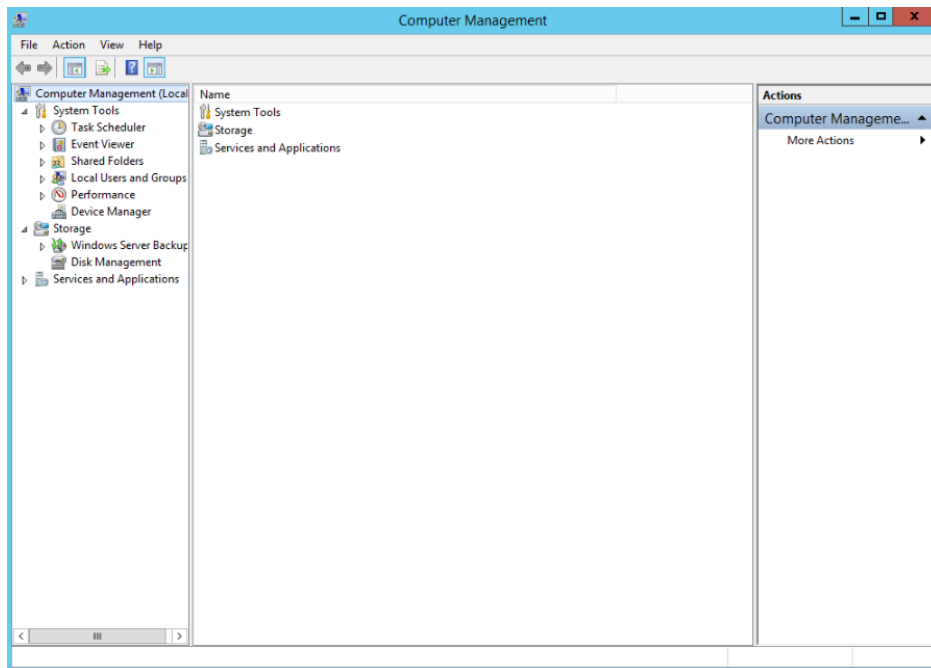
**Figure 5-19** Server Manager (Windows 2012)



**Step 2** In the upper right corner of the **Server Manager** page, choose **Tools > Computer Management**.

The **Computer Management** page is displayed.

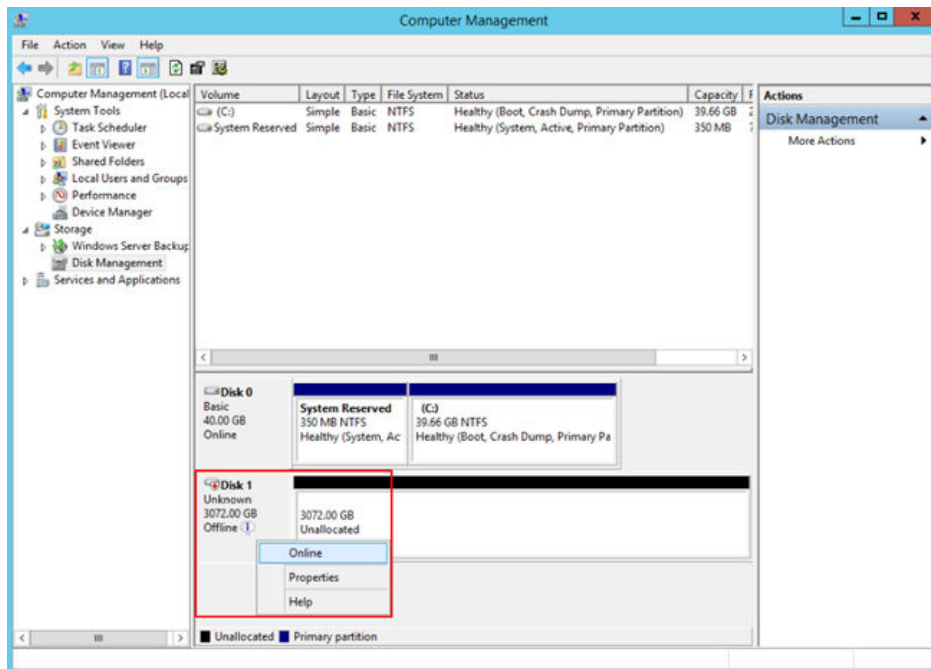
Figure 5-20 Computer Management



**Step 3** Choose **Storage > Disk Management**.

The disk list is displayed.

Figure 5-21 Disk list



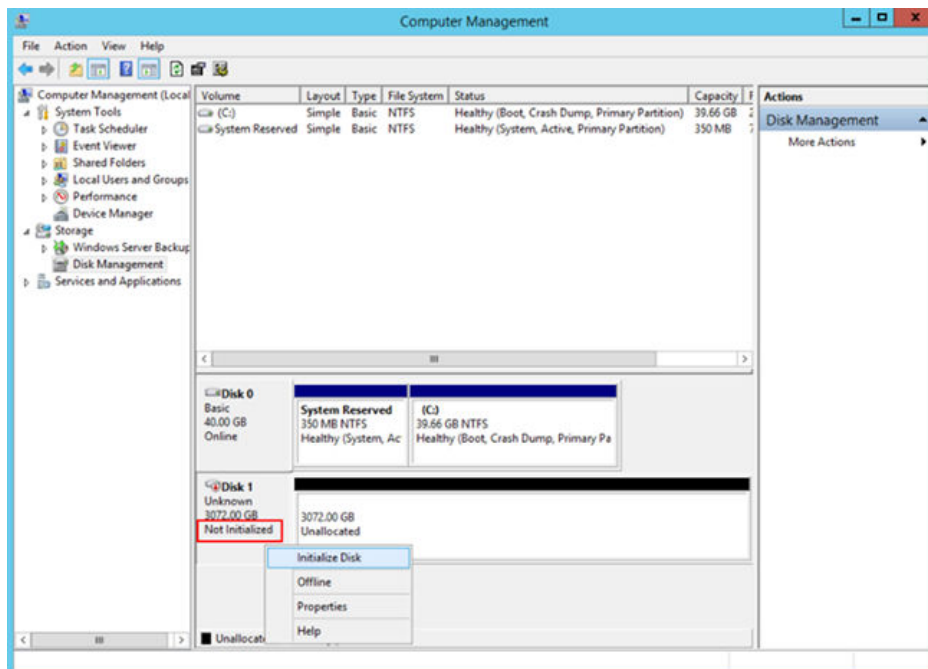
**Step 4** Disks are listed in the right pane. If the new disk is in the offline state, bring it online before initialize it.

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.

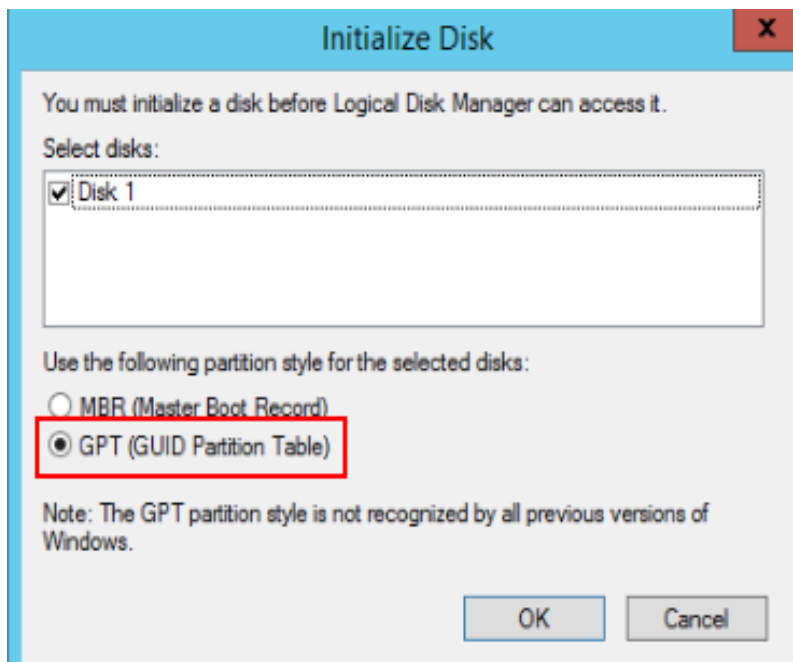


Figure 5-22 Bring online succeeded (Windows 2012)




**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 5-23 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 5-24 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.

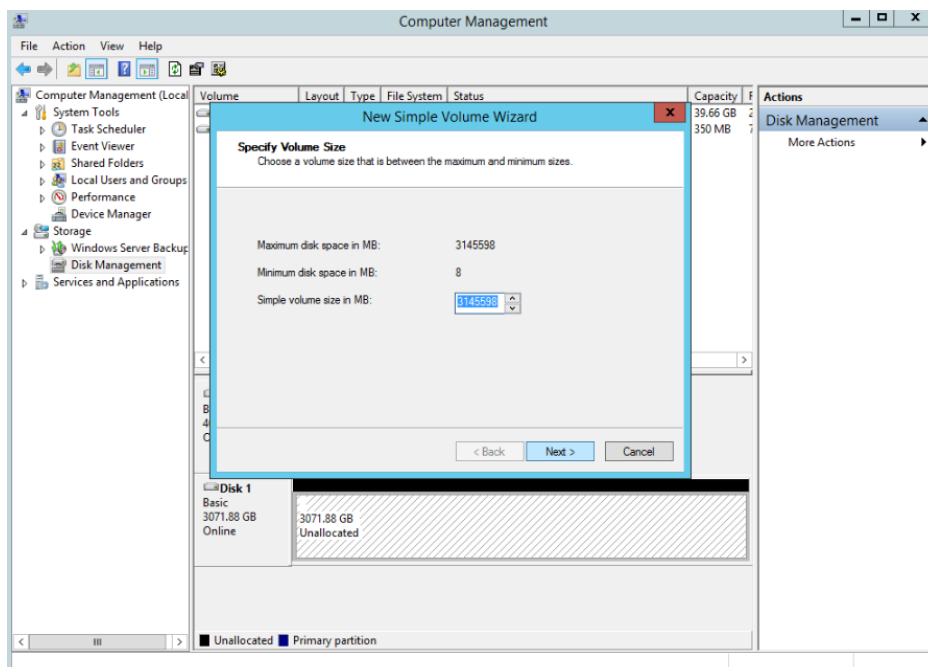
**Figure 5-25** New Simple Volume Wizard (Windows 2012)



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

The **Specify Volume Size** page is displayed.


**Figure 5-26** 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.

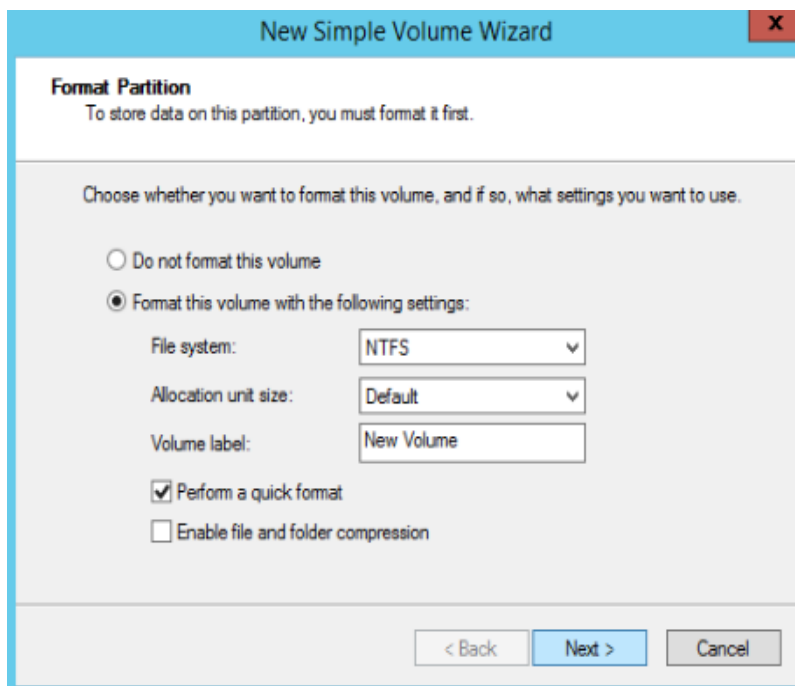
**Figure 5-27** Assign Driver Letter or Path (Windows 2012)



**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 5-28** 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 5-29** Completing the New Simple Volume Wizard (Windows 2012)




**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** Click **Finish**.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in [Figure 5-30](#).


**Figure 5-30** 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 5-31** This PC (Windows 12)



----End

## 5.3.6 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.

### 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 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

## 5.4 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**.

## 5.5 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*.

---

**NOTICE**

The system disk capacity of a Windows BMS that is quickly provisioned cannot be expanded. If you need to expand the capacity, contact technical support.

---

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 SCSI Disks (Linux)" in *Elastic Volume Service User Guide*.

# 6 Key Pair and Password

---

## 6.1 Using an SSH Key Pair

### Scenarios

To ensure system security, you are advised 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 populated 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. The file name is the specified key pair name with a suffix of .pem. Store the private key file securely. In the displayed dialog box, click **OK**.

---

**CAUTION**

You can save the private key file only once. When you create a BMS, provide the key pair name. Each time you log in to the BMS using SSH, you need to provide the private key.

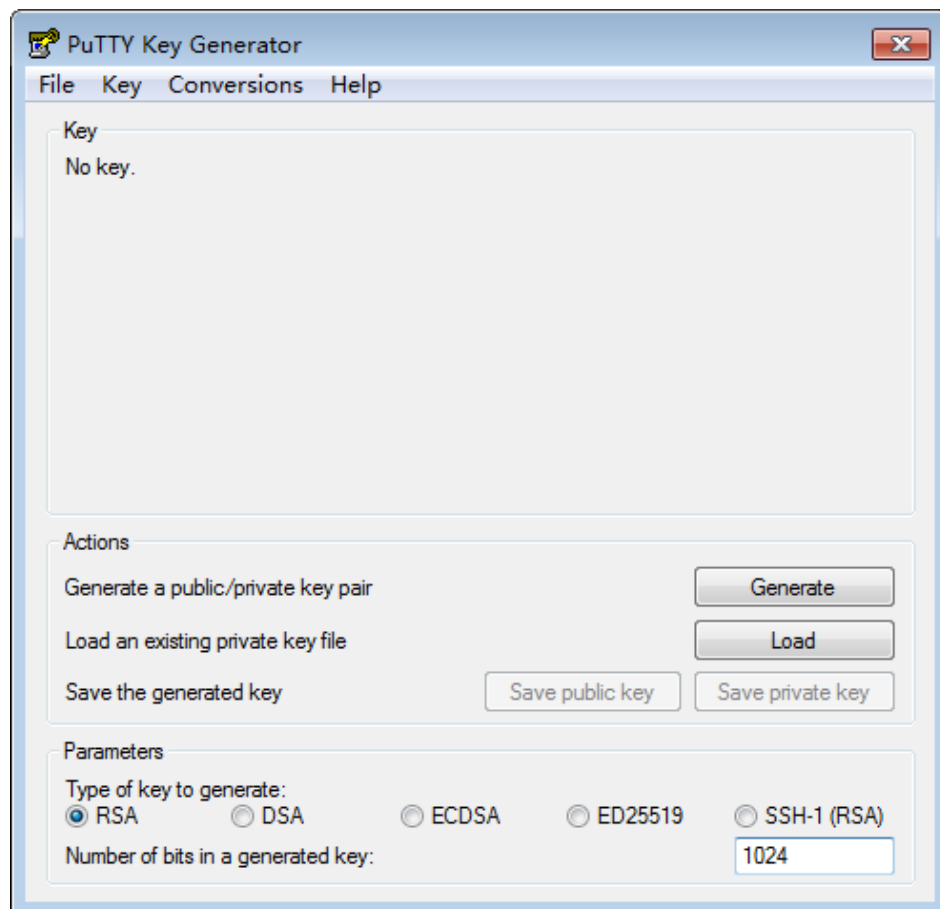
---

## 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 6-1** 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 6-2](#).

**Figure 6-2** 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**.
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.

## 6.2 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**.
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.

## 6.3 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.
5. Click **OK** to delete the password.

# 7 Network

---


## 7.1 EIP

### 7.1.1 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 7-1 Accessing the Internet through an EIP



## Helpful Links

- [Can I Bind Multiple EIPs 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?](#)

## 7.1.2 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.

### 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.

## 7.1.3 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**.

## 7.2 VPC

### 7.2.1 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.

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

[root@bms-ef79 network-scripts]# cat ifcfg-bond0.3030
PERSISTENT_DHCLIENT=1
VLAN=yesbond0
BOOTPROTO=dhcpno
TYPE=Ethernet3030
```

## 7.2.2 Binding a Virtual IP Address to a BMS


### Scenarios

You can bind a virtual IP address to a BMS for connection redundancy. This section describes how to bind a virtual IP address to a BMS.

### What Is a Virtual IP Address?

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 7-2** 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 User Guide*.

## 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 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.

## 7.2.3 Setting the Source/Destination Check for a NIC

### Scenarios

After source/destination check 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.

## 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. Expand the details of the target NIC.
5. Enable or disable **Source/Destination Check**.  
By default, **Source/Destination Check** is enabled. If the BMS functions as a NAT server, router, or firewall, you must disable the source/destination check for the BMS.

## 7.3 High-Speed Network

### 7.3.1 Overview

#### High-Speed Network

A high-speed network is an internal network among BMSs and shares the same physical plane with the 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

If high-speed networks have been upgraded to enhanced high-speed networks with higher performance, the **High-Speed Networks** tab is no longer displayed on the BMS console. For details about enhanced high-speed networks, 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.

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]#
```

## 7.3.2 Managing High-Speed Networks

### Scenarios

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.

### Constraints

- 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

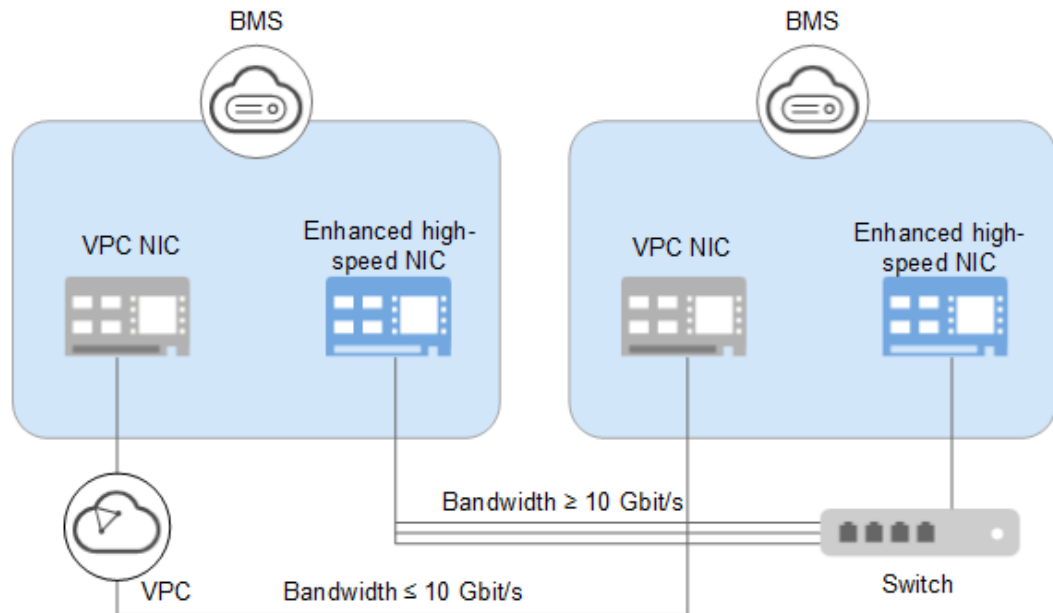
1. Log in to the management console.
2. Under **Computing**, click **Bare Metal Server**.  
The BMS console is displayed.
3. Click the **High-Speed Networks** tab. Locate the target high-speed network, click **More** in the **Operation** column, and select **Manage Private IP Address** from the drop-down 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.

## 7.4 Enhanced High-Speed Network

## 7.4.1 Overview

### Enhanced High-Speed Network

Figure 7-3 Enhanced high-speed network architecture

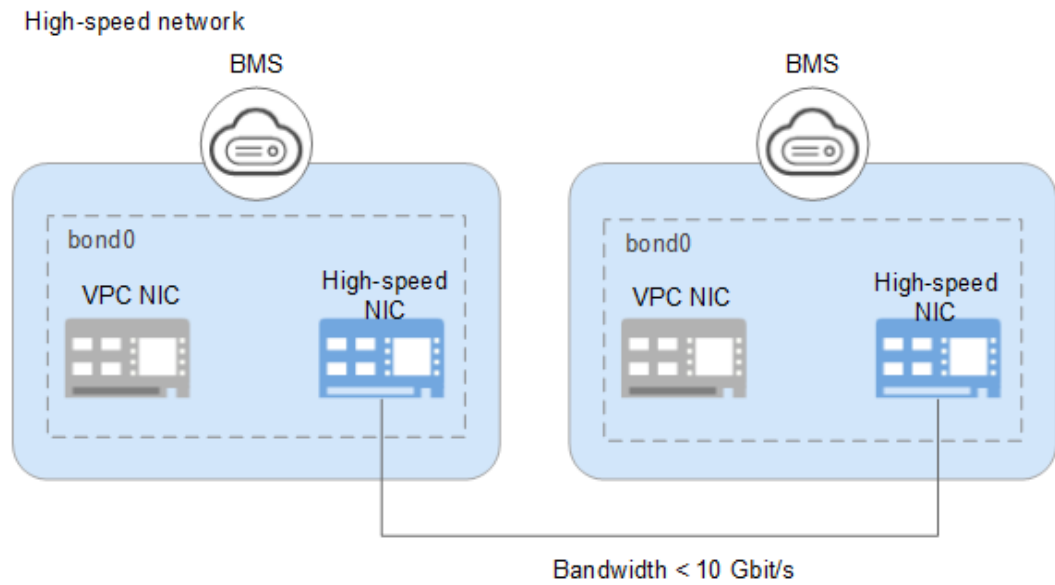


An enhanced high-speed network is a high-quality, high-speed, and low-latency internal network for BMSs to communicate with each other. It has the following features:

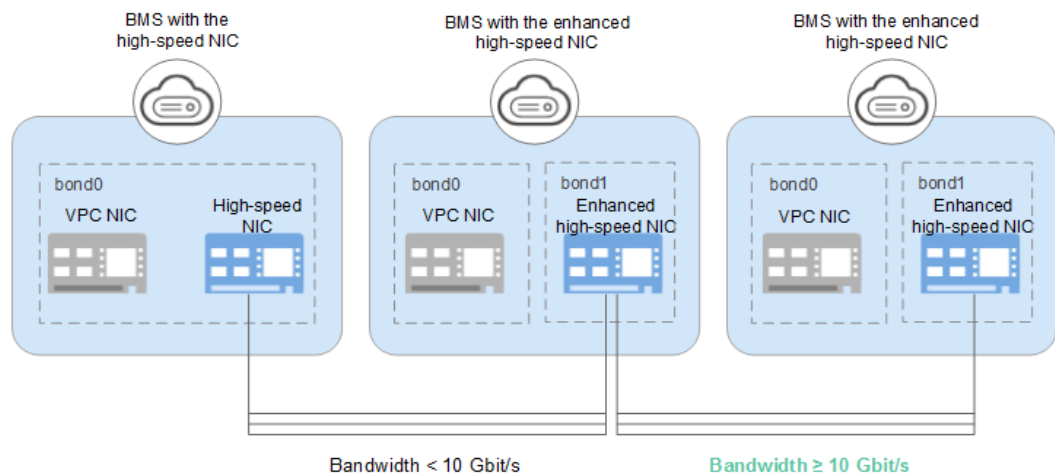
- Networks for high-speed internal interconnection
- Internal networks that you can customize
- A total bandwidth greater than 10 Gbit/s

Hardware and software in high-speed networks are upgraded to provide enhanced high-speed networks. [Figure 7-4](#) shows the architecture of the high-speed network and [Figure 7-5](#) shows a comparison between the architectures of the high-speed network and enhanced high-speed network.

**Figure 7-4** High-speed network architecture



**Figure 7-5** 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 10 Gbit/s or higher.
- The number of network planes can be customized and a maximum of 4000 subnets are supported.

## 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 7-6** 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.

## 7.4.2 Adding an Enhanced High-Speed NIC

This section describes how to add an enhanced high-speed NIC to a BMS.

### Constraints

The BMS must be in **Running** state.

## Procedure

### 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 \(Windows Server\)](#).

## 7.4.3 Deleting an Enhanced High-Speed NIC

### Scenarios

You can delete an enhanced high-speed NIC if you do not need it any longer.

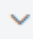
### Constraints

The BMS must be in **Running** or **Stopped** state.

### 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 **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 \(Windows Server\)](#).

## 7.4.4 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
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
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
```

## 7.4.5 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
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
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
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
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
```

**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
```

## 7.4.6 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
```

## 7.4.7 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: p3p1: <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
5: p3p2: <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: bond1.1168bond1: <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
```

## 7.4.8 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-
```

### 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 SwitchIndependent -LoadBalancingAlgorithm Dynamic -Confirm:$false  
Name : qinq  
Members : {eth3, eth2}  
TeamingMode : 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-B1-02-B4-87  
LOM4 Intel(R) Ethernet Connection X722 fo... 17 Disconnected 2C-97-B1-02-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-B1-02-B4-86  
LOM3 Intel(R) Ethernet Connection X722 ...#2 18 Disconnected 2C-97-B1-02-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
LoadBalancingAlgori : 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

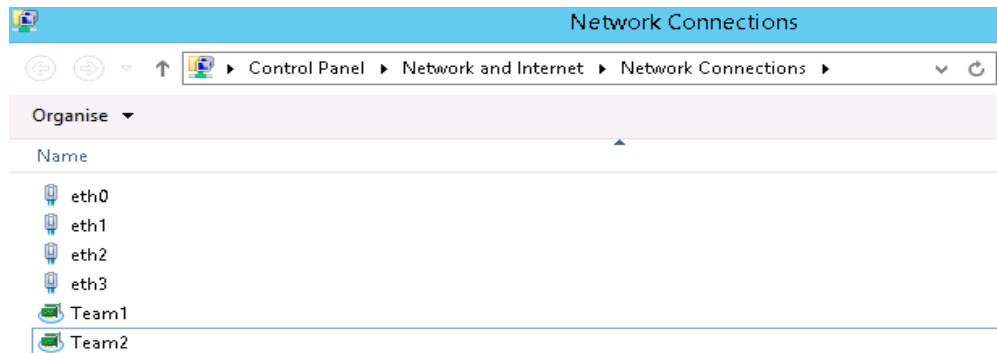


```
PS C:\Users\Administrator> Get-NetAdapter
Name InterfaceDescription IfIndex Status MacAddress LinkSpeed
-----
Team2 Microsoft Network Adapter Multiple...#2 30 Up 40-7D-0F-52-E3-AD 10 Gbps
eth2 Intel(R) 82599 10 Gigabit ????? 15 Up 2C-55-D3-C4-9C-5A 10 Gbps
eth0_198befdc-4480-49...6 Intel(R) 82599 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:

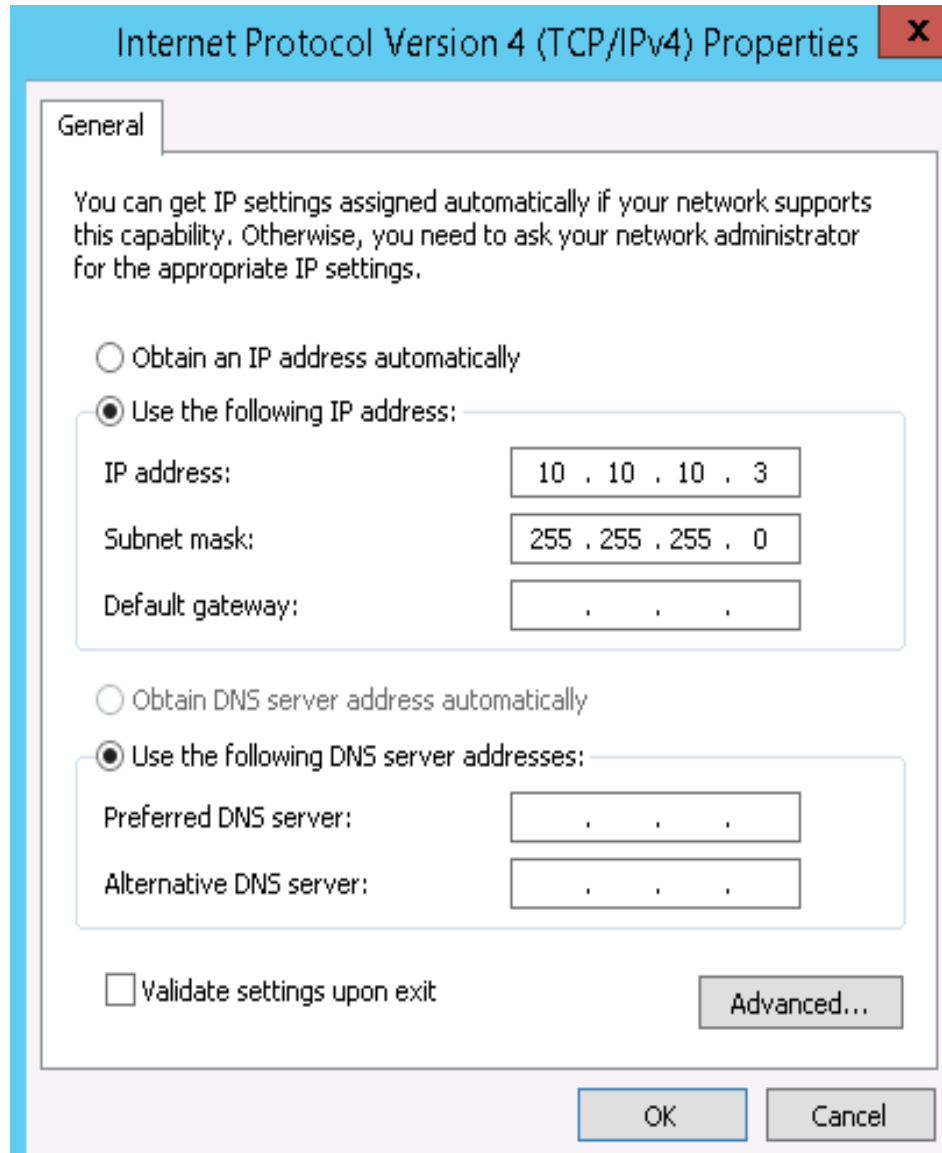
**ncpa.cpl**

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
```

## 7.5 User-defined VLAN

### 7.5.1 Overview

#### User-defined VLAN

You can use the 10GE Ethernet NICs that are not being 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. User-defined VLAN NICs are in pairs. You can configure NIC bonding to achieve high availability. User-defined VLANs in different AZs cannot communicate with each other.

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. For example, if the extended configuration of a flavor is 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.

## 7.5.2 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

## 7.5.3 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
```

```
bms-multinics-test-0002:/etc/sysconfig/network # ifup eth4
eth4 device: Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
bms-multinics-test-0002:/etc/sysconfig/network # ifup eth5
eth5 device: Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
bms-multinics-test-0002:/etc/sysconfig/network # ifup bond1
bond1
bond1 enslaved interface: eth5
bond1 enslaved interface: eth4
bms-multinics-test-0002:/etc/sysconfig/network # █
```

#### 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
```

```
bms-multinics-test-0002:/etc/sysconfig/network # ip link
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,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond1 state UP qlen 1000
    link/ether 40:7d:0f:f4:ff:5c brd ff:ff:ff:ff:ff:ff
5: eth5: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 8888 qdisc mq master bond1 state UP qlen 1000
    link/ether 40:7d:0f:f4:ff:5c 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
7: bond1: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 8888 qdisc noqueue state UP
    link/ether 40:7d:0f:f4:ff:5c brd ff:ff:ff:ff:ff:ff
```

```
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

## 7.5.4 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

## 7.5.5 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:fe4: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

## 7.5.6 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}
TeamNames           : 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
-----
gig                                  Microsoft Network Adapter Multiplexe...#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

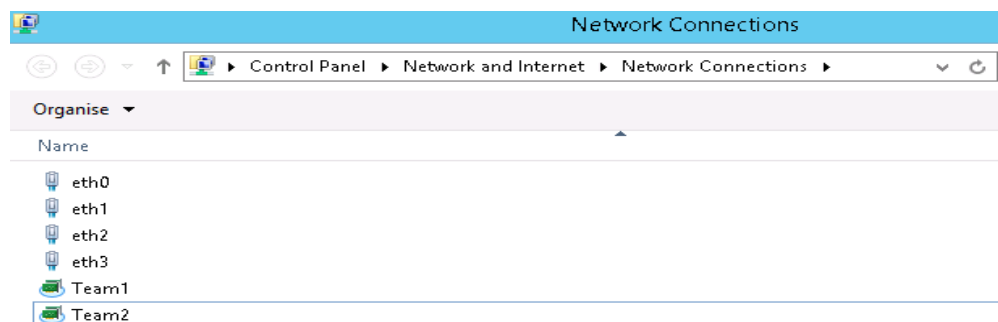
```
PS C:\Users\Administrator> Get-NetAdapter

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-C4-7B 10 Gbps
```

**Step 4** Run the following command to enter the **Network Connections** page:

**ncpa.cpl**

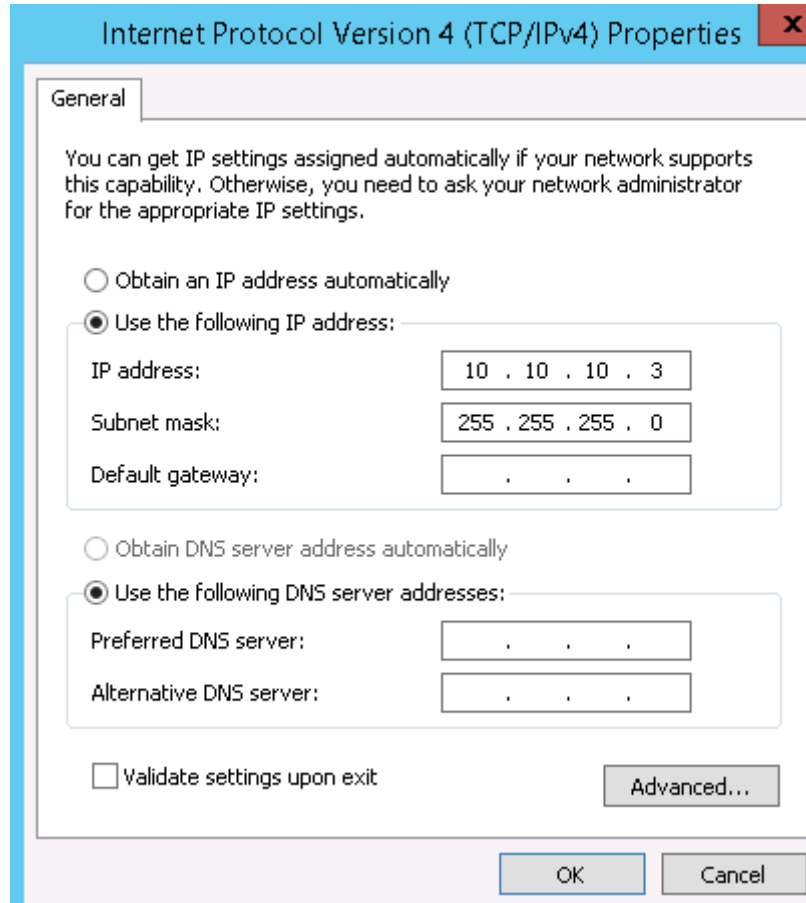
Then enter the following page.



**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.



```
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>
```

**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.

```
PS C:\Users\Administrator> Add-NetLbfoTeamNIC -Team "Team2" -VlanID 500 -Confirm:$false

Name                : Team2 - VLAN 500
InterfaceDescription : Microsoft Network Adapter Multiplexor Driver #3
Team                : Team2
VlanID              : 500
Primary             : False
Default             : False
TransmitLinkSpeed(Gbps) : 20
ReceiveLinkSpeed(Gbps) : 20
```

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

## 7.6 IB Network

### 7.6.1 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. For example, if the extended configuration of a flavor is 1\*100G IB + 2\*10GE, the BMS has IB NICs. You need to configure and plan the VLANs and IP addresses.

# 8 Security

---

## 8.1 Security Group

### 8.1.1 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 8-1** Parameter description

Parameter	Description
Protocol	Network protocol for which the security group rule takes effect. The value can be <b>All</b> , <b>TCP</b> , <b>UDP</b> , <b>ICMP</b> , or <b>GRE</b> .
Port	Port or port range for which the security group rule takes effect. The value ranges from <b>1</b> to <b>65535</b> .
Source	Traffic source (inbound rule). This parameter is required for an inbound rule. The value can be an IP address or a security group.
Destination	Traffic destination (outbound rule). This parameter is required for an outbound rule. The value can be an IP address or a security group.

### NOTE

The default source IP address **0.0.0.0/0** indicates that all IP addresses can access BMSs in the security group.

## 8.1.2 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 <b>TCP, UDP, ICMP, and All.</b>	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 4: 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

## 8.1.3 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**.  
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.

# 9 Resources and Tags

## 9.1 Tag

### 9.1.1 Overview

To facilitate your management of BMSs, 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 9-1 Example tags




Figure 9-1 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 9-1](#) provides the tag key and value requirements.

**Table 9-1** Tag key and value requirements

Parameter	Requirement	Example Value
Tag key	<ul style="list-style-type: none"><li>• Cannot be left blank.</li><li>• Can only contain letters, digits, underscores (_), and hyphens (-).</li><li>• Contains a maximum of 36 characters.</li></ul>	Organization
Tag value	<ul style="list-style-type: none"><li>• Cannot be left blank.</li><li>• Can only contain letters, digits, underscores (_), periods (.), and hyphens (-).</li><li>• Contains a maximum of 43 characters.</li></ul>	Apache

### 9.1.2 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. Click **Apply for BMS**.
2. 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 9-1](#).

 NOTE

For details about other parameters, see [Creating a Common BMS](#).

## Add Tags on the BMS Details Page


1. In the BMS list, click the name of the target BMS.  
The page showing details of the BMS is displayed.
2. 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 9-1](#).  
You can change the tag value after the tag is added.

### 9.1.3 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. Click **Search by Tag** above the upper right corner of the BMS list to expand the search area.
2. 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.
3. Click  to add multiple tags.  
You can add multiple tags. The system will display BMSs that have all the tags.
4. Click **Search**.  
The system searches for BMSs based on the specified tag keys and values.

#### Filter Resources on the TMS Console

1. On the **Resource Tags** page, set the search criteria, including **Region**, **Resource Type**, and **Resource Tag**.
2. Click **Search**.  
All the resources that meet the search criteria will be displayed in the **Search Result** area.

### 9.1.4 Deleting Tags

If you no longer need a tag, delete it in either of the following ways:

#### Procedure

1. In the BMS list, click the name of the target BMS.  
The page showing details of the BMS is displayed.

2. 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**.

## 9.2 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
Account	Global	You can use the same account in all regions.
Predefined tags	Global	You can use the same predefined tag in all regions.
Key pair	Global or regional	<p>A key pair you create on the management console is associated with the region where it is created.</p> <p>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.</p> <p>For details about key pairs, see <a href="#">Using an SSH Key Pair</a>.</p>
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.

Resource	Type	Description
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.

## 9.3 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, CPU cores, 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, click  .  
The **Service Quota** page is displayed.
4. View the used and total quota of each type of resources on the displayed page.

If a quota cannot meet service requirements, apply for a higher quota.

### How Do I Apply for a Higher Quota?

The system does not support online quota adjustment. If you need to adjust a quota, call the hotline or send an email to the customer service mailbox. Customer service personnel will timely process your request for quota adjustment and inform you of the real-time progress by making a call or sending an email.

Before dialing the hotline number or sending an email, make sure that the following information has been obtained:

- Account name, project name, and project ID, which can be obtained by performing the following operations:  
Log in to the management console using the cloud account, click the username in the upper right corner, select **My Credentials** from the drop-down list, and obtain the account name, project name, and project ID on the **My Credentials** page.
- Quota information, which includes:
  - Service name
  - Quota type
  - Required quota

[Learn how to obtain the service hotline and email address.](#)

# 10 Server Monitoring

---

## 10.1 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 monitoring for BMSs, on which Agent is installed. Agent uses less than 50 MB of memory and 1.5% of CPU resources.

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.

### Constraints

- Agent can only be installed on BMSs running a 64-bit Linux OS.
- Private images do not support this function.

**Table 10-1** lists the Linux images that support server monitoring.

**Table 10-1** Linux images that support server monitoring

OS Type (64-bit)	Version
SUSE	Enterprise11 SP4 and Enterprise12 SP1
CentOS	6.9, 7.2, 7.3, and 7.4
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 *Cloud Eye User Guide*.

# 10.2 Installing and Configuring Agent

## Scenarios

This section describes how to install and configure Agent on a BMS.

## Prerequisites

The BMS is running properly.

## Constraints

Private images do not support this function.

[Table 10-2](#) lists the Linux images that support server monitoring.

**Table 10-2** Linux images that support server monitoring

OS Type	Version
Red Hat	6.5, 6.7, 6.8, 7.2, 7.3, and 7.4
SUSE	11.4 and 12.1
Oracle Linux	6.5, 7.3, and 7.4
CentOS	6.9, 7.2, 7.3, and 7.4
EulerOS	2.2

## Procedure

1. Perform the following steps to create an agency for server monitoring of the BMS:
  - a. On the management console homepage, choose **Service List > Management & Deployment > Identity and Access Management**.
  - b. In the navigation pane on the left, choose **Agency** and then click **Create Agency** in the upper right corner.
    - **Agency Name:** Enter **bms\_monitor\_agency**.
    - **Agency Type:** Select **Cloud service**.
    - **Cloud Service:** This parameter is available if you select **Cloud service** for **Agency Type**. Click **Select**, select **ECS BMS** in the displayed **Select Cloud Service** dialog box, and click **OK**.

- **Validity Period:** Select **Permanent**.
- **Description:** This parameter is optional. You can enter **Support BMS server monitoring**.
- **Permissions:** Locate the region where the BMS resides and click **Modify** in the **Operation** column. In the displayed dialog box, enter **CES** in the **Available Policies** search box. Then select **CES (CES Administrator)** and click **OK**.

 **NOTE**

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

- c. Click **OK**.  
The operations to create an agency for server monitoring of the BMS are complete.
2. Inject the agency.
  - To inject an agency into a new BMS, select the agency created in **1** when you create the BMS.
  - To inject an agency into an existing BMS, click the BMS name to enter its details page, click **Monitoring**, and select the agency created in **1**.
3. Install and configure Agent on the BMS. For details, see "Installing and Configuring the Agent on a Linux ECS or BMS" in *Cloud Eye User Guide*.
4. Log in to the management console and choose **Management & Deployment > Cloud Eye**. On the **Server Monitoring** page, you can view the monitoring data of the BMS.

## 10.3 Monitored Metrics (with Agent Installed)

### Description

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

 **NOTE**

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

### Namespace

SERVICE.BMS

### Metrics

[Table 10-3](#) lists the metrics supported by BMS.



**Table 10-3** Metrics

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
cpu_usage	(Agent) CPU Usage	CPU usage of the monitored object  Obtain its value by checking metric value changes in the <b>/proc/stat</b> file in a collection period.  Run the <b>top</b> command to check the <b>%Cpu(s)</b> value.  Unit: percent	0-100 %	BMS	1 minute
load_averge5	(Agent) 5-Minute Load Average	CPU load averaged from the last 5 minutes  Obtain its value by dividing the <b>load5/</b> value in <b>/proc/loadavg</b> by the number of logical CPUs.  Run the <b>top</b> command to check the <b>load5</b> value in the <b>/proc/loadavg</b> file.	≥ 0	BMS	1 minute
mem_usedPercent	(Agent) Memory Usage	Memory usage of the monitored object  Obtain its value by checking the file <b>/proc/meminfo</b> . <b>Memory Usage = (MemTotal - MemAvailable)/MemTotal</b>  Unit: percent	0-100 %	BMS	1 minute
mountPointPrefix_disk_free	(Agent) Available Disk Space	Available disk space of the monitored object  Run the <b>df -h</b> command to check the data in the <b>Avail</b> 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	<p>Disk usage of the monitored object. It is calculated as follows: <b>Disk Usage = Used Disk Space/ Disk Storage Capacity.</b></p> <p><b>Disk Usage = Used Disk Space/Disk Storage Capacity</b></p> <p>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 (~).</p> <p>Unit: percent</p>	0-100 %	BMS	1 minute
mountPointPrefix_disk_ioUtils and volumePrefix_disk_ioUtils	(Agent) Disk I/O Usage	<p>Disk I/O usage of the monitored object</p> <p>Obtain its value by checking data changes in the thirteenth column of the corresponding device in the <b>/proc/diskstats</b> file in a collection period.</p> <p>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 (~).</p> <p>Unit: percent</p>	0-100 %	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mountPointPrefix_disk_inodeUsedPercent	(Agent) Percentage of Total Inode Used	Percentage of used index nodes on the disk Run the <b>df -i</b> command to check data in the <b>IUse%</b> 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
net_bitRecv	(Agent) Inbound Bandwidth	Number of bits received by this NIC per second Check metric value changes in the <b>/proc/net/dev</b> file in a collection period. Unit: bit/s	$\geq 0$ bit/s	BMS	1 minute
net_bitSent	(Agent) Outbound Bandwidth	Number of bits sent by this NIC per second Check metric value changes in the <b>/proc/net/dev</b> file in a collection period. Unit: bit/s	$\geq 0$ bit/s	BMS	1 minute
net_packetRecv	(Agent) NIC Packet Receive Rate	Number of packets received by this NIC per second Check metric value changes in the <b>/proc/net/dev</b> file in a collection period. Unit: count/s	$\geq 0$ 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 <b>/proc/net/dev</b> file in a collection period. Unit: count/s	$\geq 0$ counts/s	BMS	1 minute

<b>Metric ID</b>	<b>Metric</b>	<b>Description</b>	<b>Value Range</b>	<b>Monitored Object</b>	<b>Monitoring Interval (Raw Data)</b>
net_tcp_total	(Agent) TCP TOTAL	Total number of TCP connections of this NIC	≥0	BMS	1 minute
net_tcp_established	(Agent) TCP ESTABLISHED	Number of ESTABLISHED TCP connections of this NIC	≥0	BMS	1 minute

# 11 Troubleshooting

---

## 11.1 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?

### Symptom

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 persists after you have restarted the BMS for several times, contact the operation administrator.

## 11.2 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.

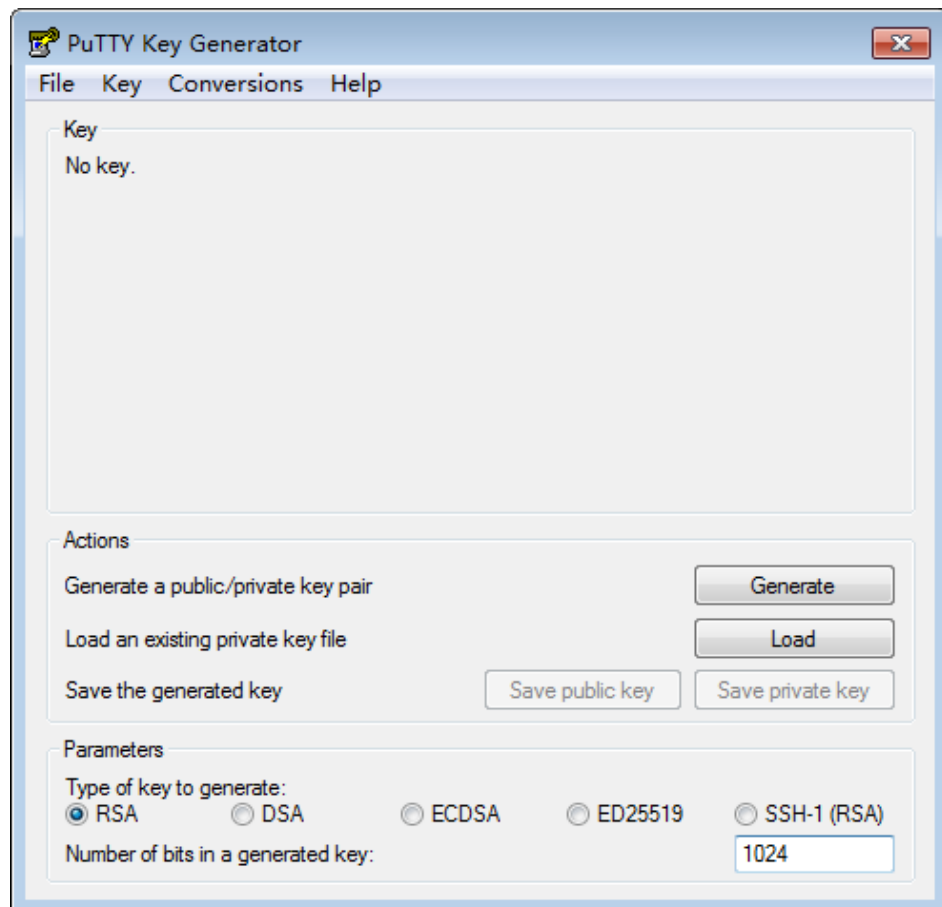
Storing a public key by clicking **Save public key** of PuTTYgen will change the format of the public key content. Such a key cannot be imported to the management console.

## 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.

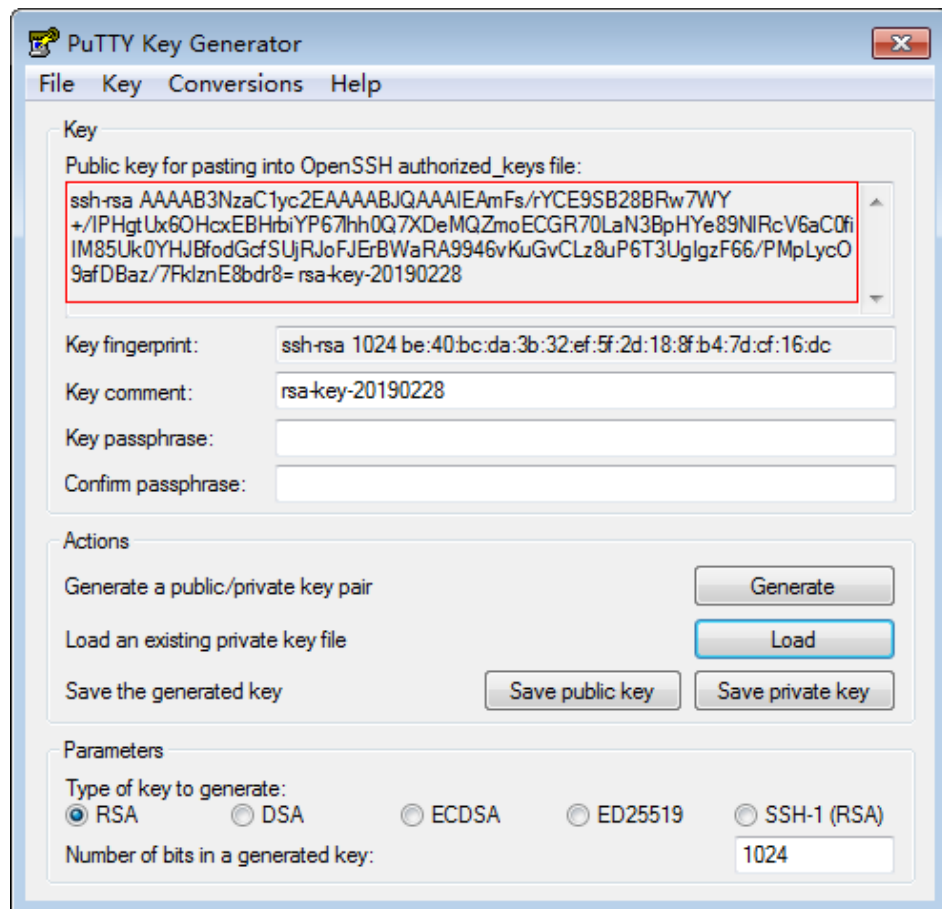
**Figure 11-1** PuTTY Key Generator



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 11-2** is the public key with the format meeting system requirements.

**Figure 11-2** 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**.

## 11.3 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 operation administrator.

## 11.4 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.

# 12 FAQs

---

## 12.1 General FAQ

### 12.1.1 What Are the Restrictions on Using BMSs?

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

## 12.1.2 How Are BMSs Different from ECSs?

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

## 12.1.3 What Are the Differences Between BMSs and Traditional Physical Servers?

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

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

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

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

For more information, see [Table 12-1](#).

**Table 12-1** Comparison between BMS and DeH

Item	BMS	DeH
Whether virtualization is provided	No	Yes
How to use	Use each BMS as a physical server, or install virtualization software on a BMS to create VMs.	Provision ECSs on a DeH.
Specifications	BMS specifications	DeH specifications and ECS specifications
Image	BMS images	ECS images

## 12.1.5 How Do BMSs Ensure Data Security?

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

## 12.1.6 Can I Use APIs to Access BMSs?

Yes.

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

For details, see .

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

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



## 12.1.8 How Do I View and Increase the BMS Quota?

### What Is Quota?

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

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

### How Do I View My Quotas?

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

## How Do I Apply for a Higher Quota?

The system does not support online quota adjustment. If you need to adjust a quota, call the hotline or send an email to the customer service mailbox. Customer service personnel will timely process your request for quota adjustment and inform you of the real-time progress by making a call or sending an email.

Before dialing the hotline number or sending an email, make sure that the following information has been obtained:

- Account name, project name, and project ID, which can be obtained by performing the following operations:  
Log in to the management console using the cloud account, click the username in the upper right corner, select **My Credentials** from the drop-down list, and obtain the account name, project name, and project ID on the **My Credentials** page.
- Quota information, which includes:
  - Service name
  - Quota type
  - Required quota

[Learn how to obtain the service hotline and email address.](#)

## 12.2 Instance FAQ

### 12.2.1 How Long Does It Take to Create a BMS?

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

### 12.2.2 What Do I Do If I Cannot Find My BMS on the Management Console?

#### Symptom

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


#### Causes

Your BMS is not in the selected region or project.

#### Solution

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

**Figure 12-1** Changing the region



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

**Figure 12-2** Switching the project



## 12.2.3 How Can I Obtain board\_type of a BMS Flavor?

### Symptom

The public images supported by BMSs vary depending on BMS flavors. You can view the public images supported by each flavor on the management console or using the API. When you use the API to query for the public images supported by a BMS flavor, you need to enter **board\_type** of the flavor. This section describes how to obtain the value of **board\_type**.

### Solution

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

In the flavor format:

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

Use **AB+First one or more letters of C** as the value of **board\_type**. For example, if the flavor is **physical.s1.large**, **board\_type** is **s1l**. For **board\_type** of more flavors, see [Table 12-2](#).

**Table 12-2** board\_type of a BMS flavor

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

BMS Flavor	board_type
physical.kl1.3xlarge	kl13xl

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

### Symptom

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

### Root Cause

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

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

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

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

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

## 12.2.6 What Are the Advanced Features of BMSs Using EVS Disks?

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

Such BMSs have the following advanced features:

- BMSs booted from EVS disks can be provisioned within about 5 minutes.
- CSBS backups ensure data security.

- BMS rebuilding upon faults is supported, enabling quick service recovery.
- An image of a BMS can be exported to apply configurations of the BMS to other BMSs, eliminating the need to repeatedly configure BMSs.

## Helpful Links

[Creating a BMS Supporting Quick Provisioning](#)

### 12.2.7 Can I Transfer a BMS to Another Account?

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

1. Use the BMS to create a private image.  
For details, see .
2. Share the image with the target account.  
For details, see .
3. Use the target account to create a BMS from the shared private image.
4. Unsubscribe from the original BMS.

### 12.2.8 Is the BMS Host Name with Suffix `novalocal` Normal?

#### Symptom

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

For example, you set the host name to `abc` during BMS creation. [Table 12-3](#) lists the host names (obtained by running the `hostname` command) of BMSs created using different images and those displayed after the BMSs are restarted.

**Table 12-3** Hostnames of BMSs created from different images

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

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

#### Troubleshooting

This is a normal phenomenon. You can ignore it.

The static host name of a Linux BMS is user-defined and injected using Cloud-Init during the BMS creation. According to the test results, Cloud-Init adapts to OSs



differently. As a result, hostnames of some BMSs have suffix **.novalocal**, whereas others do not.

If you really do not want any host names with the suffix **.novalocal**, you can change the hostname. For details, see [Changing the Name of a BMS](#)

## 12.2.9 How Do I Monitor BMS Metrics?

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

## 12.2.10 How Can I Check the BMS Monitoring Status?

The BMS monitoring software is installed in the `/usr/local/telescope` directory. Logs are in the `/usr/local/telescope/log/` directory, in which **ces.log** is the data log and **common.log** is the run log.

- If data is not sent successfully and **403** or **401** is returned, check whether **AccessKey** and **SecretKey** are specified correctly.
- If data is not sent successfully and **500** or other codes are returned, contact the customer service.

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

1. On the management console, choose **Service List > Identity and Access Management**.
2. In the navigation pane on the left, choose **Agencies** and then click **Create Agency** in the upper right corner.
  - **Agency Name:** Enter **bms\_monitor\_agency**.
  - **Agency Type:** Select **Cloud service**.
  - **Cloud Service:** Select **Elastic Cloud Server (ECS) and Bare Metal Server (BMS)** from the drop-down list.
  - **Validity Period:** Select **Unlimited**.
  - **Description:** Enter **Support BMS server monitoring**.
3. Click **Next**. On the **Select Policy/Role** page, search for and select **CES Administrator**.
4. Click **Next**. On the **Select Scope** page, select **All resources** or **Region-specific projects**.

### NOTE

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

5. Click **OK**.

## 12.3 Key Pair and Password FAQ

## 12.3.1 How Do I Change the Password of a BMS in Its OS?

You are advised to reset the password on the console. Alternatively, you can log in to the BMS and change the password in its OS.

### Change the Password for Logging In to a Linux BMS

CentOS 7.5 is used as an example.


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

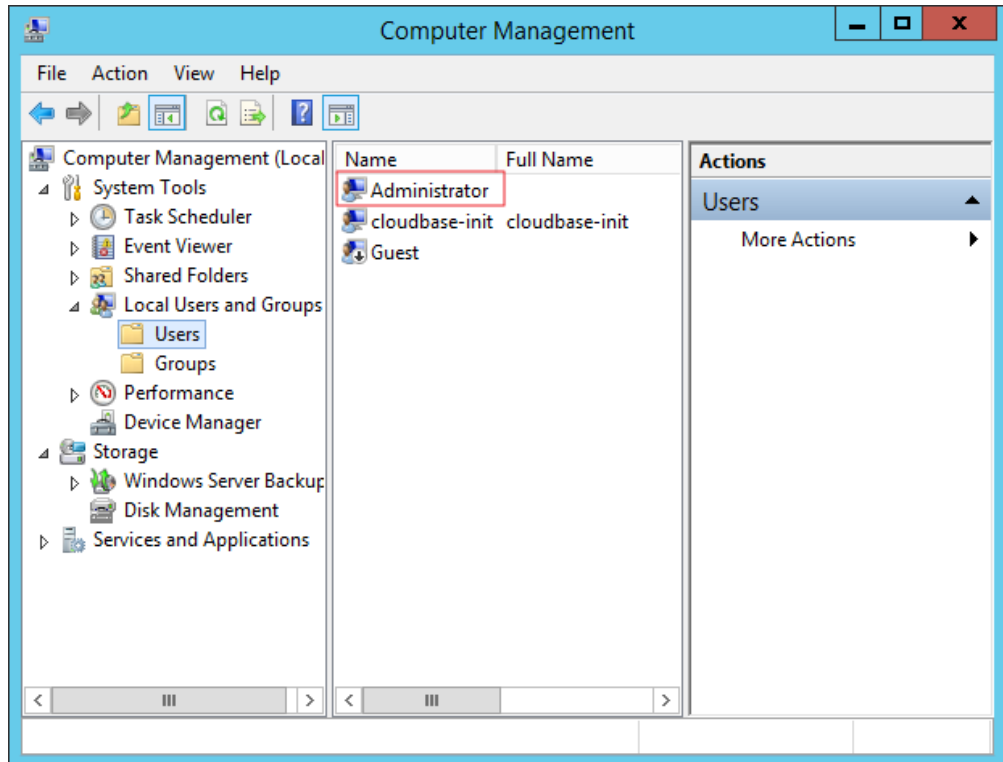
```
[root@~:~77a9 ~]# passwd root
Changing password for user root.
New password: _____
Retype new password: _____
passwd: all authentication tokens updated successfully.
[root@~:~77a9 ~]#
```

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


### Change the Password for Logging In to a Windows BMS

Windows Server 2012 R2 is used as an example.

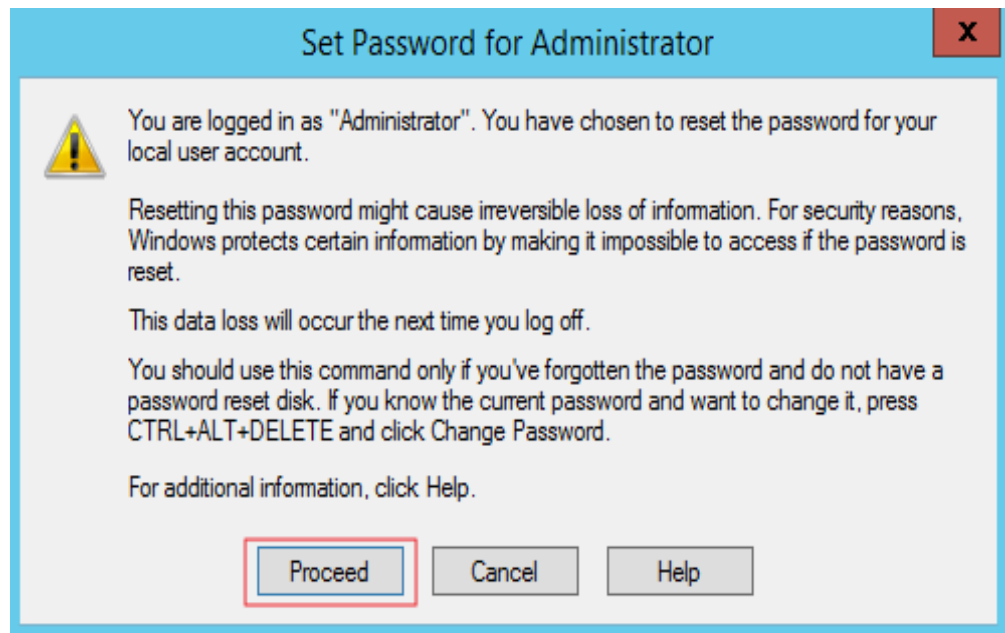
1. [Log in to the BMS remotely using MSTSC.](#)
2. Click  in the lower left corner, choose **Windows PowerShell** and enter `compmgmt.msc` to open **Computer Management**.
3. Choose **System Tools > Local Users and Groups > Users**, right-click the target username, for example, **Administrator** in the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



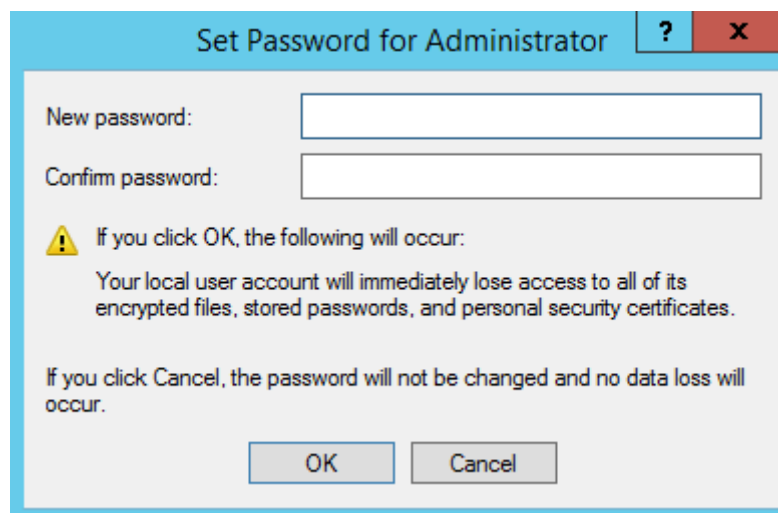
- 4. Select **Set Password**.



- 5. Click **Continue**.



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



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

## 12.4 Login FAQ

### 12.4.1 What Need to Be Prepared for Logging In to a BMS?

#### Windows

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

For details, see [Binding an EIP to a BMS](#).

## Linux

- Obtain the login password.
  - If the authentication mode is SSH key pair, use the key pair you selected when you created the BMS. If you have misplaced your private key file, enable Data Encryption Workshop to reset the key pair.
  - If the authentication mode is password, use the password you set when you created the BMS. If you have forgotten the password, you can reset it.
- Ensure that an EIP is bound to the BMS (not required for remote login).

If you want to log in to the BMS by using an SSH key pair, or a username and a password, you need to ensure that you have bound an EIP to the BMS.

For details, see [Binding an EIP to a BMS](#).

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

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

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

When you create a Linux BMS, you can select the key pair or password authentication mode. If key pair authentication is selected, you can obtain the password in either of the following ways:

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

### 12.4.3 What Do I Do If I Cannot Log In to a Windows BMS?

#### Troubleshoot

If you cannot connect to a BMS using the remote desktop, do as follows:

1. [Check Whether Login Conditions Are Met](#)
2. [Check Network Connectivity](#)
3. [Check Whether the Firewall Is Correctly Configured](#)
4. [Check the Port for Remotely Accessing the BMS](#)
5. [Restart the BMS](#)

## Check Whether Login Conditions Are Met

Check whether the login is properly prepared. For details, see [What Need to Be Prepared for Logging In to a BMS?](#).

## Check Network Connectivity

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

Protocol	Direction	Port Range	Source
ICMP	Inbound	All	0.0.0.0/0

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

## Check Whether the Firewall Is Correctly Configured

The firewall of the BMS must allow the remote connection port (3389 by default). If you have not configured the remote connection port in the inbound rule, you cannot remotely log in to the BMS. In this case, add the remote connection port to the inbound rule of the firewall.

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

## Check the Port for Remotely Accessing the BMS

Check whether port 3389 of the BMS can be accessed.

If the port is inaccessible, check whether this port is allowed by the security group. If it is not, add a security group rule to allow it.

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

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

## Restart the BMS

If the preceding configurations are correct, and you still cannot log in, [restart the BMS](#) on the console. The restart operation will stop the server and interrupt services. Exercise caution when performing this operation.

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

If you are still unable to log in after you perform the preceding operations, record the related BMS information and the time when the login failure occurred. Contact technical support.

## 12.4.4 What Do I Do If I Cannot Log In to a Linux BMS?

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

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

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

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

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

### NOTE

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

If you are still unable to log in to the BMS, record the related BMS information and the time when the login failure occurred. Contact technical support.

## Troubleshoot

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

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

## Check Network Connectivity

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

Protocol	Direction	Port Range	Source
ICMP	Inbound	All	0.0.0.0/0

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

## Check Whether the Security Group Is Correctly Configured

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

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

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

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

1. Log in to the BMS and run the following command to edit the `/etc/fstab` file:  
**`vi /etc/fstab`**
2. Comment out the data disk configuration in the `/etc/fstab` file.  
The `/etc/fstab` file contains information about the file systems and storage devices automatically attached to the BMS when the BMS starts. The configuration information about data disks automatically attached to the BMS needs to be commented out. For example, the last row shown in [Figure 12-3](#) is the data disk configuration to be commented out in the `/etc/fstab` file.

**Figure 12-3** Data disk configuration in the `fstab` file

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

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

## Check the Port for Remotely Accessing the BMS

Check the BMS settings.

1. Check whether the BMS `sshd` process is running.
2. Check whether the BMS rejects access by the local PC.
  - a. Log in to the BMS and run the following command:  
**`vi /etc/hosts.deny`**
  - b. If the IP address of the local PC is contained in the file, the IP address is rejected. In such a case, delete the IP address from the file.
3. Open the `/etc/ssh/ssh_config` file on the local PC and check the default login port. In addition, check whether the value of the `port` field (SSH-enabled port) in the `/etc/ssh/sshd_config` file of the BMS has been changed (the default port number is 22).

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

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



## Check the CPU Load

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

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

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

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

When you use a browser to remotely log in to a BMS, ensure that the browser version meets the requirements listed in [Table 12-4](#).

**Table 12-4** Browser version requirements

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

## 12.4.6 What Do I Do If the Login Page Does Not Respond?

### Symptom

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

### Possible Causes

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

### Solution

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

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

Figure 12-4 Example

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

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

Figure 12-5 Example

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


2. Update the configuration.
  - For SUSE Linux Enterprise Server 12 SP2, Oracle Linux 7.3, Red Hat Enterprise Linux 7.3, CentOS Linux 7.3, and EulerOS 2.2, run the following commands to update the configuration:  
**`stty -F /dev/ttyS0 speed 115200`**  
**`grub2-mkconfig -o /boot/grub2/grub.cfg`**  
**`systemctl enable serial-getty@ttyS0`**
  - For Ubuntu 16.04 Server, run the following commands to update the configuration:  
**`stty -F /dev/ttyS0 speed 115200`**  
**`grub-mkconfig -o /boot/grub/grub.cfg`**

**systemctl enable serial-getty@ttyS0**

3. (Optional) Modify the security configuration file.

If you log in to the BMS through the serial port as user **root**, you need to modify the security configuration file. Add the following information to the end of **/etc/securetty**:

**Figure 12-6** Example



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

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

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

### Symptom

The following symptoms occur:

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

### Possible Causes

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

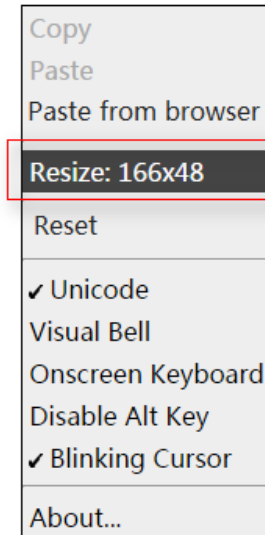
## Solution

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

**Figure 12-7** Selecting Resize: xxx

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

```
linux-8nad:~ # █
```



---

**CAUTION**

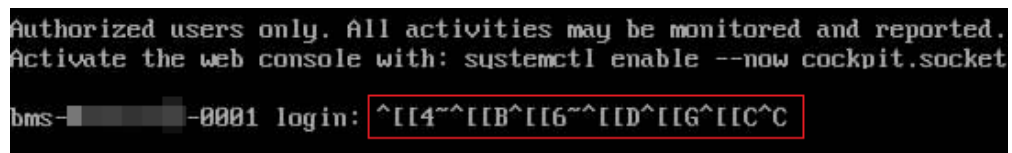
When you are using a text editor such as vim, do not adjust the window size. If you do need to adjust the window size, exit the editor first, adjust the window size, and adjust the console size based on the solution provided in this section.

---

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

### Symptom

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



### Solution

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

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

**setleds -F**

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

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

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

**setleds +num**

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

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

### Symptom

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

UseDNS is an enhanced security feature of SSH and is enabled by default. In such a case, the server obtains the host name of a client by locating the PTR record of the client IP address through a reverse DNS query. Then, the server performs a DNS query based on the obtained client host name and checks whether the obtained IP address is the same as the original IP address, preventing client spoofing. Generally, clients use a dynamic IP address and do not have a corresponding PTR record. Therefore, this feature is invalid for information comparison. However, this feature increases the delay and thereby slows down the client connection.

### Solution

1. Log in to the BMS remotely.
2. Run the following command to open the `/etc/ssh/sshd_config` file:

**vi /etc/ssh/sshd\_config**

3. Find the following field:

```
#UseDNS yes
```

Add the following information in a new line after the field:

```
UseDNS no
```

4. Save the configuration and restart SSH.

**service sshd restart**

## 12.5 Network and Security FAQ

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

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

However, if you do need to allow BMSs of different accounts to communicate with each other through an internal network, you can create a VPC peering connection between VPCs in different accounts. For details, see *Virtual Private Cloud User Guide*.

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

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

## 12.5.3 Are My BMSs in the Same Subnet?

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

## 12.5.4 Can I Associate a BMS with Multiple Security Groups?

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

## 12.5.5 Can BMSs Communicate with ECSs in the Same VPC?

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

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

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

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

They are different in the following ways:

- Generally, the OS default routes preferentially use the primary NICs. If the OS default routes use the extension NICs, network communication will be interrupted. Then, you can check the route configuration to rectify the network communication error.

- By default, only the primary NICs can communicate with the public service zone (zone where PaaS and DNS services are deployed). The extension NICs cannot communicate this zone.

### 12.5.7 Can I Bind Multiple EIPs to a BMS?

Only one EIP can be bound to a NIC. If you want to bind multiple EIPs to a BMS, you can bind them to extension NICs and then perform required operations on the BMS, such as adding policy-based routes or namespaces, to ensure network connectivity.

### 12.5.8 Can I Configure the EIP?

No. The EIP is automatically allocated from the DHCP address pool.

### 12.5.9 Will I Obtain an EIP That Has Been Released?

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

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

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

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

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

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

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

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

#### Cause

If two NICs on the same network segment are added to a BMS, communication between the primary NIC and extension NIC is abnormal because the BMS gateway strictly verifies the source MAC addresses. For example, in [Figure 12-8](#), the primary NIC and extension NIC are both on the 172.22.9.X network segment. A policy-based route needs to be configured to enable communication between the NICs.

**Figure 12-8** Network segment of the NICs

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

## Solution

1. Run the following command to add two routing table names (**net1** and **net2**) and priorities (**252** and **251**) to the **/etc/iproute2/rt\_tables** file:

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

```
252 net1
251 net2
```

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

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

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

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

## 12.5.13 How Can I Configure a Static IP Address for a BMS?

### Scenarios

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

### Procedure

1. Query the IP address and gateway of the BMS.  
Run the following command to query the IP address of the BMS:

```
ifconfig bond0
```



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

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

**ip ro**

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

2. Modify the network configuration file.

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

**Figure 12-9** Modifying the network configuration file

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

#### NOTE

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

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

## 12.5.14 How Do I Configure the DNS Server?

When installing Agent on a BMS, ensure that the DNS server of the BMS runs properly. This section describes how to configure the DNS server and how to verify the DNS server status.

### Linux

1. Log in to the BMS as user **root**.
2. Run the following command to edit the **resolv.conf** file:  
**vi /etc/resolv.conf**
3. Press **i** to enter editing mode and enter **nameserver** *DNS server IP address* before existing **nameserver** configurations.

The format is as follows:

```
nameserver DNS server IP address
```

4. Press **Esc** and enter **:wq** to save the change and exit.
5. Run the following commands to restart the network:  
**rcnetwork restart**  
**service network restart**  
**/etc/init.d/network restart**

### Windows

The following steps use Windows Server 2012 R2 as an example to describe how to configure the DNS server for Windows:


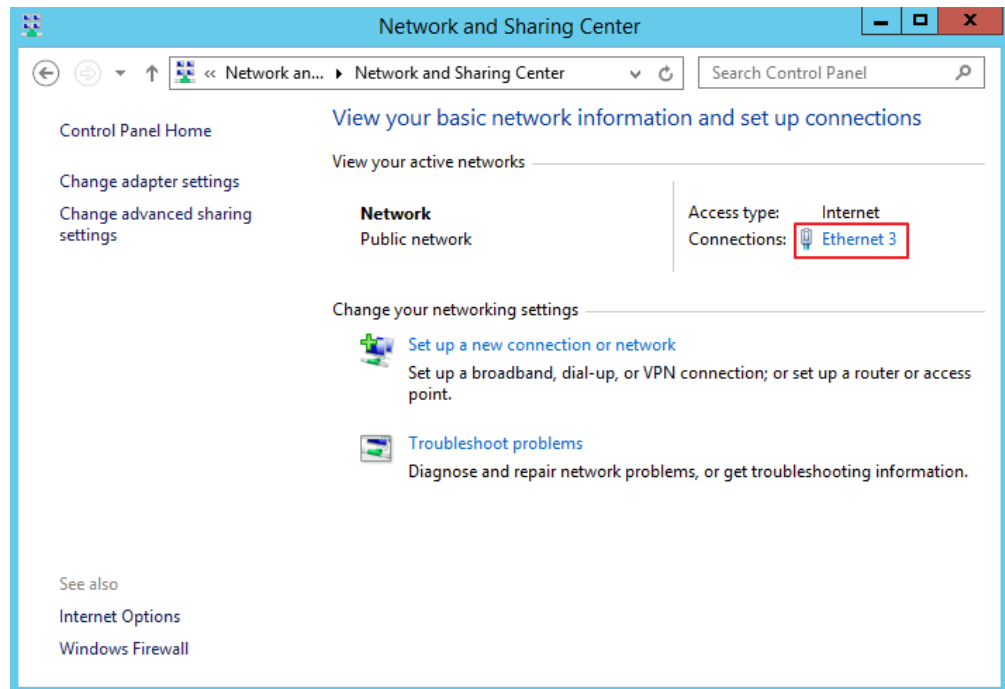

1. Log in to the BMS as user **Administrator**.
2. Click  in the lower left corner to start **Control Panel**.
3. Choose **Network and Internet > Network and Sharing Center**. Then, click the NIC for which you are to configure the DNS server, such as **Ethernet 3**.

Figure 12-10 Network and Sharing Center



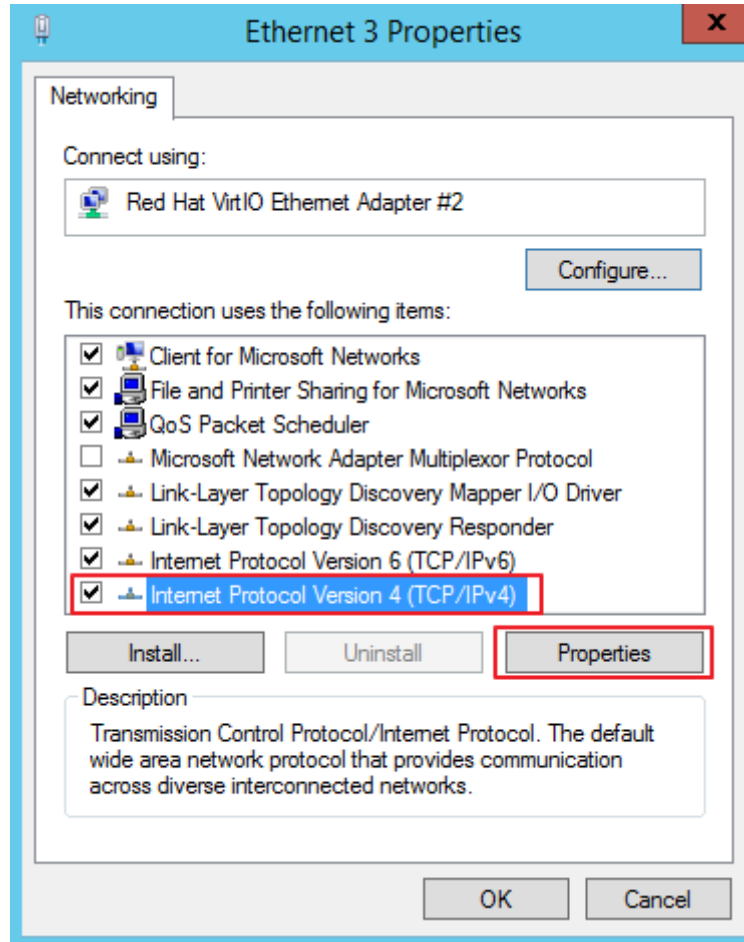
4. Click **Properties**. **Figure 12-11** shows the **Ethernet 3 Status**.

Figure 12-11 Ethernet 3 Status

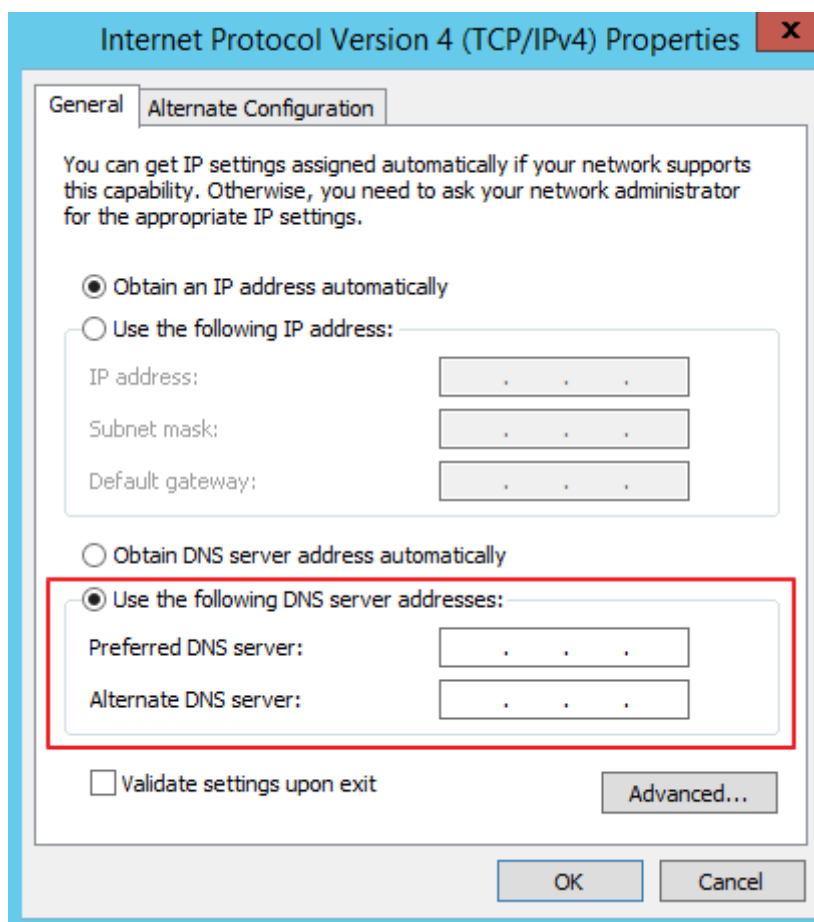



5. In the displayed **Ethernet 3 Status** dialog box, select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**.

**Figure 12-12** Ethernet 3 Properties



6. In the displayed **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, select **Use the following DNS server addresses:** and configure the required parameters shown in **Figure 12-13**. After completing the configuration, click **OK**.

**Figure 12-13** Configuring the DNS server

7. After completing the configuration, click , select **Windows PowerShell**, and enter the **ipconfig /all** command. The configured IP address is displayed in **DNS Servers**.

## 12.5.15 How Do I Change the MTU Value of a Linux BMS NIC?

Maximum Transmission Unit (MTU) specifies the largest packet of data that can be transmitted on a network and ranges from 1280 to 8888 in the unit of byte. If the MTU values of two hosts are different, the transmission may be interrupted, or packet loss may occur. This section describes how to change the NIC MTU values of BMSs running SUSE Linux, CentOS, and Ubuntu.

### SUSE Linux

The following operations use SUSE Enterprise Linux Server 11 SP4 64-bit as an example to describe how to change the MTU value:

1. Log in to the BMS as user **root**.
2. Run the **ifconfig** command to view the NIC that has a bound IP address, for example, **eth0**.
3. Run the following command to open **ifcfg-XXX**.  
**vi /etc/sysconfig/network/ifcfg-XXX**

**NOTE**

*XXX* indicates the NIC obtained in step 2, for example, **eth0**.

4. Press **i** to enter editing mode and add the following statement to set the MTU value of the NIC:  
**MTU=8888**
5. Press **Esc**, enter **:wq!**, and press **Enter** to save and exit the file.
6. Run the following command to restart the network:  
**service network restart**
7. Run the **ifconfig** command to check whether the MTU value has been changed.

## CentOS

The following operations use CentOS 7.5 64-bit as an example to describe how to change the MTU value:

1. Log in to the BMS as user **root**.
2. Run the **ifconfig** command to view the NIC that has a bound IP address, for example, **eth0**.
3. Run the following command to open **ifcfg-XXX**.  
**vi /etc/sysconfig/network-scripts/ifcfg-XXX**

**NOTE**

*XXX* indicates the NIC obtained in step 2, for example, **eth0**.

4. Press **i** to enter editing mode and add the following statement to set the MTU value of the NIC:  
**MTU=8888**
5. Press **Esc**, enter **:wq!**, and press **Enter** to save and exit the file.
6. Run the following command to restart the network:  
**service network restart**
7. Run the **ifconfig** command to check whether the MTU value has been changed.

## Ubuntu

1. Log in to the BMS as user **root**.
2. Run the following command to open the **interfaces** file:  
**vi /etc/network/interfaces**
3. Press **i** to enter editing mode and add the following statement to set the MTU value of the NIC:  
**post-up /sbin/ifconfig/ eth0 mtu 8888**

```
# This file describes the network interfaces available on your system
# and how to activate them. For more information, see interfaces(5).

# The loopback network interface
auto eth0
iface eth0 inet dhcp
post-up /sbin/ifconfig/ eth0 mtu 8888
```

4. Press **Esc**, enter **:wq!**, and press **Enter** to save and exit the file.
5. Run the following command to restart the network:  
**/etc/init.d/networking restart**
6. Run the **ifconfig** command to check whether the MTU value has been changed.

## 12.5.16 How Do I Change the MTU Value of a Windows BMS NIC?

Maximum Transmission Unit (MTU) specifies the largest packet of data that can be transmitted on a network and ranges from 1280 to 8888 in the unit of byte. If the MTU values of two hosts are different, the transmission may be interrupted, or packet loss may occur. This section uses Windows Server 2012 R2 as an example to describe how to change the MTU value of a Windows BMS NIC.


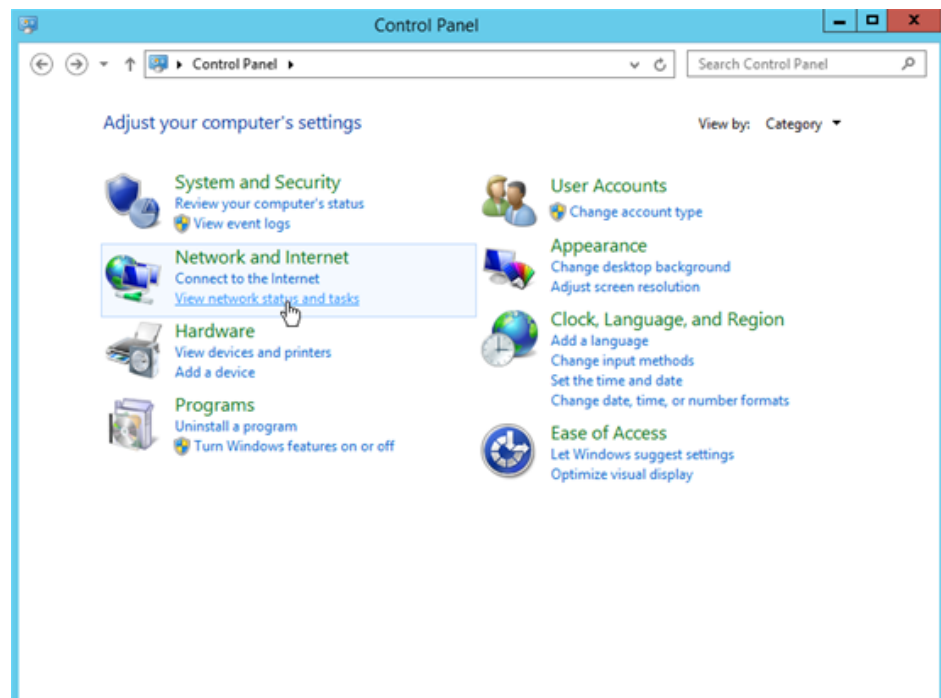
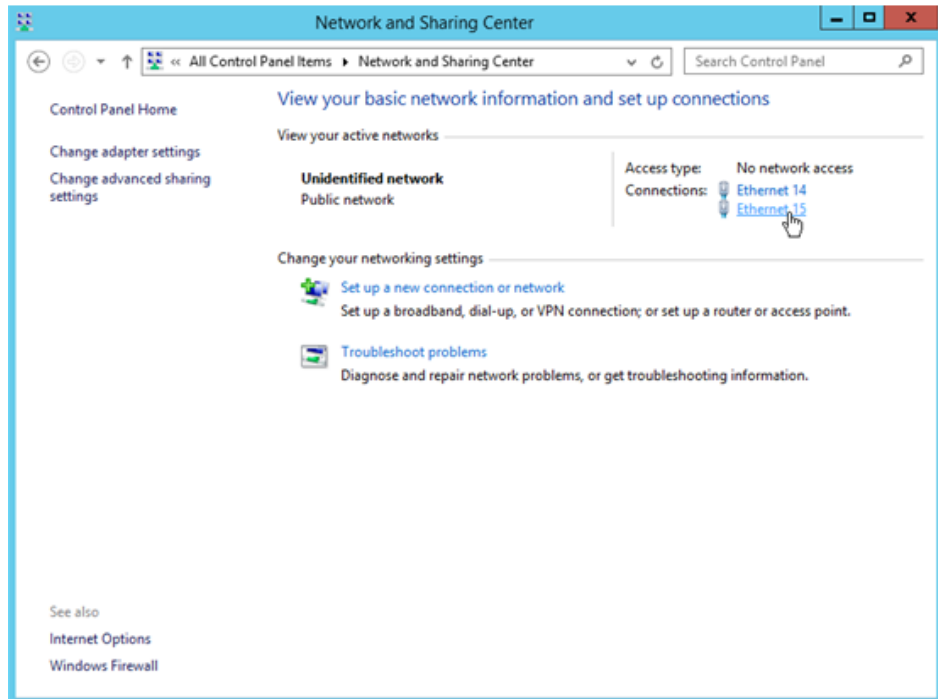
1. Enable Jumbo Packet on the NIC.
  - a. Click  in the lower left corner to start **Control Panel**.

Figure 12-14 Control Panel



- b. Click **View network status and tasks** under **Network and Internet**.


**Figure 12-15** Network and Sharing Center



- c. In the **View your active networks** area, click the target NIC, for example, **Ethernet 15** shown in **Figure 12-15**.  
The page showing the Ethernet 15 NIC status is displayed.
- d. Click **Properties**.  
The page showing the Ethernet 15 NIC properties is displayed.
- e. Click **Configure**. In the displayed dialog box, click the **Advanced** tab.



**Figure 12-16** Inter(R) 82599 Virtual Function #13 Properties




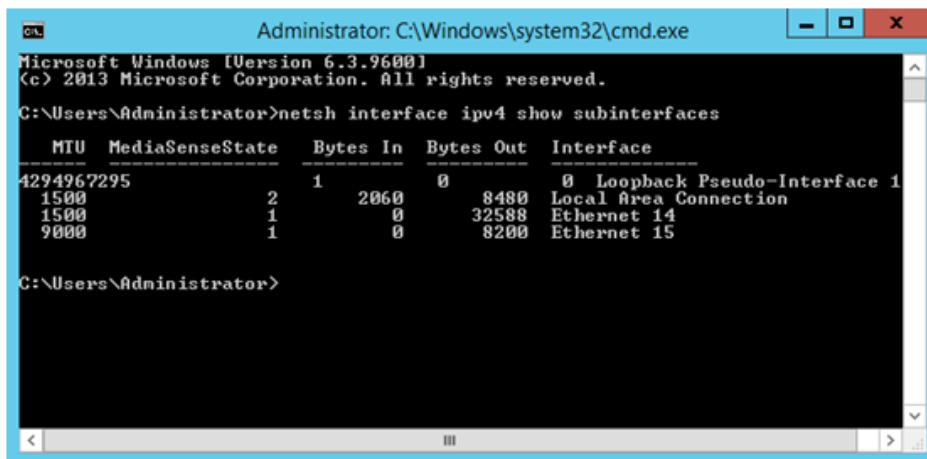
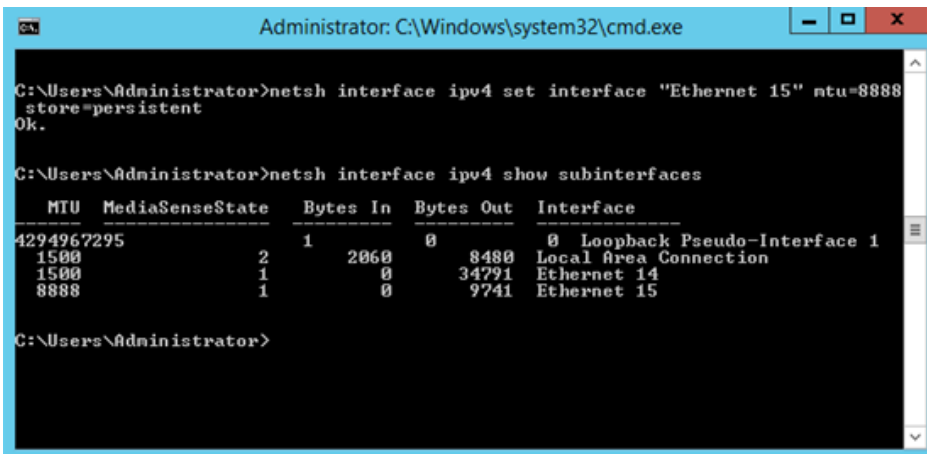
- f. In the **Settings** area, select **Jumbo Packet**. In the **Value** area, select **9014 Bytes**.
  - g. Click **OK**.
2. Change the MTU value.
- a. Click  in the lower left corner, choose **Windows PowerShell**, and run the following command to query the MTU value of the current NIC:  
**netsh interface ipv4 show subinterfaces**
  - b. Obtain the result. The MTU value of the NIC with Jumbo Packet enabled is 9000.

Figure 12-17 Obtaining the NIC MTU value



- c. Run the following command to change the NIC MTU value:  
**netsh interface ipv4 set interface "NIC name" mtu=Changed MTU value store=persistent**  
 For example, to change the MTU value of the Ethernet 15 NIC to 8888, run the following command:  
**netsh interface ipv4 set interface "Ethernet 15" mtu=8888 store=persistent**
- d. Run the following command to query the changed MTU value:  
**netsh interface ipv4 show subinterfaces**

Figure 12-18 Obtaining the changed MTU value



## 12.6 Disk FAQ

### 12.6.1 Can EVS Disks Be Attached to BMSs?

Yes. Ultra-high I/O, high I/O, and common I/O EVS disks can be attached to BMSs.

 NOTE

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

## 12.6.2 What Are the Restrictions for Attaching a Disk to a BMS?

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

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

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

You can call the API to check whether a BMS flavor supports EVS disks. If the value of **baremetal:\_\_support\_evs** in the response is **true**, EVS disks are available in the flavor. If the parameter value is **false** or the parameter is not displayed in the response, EVS disks are unavailable in the flavor.

## 12.6.4 How Do I Change the Disk Identifier in the fstab file to UUID?

### Scenarios

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

 NOTE

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

### Procedure

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

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

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

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

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

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

- If the disk identifier is UUID, no further action is required.
- If the disk identifier is the device name, go to **4**.

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

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

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

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

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

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

Figure 12-19 Checking the mapping between the WWN and drive letter

```
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c8088c61 -> ../../sdb
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c80b2539 -> ../../sde
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c80b685d -> ../../sdc
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c80ba3e9 -> ../../sdg
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c80bb905 -> ../../sdf
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x50000397c810e531 -> ../../sdd
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x600508e000000002ab14603b88fa90b -> ../../sda
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e000000002ab14603b88fa90b-part1 -> ../../sda1
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e000000002ab14603b88fa90b-part2 -> ../../sda2
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e000000002ab14603b88fa90b-part3 -> ../../sda3
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e000000002ab14603b88fa90b-part4 -> ../../sda4
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x600508e000000002ab14603b88fa90b-part5 -> ../../sda5
rwXgwxgwx. 1 root root 9 Mar 20 17:20 wwn-0x68886030000369fafaf17a17502223655 -> ../../sdh
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x68886030000369fafaf17a17502223655-part1 -> ../../sdh1
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x68886030000369fafaf17a17502223655-part2 -> ../../sdh2
rwXgwxgwx. 1 root root 10 Mar 20 17:20 wwn-0x68886030000369fafaf17a17502223655-part3 -> ../../sdh3
rwXgwxgwx. 1 root root 9 Mar 21 14:16 wwn-0x6888603000036b61fa17a17502223655 -> ../../sdo
```

### NOTE

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

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

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


### Local System Disk

The EVS disk device names displayed on the BMS details page on the VPC console are inconsistent with the device names displayed in the BMS OS. To prevent

impact of device name changes on services, you are advised to use EVS disks by UUID.

If EVS disks are specified during BMS allocation, the EVS disk device names displayed on the BMS details page start from **/dev/sdb** and the device names displayed in the BMS OS start after the BMS local disk names, as shown in [Figure 12-20](#).

**Figure 12-20** Device names in the BMS OS



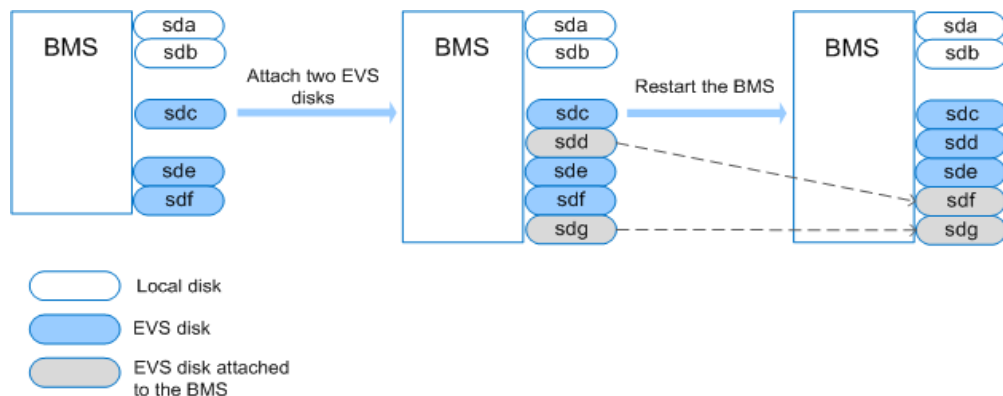
If EVS disks are attached to an allocated BMS, the device names displayed on the BMS details page are those specified by the tenant during disk attaching. After the EVS disks are detached from the BMS, the disks will not be displayed on the BMS details page, and the device names will be released.

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


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

If the BMS OS restarts, the device names displayed in the BMS OS will change based on the number of disks the BMS has and the disk attaching sequence. [Figure 12-21](#) shows the device names displayed in the BMS OS after EVS disks are attached to the BMS (before and after BMS restart). [Figure 12-22](#) shows the device names displayed in the BMS OS after EVS disks are detached from the BMS (before and after BMS restart).

**Figure 12-21** Attaching EVS disks to a BMS



**Figure 12-22** Detaching EVS disks from a BMS




## EVS System Disk

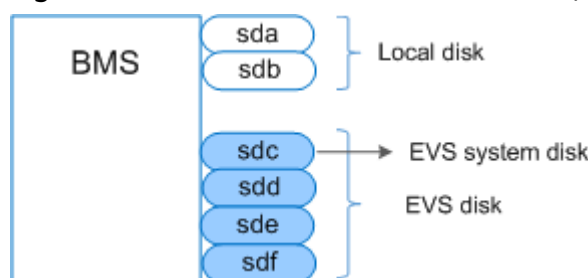
The EVS disk device names displayed on the BMS details page on the VPC console may be inconsistent with the device names displayed in the BMS OS.

If EVS disks are specified during BMS allocation, the EVS disk device names displayed on the BMS details page start from `/dev/sda` and the device names in the BMS OS are displayed in a sequence determined by system scanning. There are two situations as shown in [Figure 12-23](#) and [Figure 12-24](#), and the EVS system disk always has the smallest drive letter of all the EVS disks.

**Figure 12-23** Device names in the BMS OS (situation 1)



**Figure 12-24** Device names in the BMS OS (situation 2)



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


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

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

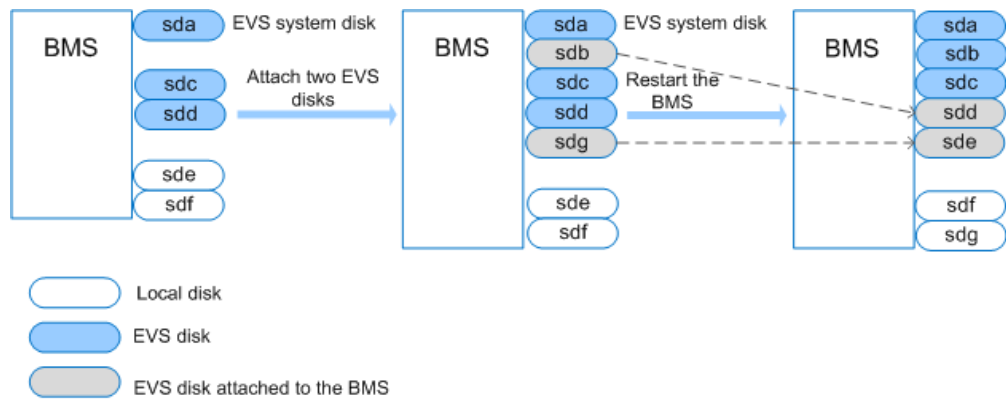
used by other devices. For example, if device names `/dev/sda` and `/dev/sdc` are in use, the device names will start from `dev/sdb`. After EVS disks are detached from the BMSs, if the BMS OS does not restart, the BMS OS will release the device names.

If the BMS OS restarts, the device names displayed in the BMS OS will change based on the number of disks the BMS has and the disk attaching sequence. [Figure 12-25](#) and [Figure 12-26](#) show the device names displayed in the BMS OS after EVS disks are attached to the BMS (before and after BMS restart). [Figure 12-27](#) and [Figure 12-28](#) show the device names displayed in the BMS OS after EVS disks are detached from the BMS (before and after BMS restart).


**Figure 12-25** Attaching an EVS disk (before the BMS restart)



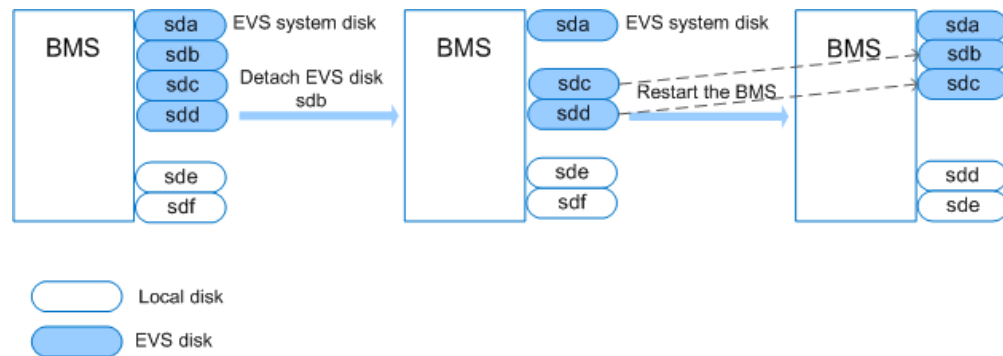
**Figure 12-26** Attaching an EVS disk (after the BMS restart)



**Figure 12-27** Detaching an EVS disk (before the BMS restart)



**Figure 12-28** Detaching an EVS disk (after the BMS restart)



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

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

## 12.6.8 How Can I Restore System Disk Data Using the Snapshot?

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

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

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

Attaching or detaching the system disk is a high-risk operation. You can attach or detach the system disk only when you need to restore the system disk data using



the snapshot. In other cases, you are forbidden to attach or detach the system disk.

## 12.6.10 How Should I Select Storage?

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

- Elastic Volume Service (EVS): provides EVS disks of different QoS configurations to meet performance requirements in various scenarios.

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

Possible causes of this issue are as follows:

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

## 12.7 OS FAQ

### 12.7.1 Can I Install or Upgrade BMS OSs by Myself?

You can reinstall a BMS OS. If an upgrade or patch installation is involved and the kernel version changes, confirm with the cloud service vendor whether drivers, such as RAID controller card drivers and NIC drivers, need to be reinstalled. If the required drivers of the corresponding kernel version are not installed, the OS may fail to start or basic functions of the OS may be unavailable.

### 12.7.2 Can the BMS OS Be Replaced?

No. The BMS OS cannot be replaced.

### 12.7.3 Is a GUI Provided for BMS OSs?

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

### 12.7.4 Is an Upload Tool Delivered with BMS OSs?

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

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

### Symptom



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

### Automatically Updating the Host Name (Recommended)

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

This method has the following restrictions:

- The host name contains a maximum of 63 characters.
- Special characters except hyphens (-), underscores (\_), and periods (.) are not supported.
- Uppercase letters are not supported.
- This method does not apply to Windows BMSs.

1. Log in to the management console, click **Bare Metal Server** under **Computing**.
2. Click the name of the BMS whose name is to be changed.
3. On the displayed page, click  next to **Name**, enter a new name that meets the preceding requirements, and click  to save the change.
4. Log in to the BMS OS and run the following command to enable automatic hostname synchronization:

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

Check that automatic synchronization is enabled.

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

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

[METADATA]
enable_preserve_hostname = False
auto_synchronize_hostname = True

[IB]
enable_ib = True

[ROCE]
enable_roce = True
```

5. Log in to the management console again. Locate the row that contains the BMS, click **More** in the **Operation** column, and select **Restart**.

After about 10 minutes, verify that the BMS is restarted and its hostname is automatically updated.

 NOTE

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

## Manually Updating the Host Name

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

For example, if the changed host name is *new\_hostname*, perform the following steps:

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

4. Change the BMS network configuration script **bms-network-config.conf**.  
The value of parameter **enable\_preserve\_hostname** in the **bms-network-config.conf** file is **False** by default, indicating that the host name is updated each time the board resets. To disable this function, change its value to **True**.
  - a. Change the value of **enable\_preserve\_hostname** in the **bms-network-config.conf** file to **True**:

```
sed -i 's/enable_preserve_hostname.*/enable_preserve_hostname = True/g' `find / -name bms-network-config.conf`
```
5. (Optional) For SUSE, modify the configuration file **/etc/sysconfig/network/dhcp**.
  - a. Run the following command to edit the **/etc/sysconfig/network/dhcp** configuration file:

```
sudo vim /etc/sysconfig/network/dhcp
```
  - b. Set the value of **DHCLIENT\_SET\_HOSTNAME** to **no** to ensure that DHCP does not automatically allocate host names.

```
DHCLIENT_SET_HOSTNAME="no"
```
  - c. Run the following command to save and exit the configuration file:

```
:wq
```
6. Run the following command to restart the BMS:

```
sudo reboot
```
7. Run the following command to check whether the static host name is changed:

```
sudo hostname
```

If the changed host name *new\_hostname* is displayed in the command output, the host name is changed and the new name permanently takes effect.

## 12.7.6 How Do I Set the Password Validity Period?

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

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

1. Log in to the BMS OS and run the following command to query the password validity period:

```
vi /etc/login.defs
```

The value of parameter **PASS\_MAX\_DAYS** indicates the password validity period.
2. Run the following command to change the value of parameter **PASS\_MAX\_DAYS** in **1**:

```
chage -M 99999 user_name
```

*99999* is the validity period of the password, and *user\_name* is a system user. You are advised to set the password validity period as needed and change it on a regular basis.

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

Figure 12-29 Configuration verification

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

## 12.7.7 How Do I Set SSH Configuration Items?

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

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



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

- a. Modify the Cloud-Init configuration file.

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

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

disable_root: 0
ssh_pwauth: 1
```

In the preceding information:

- If the value of `lock_passwd` is set to `false`, user password is not locked.
- `disable_root` specifies whether to disable remote SSH login as user `root`. Set the value to `0`, indicating that the remote SSH login as user `root` is enabled (In some OSs, value `true` indicates disabled and `false` indicates enabled).

- **ssh\_pwauth** specifies whether to support SSH password login. Set this parameter to **1**, indicating that SSH password login is supported.
- b. Run the following command to open the **/etc/ssh/sshd\_config** file using the vi editor:

**vi /etc/ssh/sshd\_config**

Change the value of **PasswordAuthentication** in the **sshd\_config** file to **yes**.

 **NOTE**

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

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

**vim /etc/shadow**

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

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

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

 **NOTE**


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

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

### Cause

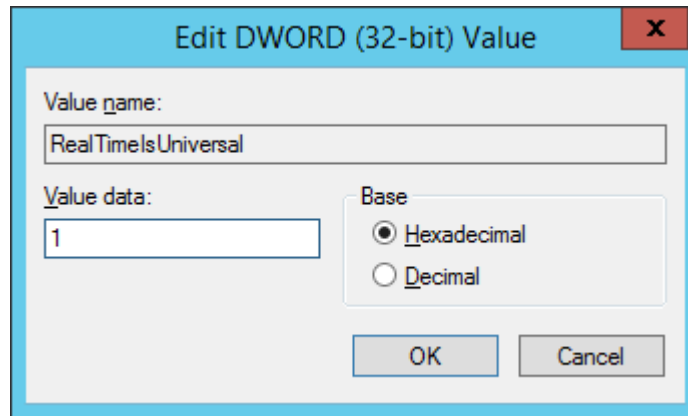
Linux uses the time of the motherboard CMOS chip as the Coordinated Universal Time (UTC) and determines the system time based on the configured time zone. However, Windows uses the CMOS time as the system time directly without converting it based on the time zone.

### Solution

1. Log in to the Windows BMS.
2. Click  in the lower left corner, choose **Windows PowerShell**, and enter **regedit.exe** to open the registry.
3. In the displayed **Registry Editor** window, choose **HKEY\_LOCAL\_MACHINE > SYSTEM > CurrentControlSet > Control > TimeZoneInformation**.

4. In the right pane, right-click a blank area and choose **New > DWORD (32-bit) Value** to add a REG\_DWORD code. Set its name to **RealTimeIsUniversal** and value to **1**.


**Figure 12-30** Adding a code



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

## 12.7.9 How Can I Activate a Windows BMS?

Perform the following operations to manually activate a Windows BMS:

1. Log in to the Windows BMS.
2. Click  in the lower left corner and choose **Windows PowerShell**.
3. Run the following command to configure the IP address of the KMS server:  
**slmgr -skms x.x.x.x**  
x.x.x.x indicates the IP address of the KMS server. Contact the operation administrator to obtain the IP address.
4. Run the following command to check whether the BMS has been activated:  
**slmgr -ato**  
If error 0xC004F074 occurs, the BMS cannot be activated. In such an event, go to [5](#).
5. Verify that the time in the BMS is the same as the standard time. If the time is significantly different, the BMS cannot be activated.
6. Run the following command on the BMS to check whether the link between the BMS and the KMS server port is reachable:  
**telnet x.x.x.x 1688**  
If the link is unreachable, port 1688 is not enabled on the BMS firewall. You must disable the firewall or enable port 1688 on the firewall. If the BMS has any security software such as safedog, stop using it.
7. Run the following command to check whether the BMS has been activated:  
**slmgr -ato**

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

### Scenarios

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

A complete SID contains:


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

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

S	1	5	21-287469276-4015456986-3235239863	500
The string is a SID.	SID version	SID authority, which is NT in this example	SID sub-authorities	Accounts and groups in the domain

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

### Procedure

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

**Figure 12-31** Querying the original SID

```

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

```

3. Modify the Cloudbase-Init configuration files.
  - a. Open the **cloudbase-init.conf** and **cloudbase-init-unattend.conf** files.



File directory: C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf

- b. Add **first\_logon\_behaviour=no** to both files.

```
[DEFAULT]
username=Administrator
groups=Administrators
first_logon_behaviour=no
netbios_host_name_compatibility=false
metadata_services=cloudbaseinit.metadata.services.httpserv
inject_user_password=true
...
```

- c. Delete **cloudbaseinit.plugins.common.sethostname.SetHostNamePlugin** from the **cloudbase-init-unattend.conf** configuration file.

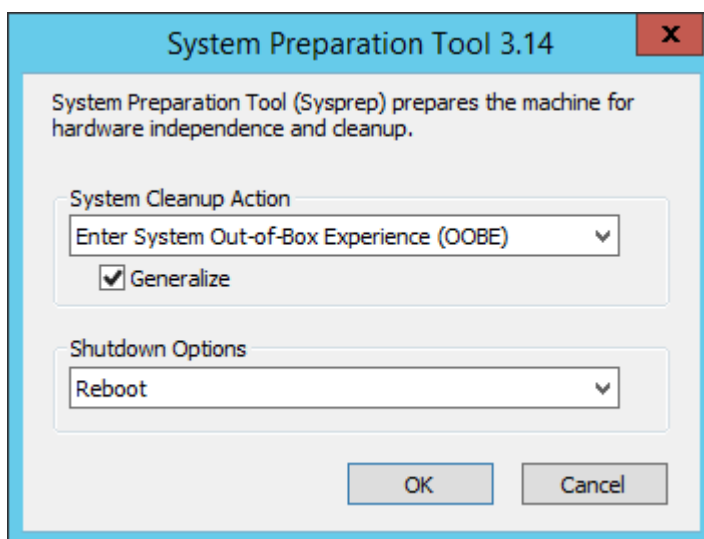
**Figure 12-32** Modifying the configuration file

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

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

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

**Figure 12-33** System Preparation Tool settings



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

**Figure 12-34** Querying the new SID

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

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

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

### Scenarios

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

### Solution

1. Log in to the BMS OS.
2. Run the following command to check the current kernel version:

```
uname -r
```

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

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

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

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

4. Run the following command to set the kernel that is started by default, for example, CentOS Linux (3.10.0-327.el7.x86\_64) 7 (Core):

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

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

```
grub2-editenv list
```

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

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

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

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

### Scenarios

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

### Procedure

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

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

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

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

3. Run the following command to create a 200 GB image file for storing logs.  
**dd if=/dev/zero of=disk.img bs=1M count=200000**

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

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

```
losetup /dev/loop0 disk.img
mkfs.ext4 /dev/loop0
```

```
root@bms-ubuntu-0001:~# losetup /dev/loop0 disk.img
root@bms-ubuntu-0001:~# mkfs.ext4 /dev/loop0
mke2fs 1.42.9 (4-Feb-2014)
Discarding device blocks: done
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
12804096 inodes, 51200000 blocks
2560000 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
1563 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872

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

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

```
root@bms-ubuntu-0001:~# mount disk.img log/
root@bms-ubuntu-0001:~# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda3       819G  1.7G  784G   1% /
none            4.0K   0    4.0K   0% /sys/fs/cgroup
udev            158G   12K  158G   1% /dev
tmpfs           32G   1.1M   32G   1% /run
none            5.0M   0    5.0M   0% /run/lock
none            158G   0    158G   0% /run/shm
none            100M   0    100M   0% /run/user
/dev/sda1        922M   54M   806M   7% /boot
/dev/loop1      193G   60M   183G   1% /root/log
```

6. Create a file in the log directory.

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

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

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

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

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

Connection closed by foreign host.

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

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

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

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

### Scenarios

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

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

### Upgrade Scenario

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

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

As shown in the following figure, **3.10.0-327.el7.x86\_64** is the default kernel version and **3.10.0-862.3.2.el7.x86\_64** is the upgraded kernel version.

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

### Emergency Settings After Kernel Upgrade

1. Run the following commands to set the original kernel version as the default startup kernel and verify the modification result:  
**grub2-set-default "CentOS Linux (3.10.0-327.el7.x86\_64) 7 (Core)"**



### grub2-editenv list

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

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

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

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

## 12.7.14 How Do I Increase the Swap Partition Size?

### Scenarios

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

#### NOTE

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

### Procedure

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

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

The size of the swap partition is 3 GB.

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

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

```
chmod 600 /swapfile
```

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

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

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

The size of the swap partition is 8 GB.

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

### Scenarios

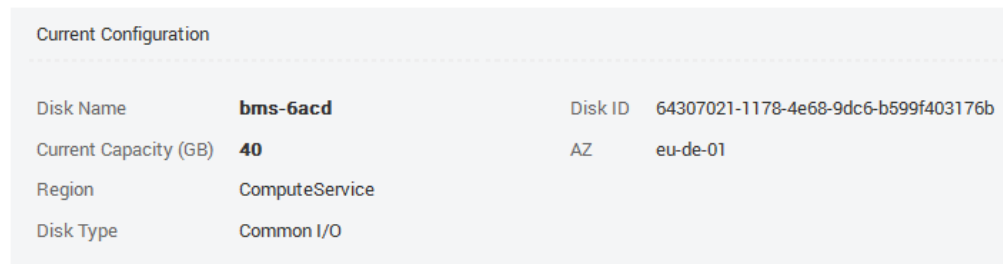
If the root partition size of a BMS which is quickly provisioned cannot meet service requirements, you can increase the root partition size by performing the operations in this section.

### Procedure

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

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

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



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

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

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

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

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

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

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



```

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

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

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

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

Calling ioctl() to re-read partition table.

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

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

```

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

4. Create a 100 MB configdriver partition in the system disk.

Assume that the available sector range is 83755008-293601279, the initial value of the new partition is equal to the maximum value of the available sector minus 200000, that is, 293401279. The maximum value of the new partition is 293601279.

```

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

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

Calling ioctl() to re-read partition table.

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

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

```

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

5. Run the following command to expand the root partition:

**growpart /dev/sdf 3**

```

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

The size of the extended root partition is as follows:

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

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

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

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

- Run the following command to restore the configdriver partition:

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

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

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

The capacity expansion of the BMS root partition is complete.

## 12.7.16 Common Linux Commands

### lsblk

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

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

```
lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda  202:0  0  40G  0 disk
├─sda1 202:1  0   4G  0 part [SWAP]
└─sda2 202:2  0  36G  0 part /
sdb  202:16  0  10G  0 disk
```

The parameters are as follows:

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

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

### Scenarios

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

### Procedure

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

```
lvm vgcfgrestore sysvg
```

```
lvm pvscan
```

```
lvm vgscan
```


```
lvm vgchange -ay
```

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

### Scenarios


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

### Procedure

1. Log in to the Windows BMS and create a reserved administrator account. Choose **Control Panel > User Accounts > Change account type > Add** to add an administrator account.
2. Click  in the lower left corner, choose **Windows PowerShell**, and enter **regedit.exe** to open the registry.

In the **Registry Editor** window, choose **HKEY\_LOCAL\_MACHINE > SOFTWARE > Cloudbase Solutions > Cloudbase-Init** and delete all folders from **Cloudbase-Init**.

**Figure 12-35** Deleting Cloudbase-Init folders



3. Manually delete bond information.
  - a. Delete Team1.

Run the **Get-NetAdapter** command to check whether Team1 is there.

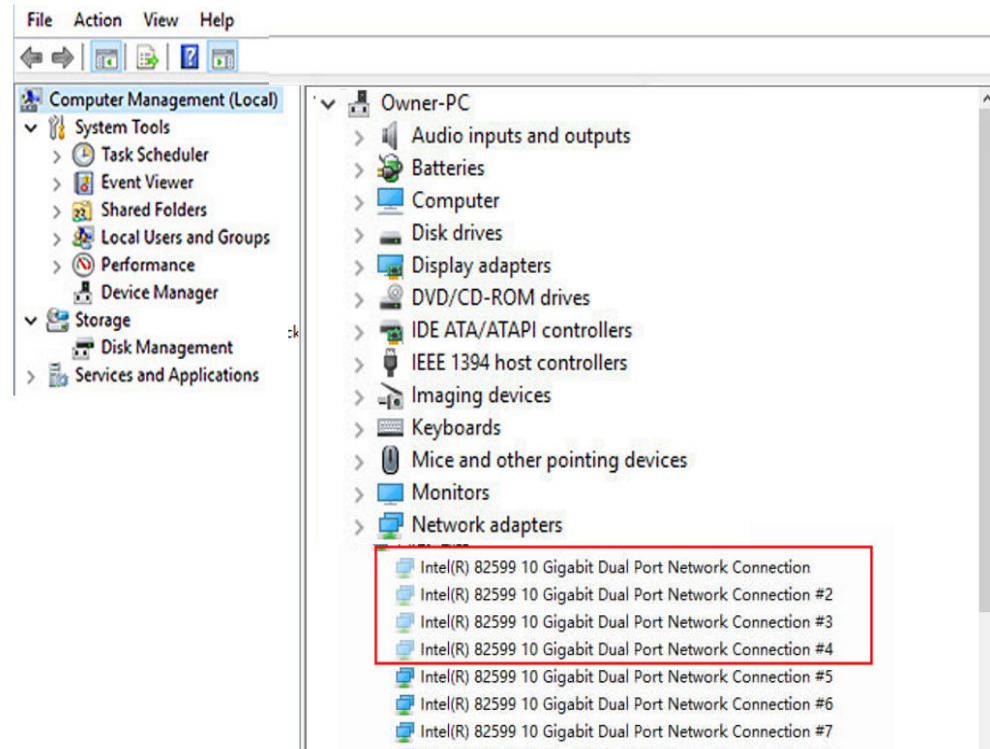
If there is a Team1, open **Windows PowerShell** as an administrator and run the following command to delete it:

**Remove-NetLbfoTeam -Name Team1**

- b. Delete hidden network adapters.

Choose **Computer Management > Device Manager > View > Show hidden devices > Network adapters**. Right-click the adapter (it should be displayed in gray) and choose **Uninstall device**. Uninstall other grayed out adapters, one by one, in the same way.

Figure 12-36 Network adapters



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

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

# A Change History

Released On	Description
2023-3-30	<p>This issue is the fourth official release.</p> <p>Added the following content:</p> <ul style="list-style-type: none"><li><a href="#">Features and Constraints</a></li><li><a href="#">Introduction</a></li><li><a href="#">Disk Types</a></li><li><a href="#">Overview</a></li><li><a href="#">Tag</a></li><li><a href="#">Resource Location</a></li><li><a href="#">Overview</a></li></ul> <p>New questions and answers in <a href="#">FAQs</a>.</p>
2023-02-08	<p>This issue is the third official release.</p> <p>Added <a href="#">Billing</a>.</p>
2022-11-30	<p>This issue is the second official release.</p> <p>Modified the following content:</p> <p>Moved security groups, key pairs, and passwords from other sections in <a href="#">Security</a> to <a href="#">Identity Authentication and Access Control</a>.</p>
2021-08-25	<p>This issue is the first official release.</p>