

Copyright © Huawei Technologies Co., Ltd. 2020. All rights reserved.

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

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Contents

1 Virtualization.....	1
1.1 Overview.....	1
1.2 VMware on BMS.....	2
1.2.1 Solution Overview.....	2
1.2.2 Environment Preparations.....	3
1.2.3 Restrictions.....	5
1.2.4 Deployment Process.....	5
1.2.5 Creating and Configuring the First ESXi BMS.....	6
1.2.6 Creating a Windows Jump VM.....	6
1.2.7 Deploying the DNS and NTP Service VMs.....	7
1.2.8 Deploying vCenter Server Appliance.....	7
1.2.9 Creating Other ESXi BMSs.....	9
1.2.10 Deploying VMware vSAN.....	10
1.2.11 Deploying NSX.....	11
1.2.12 Deploying Other VMware Components.....	16
1.2.13 FAQs.....	18
1.3 XenServer on BMS.....	20
1.3.1 Solution Overview.....	20
1.3.2 Preparing for the Deployment.....	21
1.3.3 Buying a BMS.....	21
1.3.4 Creating a User-defined Network.....	22
1.3.5 Configuring a VPC NIC for the BMS.....	23
1.3.6 Configuring the User-defined VLAN NIC for the BMS.....	23
1.3.7 Configuring the vNIC for the Service VM.....	29
1.3.8 (Optional) Buying a NAT Gateway and Adding a DNAT Rule.....	30
1.4 Hyper-V on BMS.....	31
1.4.1 Solution Overview.....	31
1.4.2 Preparing for the Deployment.....	31
1.4.3 Purchasing a BMS.....	32
1.4.4 (Optional) Creating a User-defined Network and NAT Gateway and Adding a DNAT Rule.....	33
1.4.5 Deploying the Hyper-V Role.....	33
1.4.6 Creating a Hyper-V vSwitch.....	37
1.4.7 Creating a Hyper-V VM.....	40

1.4.8 Configuring the vNIC for the Service VM.....	44
2 Monitoring.....	46
2.1 Overview.....	46
2.2 Installing and Configuring the Agent for an Existing BMS.....	47
2.2.1 Installing the Agent.....	47
2.2.2 (Optional) Managing the Agent.....	50
2.3 Monitoring Data.....	51
2.4 Supported Monitoring Metrics (with Agent Installed).....	52
2.5 Supported Monitoring Metrics (with Agent Installed and Simplified Metrics).....	73
2.6 FAQs.....	77
2.6.1 Why Does Not the Cloud Eye Console Display Monitoring Data or Why Is There a Delay in Data Display After Agent Is Installed and Configured?.....	77
2.6.2 How Do I Create an Agency for Server Monitoring of the BMS?.....	77
3 Backup.....	79
3.1 Overview.....	79
3.2 Creating a Backup Policy.....	81
3.3 Purchasing a BMS.....	84
3.4 Creating a BMS Backup.....	85
3.5 Viewing Backups and Restoring Data.....	86
A Change History.....	89

1 Virtualization

1.1 Overview

BMSs provide both the stability of traditional physical servers and the high scalability of cloud resources. BMSs meet your requirements for high-performance computing and help you build hybrid clouds. Bare Metal Servers (BMSs) have all the features and advantages of physical servers and support secondary virtualization.

Advantages

- Easing your concern about migrating services to the cloud
You need to use computing resources of various types and forms in the cloud. Only a combination of physical servers and VMs, instead of VMs alone, can meet requirements in complex application scenarios. Some special application scenarios have high requirements on host performance and stability, or have strict requirements on data security and supervision. VMs cannot meet customer requirements and exclusively used physical servers are required. AnyStack on BMS provides you with a more reliable, comprehensive, and flexible way to migrate your services to the cloud.
- Enhancing your sense of security
Some enterprises have successful experience in virtualization services deployed in private cloud data centers and have many talents in virtualization. When migrating services to the public cloud, these enterprises attach great importance to security and transformation of their talent skills. With AnyStack on BMS, you can smoothly migrate VMs, service load, or your data center to the public cloud without making any change to your services, personnel skills, or existing O&M tools and experience.
- Providing more flexible load distribution and quicker service expansion
Service VMs running on BMSs can access various public cloud products, including computing, database, monitoring, and security services. Expanding and reducing the host capacity takes as less as several minutes, making it easy to develop services fast and meet demands in peak and off-peak periods. In the public cloud environment, different workload is processed separately. Infrastructure resources are scalable to cope with burst traffic. Using public

cloud resources for DR and backup ensures service data security and reliability.

Supported Hypervisors

- XenServer
- VMware
- Hyper-V

Constraints

This feature is available only for customers with a large business volume.

1.2 VMware on BMS

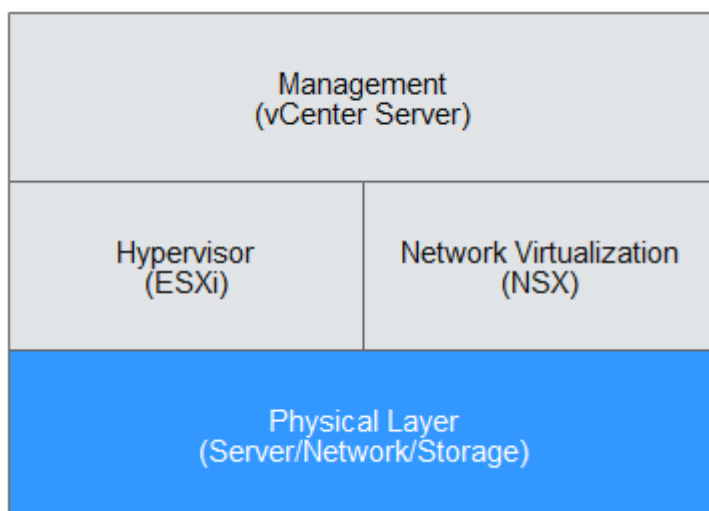
1.2.1 Solution Overview

Bare Metal Server (BMS) provides dedicated physical servers with excellent computing performance to meet your requirements for high performance and stability in core application scenarios. The BMS service allows you to provision VMware ESXi BMSs. You can build a VMware private cloud based on the BMS cluster or migrate your VMware data center to the public cloud. You can enjoy the scalable resources and various services of the public cloud as well as the mature computing and network virtualization functions and the resource management system of VMware. This allows you to expand or migrate your services to the public cloud, or back up your resources in the public cloud.

This document provides a reference solution for deploying the VMware system using BMSs. You can also deploy the VMware system in a different mode based on the capabilities provided by the public cloud. This document applies to users with experience in deploying the VMware system.

Logical Architecture

Figure 1-1 Logical architecture of VMware



As shown in **Figure 1-1**, the management plane, hypervisor, and network virtualization layer are provided by VMware (VMware vCenter Server®, VMware ESXi®, and VMware NSX®, respectively). The public cloud provides the physical infrastructure required by the VMware system.

- The BMS service provides BMSs using the **physical.io2.4xlarge** flavor. It has 44 physical cores, 384 GB memory, and 9600 GB local disks, providing excellent computing and storage capabilities.
- The user-defined network is designed for the Virtualization on BMS solution and provides functions similar to Virtual Private Cloud (VPC). Through the user-defined network, service VMs provisioned by VMware can access:
 - Elastic Cloud Server (ECS)
 - BMS
 - Direct Connect
 - Internet

 **NOTE**

VMware vSphere Replication is used for VM replication, and VMware vCenter Site Recovery Manager for restoration and disaster recovery (DR). Deploy them based on site requirements.

Licenses and Technical Support

Currently, the BMS service only supports automatic provisioning of ESXi BMSs. You need to install upper-layer management suites and their licenses.

If you have any problem with VMware products, contact VMware technical support.

1.2.2 Environment Preparations

Software Environment

Table 1-1 Software versions

Component	Version
VMware ESXi	6.5.0
VMware vCenter Server Appliance	6.5.0
VMware NSX Manager	6.4.3
VMware vSphere Replication	6.5.1
VMware vCenter Site Recovery Manager	6.5.1

Network Environment

NOTE

The IP address segments are examples only. You need to obtain the information from the BMS metadata.

Eight network ports are configured for the provisioned BMS, two for accessing the VPC network and six for configuring the QinQ network. vmnic0 and vmnic1 in the example are two network ports for connecting to the VPC network. Different from two network ports connecting to the VPC network, these two ports become uplinks of vSwitch0 by default after provisioned and enabled on the BMS.

The following example is a typical VMware architecture, which consists of three BMSs in cluster mode on which the management and service VMs are deployed. Each server has eight physical network ports configured and each BMS connects to a VPC subnet. You need to plan and configure the network planes used for management, vmotion, vSAN, and VXLAN on the NICs. The following method is for reference only.

Table 1-2 Network plane planning

Type	Name	IP Address Segment	Gateway	Description
VPC subnet	esx-primary	192.168.0.0/24	192.168.0.1	BMS primary NIC
User-planned VLAN	DPortGroup-mgmt	11.11.11.0/24	11.11.11.1	Management and vMotion plane
	DPortGroup-vxlan	11.11.13.0/24	11.11.13.1	vxlan
	DPortGroup-vsan	11.11.12.0/24	11.11.12.1	vSAN plane
	hb-edge-internal	11.11.8.0/24	11.11.8.1	Service plane

The following table lists the vNICs of the ESXi BMS.

Table 1-3 vNICs

Network	IP Address	vSwitch uplink
esx-primary	192.168.0.10	vmnic0/vmnic1
DPortGroup-mgmt	11.11.11.101	vmnic2/vmnic3
DPortGroup-vxlan	11.11.13.101	vmnic4/vmnic5
DPortGroup-vsan	11.11.12.101	vmnic6/vmnic7

The following table lists the IP addresses of the management components connected to the DPortGroup-mgmt network.

Table 1-4 IP addresses of components

Component	IP Address	Domain Name
vCenter Server	11.11.11.3	Optional
NSX manager	11.11.11.4	-
NSX controller	11.11.11.200 11.11.11.201 11.11.11.202	-
DNS/NTP	11.11.11.6	-
Nsx-edge	11.11.11.30	-
Jump VM	11.11.11.2	-

 **NOTE**

- The VPC name, subnet, and gateway in [Table 1-2](#) are defined and input on the console by the tenant during VM creation.
- You can plan the IP addresses of the vNICs in [Table 1-3](#).
- You can plan and allocate the IP addresses in [Table 1-4](#).

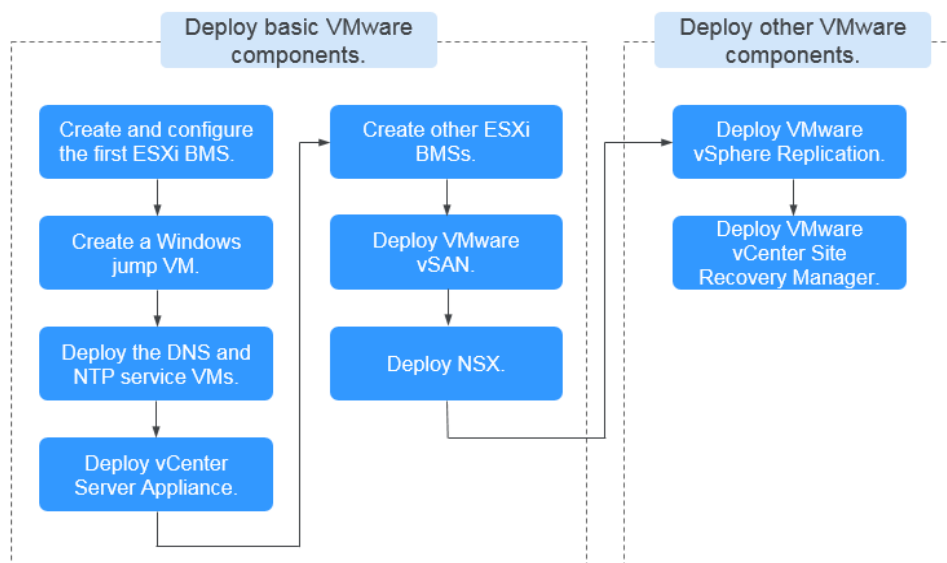
1.2.3 Restrictions

- The MAC address of the primary ESX NIC cannot be changed. If you change it, network connection of the primary ESX NIC will be interrupted.
- To enable the VMware VM to communicate with other VMs or BMSs in the VPC, ensure that the VMware VM IP address is different from the IP CIDR of the VPC.
- The host for provisioning ESXi images cannot have EVS disks.
- When forming a NIC group of the QinQ high-speed network, you cannot use ToR for configuration, such as etherchannel and LACP.

1.2.4 Deployment Process

To deploy VMware in the BMS cluster and configure public cloud services and VMware, perform the following operations:

Figure 1-2 Deployment process of VMware on BMS



Basic VMware components are mandatory, including vCenter Server Appliance, VMware vSAN, and NSX. Deploy them as instructed. Other VMware components are optional, including VMware vSphere Replication and VMware vCenter Site Recovery Manager. You can deploy them as needed.

1.2.5 Creating and Configuring the First ESXi BMS

- Step 1** Log in to the management console.
- Step 2** Create a BMS. Select a flavor and ESXi image (see [Table 1-3](#)) and configure the VPC to which it connects. Set other parameters as required.
- Step 3** Apply for an EIP and bind it to the BMS port connected to the esx-primary network.
- Step 4** Use the EIP and key pair to log in to the ESXi BMS and change the user **root** password of the host.
- Step 5** Enter the EIP of the primary NIC of the ESXi BMS in the browser to log in to the ESXi host.
- Step 6** Add datastores to the ESXi host.

----End

1.2.6 Creating a Windows Jump VM

Create a jump VM that connects to the hb-mgmt network and VPC. Configure VMware components manually or enable automatic configuration.

- Step 1** Create a Windows VM as a jump server and connect the VM to the hb-mgmt network.
 1. On the **vSwitch** page, click **Add Standard vSwitch**. In the displayed dialog box, enter **vSwitch1** for **vSwitch Name** and add uplinks **vmnic2** and **vmnic3**.

2. Click **Add Port Group**. In the displayed dialog box, select **vSwitch1** for **vSwitch**, add port group **hb-mgmt**, and set **VLAN ID** to **0**.
3. Create a VM and configure a network adapter to connect the VM to the **hb-mgmt** port group.
4. Start the VM and install Windows on the VM. After the installation is complete, configure IP address 11.11.11.2 for the VM.
5. Use the remote desktop to log in to the Windows jump VM and enable the remote desktop service (set the gateway address to 11.11.11.1 and disable the firewall of the jump VM).
6. Upload the downloaded software such as the vCenter ISO and NSX-manager ovf templates to the jump VM.

Step 2 On the **VMkernel NIC** page, click **Add VMkernel NIC**. In the displayed dialog box, select **vSwitch1** for **vSwitch**, create **VMkernel** port group and NIC, and set **VLAN ID** to **0** and **Address** to **11.11.11.101**.

The configurations of vSwitch and port groups on the ESXi host are as follows.

Table 1-5 vSwitch and port group configurations

Port Group	vSwitch	Uplink	VMkernel NIC
management network	vSwitch0	vmnic0/vmnic1	vmk0
hb-mgmt	vSwitch1	vmnic2	-
vmk-hb-mgmt		vmnic3	vmk1

----End

1.2.7 Deploying the DNS and NTP Service VMs

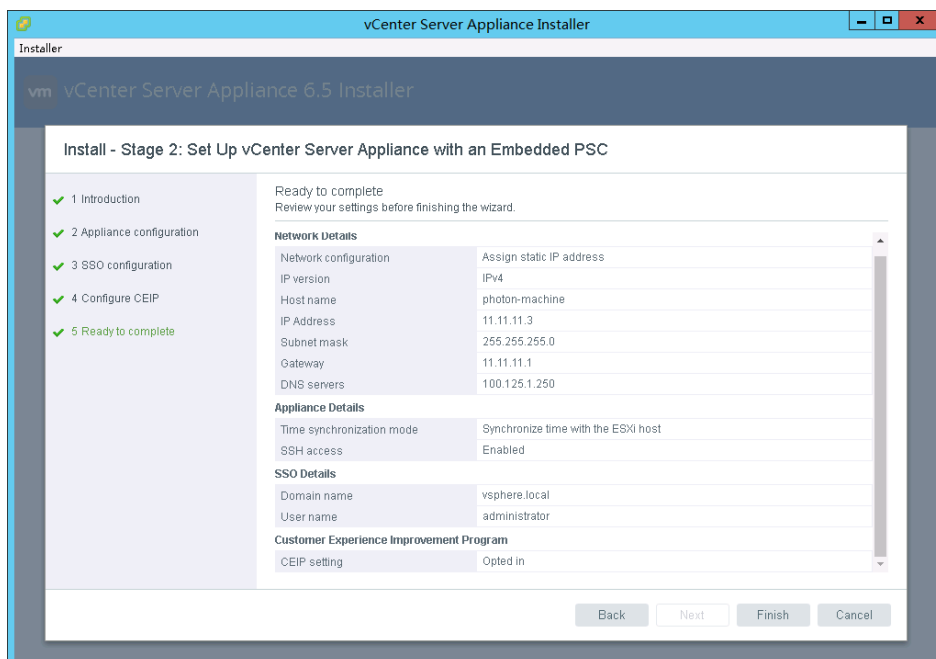
- Step 1** Create a VM and configure a vNIC which connects to the hb-mgmt port group.
- Step 2** Install Linux and the NTP service on the VM.
- Step 3** Log in to the VM from the remote desktop and set the IP address of the vNIC to 11.11.11.6. To configure the NTP service, enable the firewall ports.

----End

1.2.8 Deploying vCenter Server Appliance

- Step 1** Deploy vCenter Server Appliance on the ESXi host. For details, see section "Deploying the vCenter Server Appliance and Platform Services Controller Appliance" in [vSphere Installation and Setup](#).
- Step 2** Some example configurations are as follows:
1. When importing the OVA file during the first installation phase, set a static IP address 11.11.11.3 for vCenter Server Appliance, and set the subnet mask to 255.255.255.0, default gateway to 11.11.11.1, and DNS to 11.11.11.6.

2. In phase 2, set the username and password of the new vCenter Single Sign-on domain. Assume that the username is **administrator**, domain name is **vsphere.local**, and site name is **photon-machine**.



Step 3 Log in to vCenter Web Client and create a data center and a cluster.

Step 4 On the vCenter Web Client, create three distributed virtual switches (DVSs), create one or more distributed port groups on each DVS, and configure VLAN IDs for the distributed port groups based on DPortGroup-mgmt, DPortGroup-vxlan, DPortGroup-edge-internal, and DPortGroup-vsan. [Table 1-6](#) lists the DVSs and distributed port groups.

Table 1-6 Configuration information

DVS	Distributed Port Group	VLAN ID
dvSwitch-0	DPortGroup-mgmt	0
dvSwitch-1	DPortGroup-vxlan	200
dvSwitch-1	DPortGroup-edge-internal	300
dvSwitch-2	DPortGroup-vsan	400

Step 5 Create the VMkernel NIC on the hb-vsan port group. Select **Virtual SAN traffic** for the port attribute and configure the IP address based on [Table 1-3](#).

Migrate the VMkernel NIC and uplink on vSwitch1 to the distributed port group and DVS of dpd-hb-mgmt.

Step 6 Add the first ESXi host to vCenter through IP address 11.11.11.101.

Configure a distributed switch and migrate the vmknic, DNS/NTP, vCenter, and jump VM to the switch. For details, see section [Setting Up Networking with vSphere Distributed Switches](#).

- Step 7** Add the ESXi host to the DVSs dvSwitch-1 and dvSwitch-2. Use uplinks vmnic4/vmnic5 and vmnic6/vmnic7 in active/standby mode.
- Step 8** Add the ESXi host to the distributed switch dvSwitch-0 using the uplink vmnic3.
- Step 9** Migrate the VMkernel NIC vmk1 to the distributed port group dpg-hb-mgmt.
- Step 10** Connect the DNS/NTP on the ESXi host and the vNIC of the vCenter VM to the distributed port group dpg-hb-mgmt.
- Step 11** Migrate the Windows jump VM on the ESXi host and connect the vNIC that connects to vSwitch1/hb-mgmt to the port group dpg-hb-mgmt on dvSwitch0.
- Step 12** Add vmnic2 to dvSwitch-0 and delete the vSwitch1 VM.

Configurations of the vSwitch on the first ESXi host are as follows.

Table 1-7 vSwitch configuration (first ESXi host)

Port Group	vSwitch	Uplink	vlan ID	vmkernel NIC
management network	vSwitch0	vmnic0 vmnic1	-	vmk0
hb-mgmt	dvSwitch-0	vmnic2 vmnic3	0	vmk1
DPortGroup - vxlan	dvSwitch-1	vmnic4 vmnic5	100	-
hb-edge-internal			300	-
DPortGroup - vsan	dvSwitch-2	vmnic6 vmnic7	400	vmk2

----End

1.2.9 Creating Other ESXi BMSs

- Step 1** Create a BMS. Select a flavor and an image (see [Table 1-3](#)) and configure the VPC subnet to which it connects.
- Step 2** Use the key to log in to the ESXi BMS and change the user **root** password.
- Step 3** Add a local datastore to the ESXi host.
- Step 4** Create vSwitch1 for the second and third ESXi hosts, set vmnic2 to uplink, and create vmkernel NICs with IP addresses 11.11.11.102 and 11.11.11.103 on vSwitch1.
- Step 5** Log in to the Windows jump VM and vCenter. Add the ESXi hosts to the vCenter host cluster and DVS through IP addresses 11.11.11.102 and 11.11.11.103.

Step 6 Migrate the VMkernel NIC and uplink on vSwitch1 to the distributed port group and DVS of dvSwitch-0/hb-mgmt. Configurations of the vSwitch on the second and third ESXi hosts are as follows.

Table 1-8 vSwitch configurations of the second and third ESXi hosts

Port Group	vSwitch	Uplink	vmkernel NIC
management network	vSwitch0	vmnic0 vmnic1	vmk0
DPortGroup-mgmt	dvSwitch-0	vmnic2 vmnic3	vmk1
DPortGroup-vxlan	dvSwitch-1	vmnic4	-
DPortGroup-edge-internal		vmnic5	-
DPortGroup-vsan	dvSwitch-2	vmnic6 vmnic7	vmk2

----End

1.2.10 Deploying VMware vSAN

The vSAN cluster requires at least three ESXi hosts.

- Step 1** Log in to the vSphere Web Client. In the navigation pane, choose **Host**. On the **Configuration** page, click **VMkernel Adapter** under **Network** and create a VMkernel.
- Step 2** On the **Select Target Device** page, select **DPortGroup-vsan** for **Select an existing network** and click **Next**.
- Step 3** On the **Port Attributes** page, select **vSAN** and click **Next**.
- Step 4** On the **Set IPv4 Address** page, select **Use a static IPv4 address**. Set the IP address of the three ESXi hosts to **11.11.12.101**, **11.11.12.102**, and **11.11.12.103**, respectively. Then click **Next**.
- Step 5** You can create a vSAN cluster or enable vSAN for an existing cluster. For details, see [Creating a vSAN Cluster](#).

 **NOTE**

- A disk on each ESXi host is reserved as the local datastore. Other SSDs form a disk group and are added to the vSAN cluster.
- VM disks that have been created on the first ESXi host can be migrated to the vSAN cluster.

----End

1.2.11 Deploying NSX

Step 1 Log in to vCenter (Mozilla Firefox is recommended) and deploy NSX Manager.

1. Right-click **Data Center** and select **Deploy OVF Template** from the drop-down list. Select the VMware-NSX-Manager OVA file that has been uploaded to the jump VM.
2. The configurations are as follows:
 - **Download Size: 2.5 GB**
 - **Disk Space: 60.0 GB**
 - **Name: NSX Manager**
 - **Data Store: Datastore**
 - **Destination: 11.11.11.101**
 - **Disk Storage: thick provisioning lazy zeroed**
 - **Network Mapping: Management Network to DPortGroup-mgmt**
 - **IP Address Assignment Mode: static-manual, IPv4**
 - Properties:
 - Hostname=NSX-manager
 - Network 1 IPv4 address=11.11.11.4
 - Network 1 subnet mask=255.255.255.0
 - Default IPv4 gateway: 11.11.11.1
 - DNS server list=11.11.11.6
 - NTP server list=11.11.11.6
 - Enable SSH=False

Step 2 Open a browser and connect to the NSX Manager GUI.

The login address is `https://nsx-manager-ip` or `https://nsx-manager-hostname`.

Use the password configured during installation to log in as user **admin** and click **View Summary**. Ensure that the vPostgres, RabbitMQ, and NSX management services are running.

NSX Manager Virtual Appliance

DNS Name: NSX-manager
 IP Address: 11.11.11.4
 Version: 6.4.3 Build 9927516
 Uptime: 9 days, 18 hours, 34 minutes
 Current Time: [blurred]

Common components

Name	Status	
vPostgres	Running	Stop
RabbitMQ	Running	Stop

NSX Management Components

Name	Status	
NSX Universal Synchronization Service	Stopped	Start
NSX Management Service	Running	Stop

Step 3 Register vCenter Server on NSX Manager.

1. Under **NSX Manager Virtual Appliance Management**, click **Manage vCenter Registration**.

NSX Manager Virtual Appliance Management

2. Edit the vCenter Server element to point to the IP address or host name of the vCenter Server, and enter the username and password of the vCenter Server. You are advised to enter username **administrator@vsphere.local** or your secondary account instead of user **root**.

Step 4 Configure Single Sign On.

1. On the **NSX Management Service** page, click **Edit** of **Lookup Service**. In the displayed dialog box, enter the following information.

Figure 1-3 Lookup Service page

Lookup Service URL ✕

For vCenter versions 6.0 and above, you may configure Lookup Service and provide the SSO administrator credentials to register NSX Management Service as a solution user. It is also recommended to set the NTP server for SSO configuration to work correctly.

Lookup Service Host:

Lookup Service Port:

Enter port 443 for vSphere 6.0 and above.

Lookup Service URL:

SSO Administrator User Name:

Password:

- **Lookup Service Host:** Enter the vCenter Server IP address or host name and its username and password.
 - **Lookup Service Port:** Enter **443**.
2. Verify that the **Status of Lookup Service** is **Connected**.

Figure 1-4 Lookup Service status

Lookup Service URL

For vCenter versions 6.0 and above, you may configure Lookup Service and provide the SSO

Lookup Service URL:	https://11.11.11.3:443/lookupservice/sdk
SSO Administrator User Name:	administrator@vsphere.local
Status:	● Connected


- Step 5** Install and allocate the NSX for vSphere license (this solution relies on vSphere 6.5).
1. Log in to the vSphere Web Client.
 2. Choose **Administration > Licenses > Assets > Solutions** and select **NSX for vSphere** in the **Solution** list. Select **Assign license...** from the **All Actions** drop-down list.
 3. Click **Add**, enter the license key, and click **Next**.
 4. Add the license name, click **Next**, and click **Finish**.
 5. Select a new license and click **View Features** to view the functions that can be enabled by the license.
 6. Click **OK** to allocate the new license to NSX.

- Step 6** Deploy the NSX Controller cluster.
1. Log in to the vSphere Web Client.
 2. Choose **Home > Network and Security > Installation and Upgrade** and click the **Management** tab.

3. Click the **NSX Controller Node** tab and click the + icon under the controller node.
4. Specify the following parameters.
 - **Name: NSX-controller01**
 - **Data Center: DataCenter**
 - **Cluster/Resource Pool: test**
 - **Data Store: vsanDatastore**
 - **Connected To: DPortGroup-mgmt**
 - **IP Address Pool: ippool**

Connect NSX Controller to the vSphere Distributed Switch port group of NSX Manager, other controllers and hosts.
5. If you have not configured an IP address pool for the controller cluster, click **Create IP Address Pool** and create one. Use the IP network segment planned in [Table 1-2](#).
6. Enter the password of the controller. If the password does not comply with the password requirements, a prompt will be displayed.
7. After you finish deploying the first controller, deploy the other two using the same method.

Step 7 Exclude VMs from the firewall protection (the network of vCenter Server will be interrupted in case of a misoperation). NSX Manager, NSX Controller, and NSX Edge will be automatically excluded from the protection of the NSX distributed firewall. You need to add the vCenter, Windows jump VMs, and DNS VMs to the exclusion list.

1. On the vSphere Web Client, click **Network and Security**. In **Security**, click **Firewall Configuration**.
2. Click the **Exclusion list** tab.
3. Click +, select the VMs that you want to exclude, and click .
4. Click **OK**.

Step 8 Prepare a host cluster for NSX. This step is applicable to preparing for hosts for the first time. If NSX nodes have been added to the cluster before the ESX nodes are created, you are advised to reinstall ESX and then perform this step.

1. On the vCenter Web Client, choose **Network and Security > Installation and Upgrade**. Then click the **Host Preparation** tab.
2. Click the cluster that requires the NSX logical switching, routing, and firewall functions, click **Operation**, and select **Install** from the drop-down list.

 **NOTE**

A computing cluster (also referred to as payload cluster) uses application program VMs (such as web and database VMs). If a computing cluster requires the NSX logical switching, routing, and firewall functions, you must perform installation operations for the computing cluster.

In this example, in the shared **Management and Edge** cluster, NSX Manager VMs and controller VMs share a cluster containing Edge devices, including distributed logical routers (DLRs) and Edge service gateways (ESGs). In this case, you must perform installation operations for the shared cluster. On the contrary, if **Management and Edge** has a specified cluster that is not shared (recommended in the production environment), perform installation operations on the Edge and management clusters, respectively.

3. If a green tick is displayed in the **NSX Installation** column, the installation is complete.
4. Verify the installation. Log in to each ESX host and run the **esxcli software vib list | grep esx** command. If **esx-nsxv** is displayed, the installation is successful.

Step 9 Configure the VXLAN transmission parameters.

1. On the vCenter Web Client, choose **Network and Security > Installation and Upgrade > Host Preparation**, and click **Configure** next to **VXLAN**.
2. Configure the logical network.

Specify the parameters shown in the preceding figure. Set **MTU** to **1550** or a greater value for each switch. By default, the value is **1600**.

 **NOTE**

If the MTU value is greater than that of the VXLAN MTU, the value of MTU will not be adjusted. If the value is set to a small one, it will be adjusted to match VXLAN MTU. For example, if the value of MTU is set to 2000 and you accept the default VXLAN MTU value 1600, the value of MTU will not be changed. If the value of MTU is 1500 and that of VXLAN MTU is 1600, the value of MTU will be changed to 1600.

3. Add an IP address pool. Ensure that the selected VLAN does not contain IP address segments used by other resources.
4. When you configure the VXLAN, a new distributed port group will be created. You can view its information on the **Summary** page.

Step 10 Allocate a segment ID pool and the multicast address range.

1. On the vCenter Web Client, choose **Home > Network and Security > Installation and Upgrade** and click the **Configure Logical Network** tab.
2. Click **VXLAN Configuration**, click **Edit** next to **Segment ID**, and set the segment ID range to **5000-5999**. Disable the multicast addressing.

Step 11 Add a transmission area (select the unicast mode).

1. On the vCenter Web Client, choose **Home > Network and Security > Installation and Upgrade** and click the **Configure Logical Network** tab.
2. Click the **Transport Zones** tab and then click the + icon to create a transport zone.
3. In the **Create Transport Zone** dialog box, specify **Name** and **Description** (optional), select **Unicast** for **Replication Mode**, select the clusters to be added, and click **Add**.

- Step 12** Create a logical switch and two VMs. Connect the VMs to the logical switch to verify the connectivity. For details, see [Add a Logical Switch](#).
- Step 13** Create an NSX Edge to enable VMware VMs to communicate with the external network.
1. On the vCenter Web Client, choose **Home > Networking & Security > NSX Edges**, create a new logical router and uplink interface to access the hb-edge-internal port group. Set the IP address to 11.11.8.2/24. Create internal to access logical switch to which the VM belongs. For details, see [Add a Distributed Logical Router](#).
 2. On the **NSX Edge** page, create an NSX Edge device as the Edge service gateway.
 3. Connect the uplink to the DPortgroup-mgmt port group and set the IP address to 11.11.11.30, which is the IP address of the port reserved by edge. Add an internal interface to connect to the edge-internal port group and set the IP address to 11.11.8.1/24.
For details, see [Add an Edge Services Gateway](#).
 4. Configure an OSPF dynamic route between the created logical router and Edge service gateway (the Edge gateway detects the routing topology on the logical router). For details, see [Configure OSPF on a Logical \(Distributed\) Router](#) and [Configure OSPF on an Edge Services Gateway](#).
After the preceding configurations are complete, VMs in the VXLAN network connected to the BMS router can communicate with each other through the logical router. The IP requests for accessing the Internet from VMs will be routed to the uplink of the Edge service gateway and then to the VPC.
 5. (Optional) If a VMware VM needs to access the Internet through EIP, you need to configure NAT rules on the Edge gateway to convert the internal network IP address of the VM into an uplink IP address of the Edge gateway.

----End

1.2.12 Deploying Other VMware Components

Deploying VMware vSphere Replication

- Step 1** Log in to vCenter (Mozilla Firefox is recommended) and deploy vSphere Replication.
1. Mount the downloaded ISO image to the jump VM. Right-click **Data Center** and select **Deploy OVF Template**. Select files **vSphere_Replication_OVF10.ovf**, **vSphere_Replication-system.vmdk**, and **vSphere_Replication-support.vmdk** in the **/bin** directory.
 2. The configurations are as follows:
 - **Download Size: 868.3 MB**
 - **Disk Space: 18.0 GB**
 - **Name: vSphere_Replication_SRM**
 - **Data Center: DataCenter**
 - **Destination: 11.11.11.101**

- **Disk Storage: thick provisioning lazy zeroed**
- **Network Mapping: Management Network to DPortGroup-mgmt**
- **IP Address Assignment Mode: static-manual, IPv4**
- Properties:
 - Domain name server=replication
 - Network 1 IPv4 address=11.11.11.5
 - Network 1 subnet mask=255.255.255.0
 - Default IPv4 gateway: 11.11.11.1
 - NTP server list=172.16.0.102

Step 2 Power on the vSphere Replication device and log in to the NSX Manager GUI through a browser.

Login address: **https://vr-appliance-address:5480**


Log in to the GUI as user **root** and using the password you set during installation.

Step 3 Register vCenter Single Sign-On for vSphere Replication.

On the **Configuration** page, enter an IP address for **LookupService Address**, enter **administrator@vsphere.local** for **SSO Administrator** and its password, save the change, and restart the service.

Step 4 Log in to the vSphere Web Client again. vSphere Replication is displayed on the homepage.

Step 5 Repeat the preceding steps at the target site to install vSphere Replication. Click the **Configuration** tab. In the navigation pane on the left, choose **vSphere**

Replication > Target Site and click . In the displayed dialog box, enter the username and password of the target site and click **Log In**. Select the remote site you want to connect to and click **OK**.

Step 6 If the target site is in the **Connected** state, the source and target sites are connected successfully.

Step 7 To isolate vSphere Replication and other networks, see [Isolating the Network Traffic of vSphere Replication](#).

----End

Deploying VMware vCenter Site Recovery Manager

Step 1 Upload the SRM installation package through the jump VM, double-click the installation program, select the language, and click **OK**.

Step 2 Select the target folder. During the installation, register Site Recovery Manager with vSphere Platform Services Controller. Enter the address, username, and password, and click **Next**.

Step 3 After the installation is complete, log in to vCenter again. **Site Restoration** is displayed on the homepage.

Step 4 Select **Site Restoration**. On the displayed page, pair Site Recovery Manager of the target site, enter the IP address of the vCenter Server that matches SRM extension, and enter the username and password.

Step 5 After connecting to the Site Recovery Manager Server, you need to establish a connection between the SRM and the remote SRM Server. Log in to the vSphere Web Client at one site, choose **Site Restoration** > **Sites**, right-click the remote site, enter the SSO username and password, and click **Log In**.

----End

1.2.13 FAQs

How Do I Enable Layer-3 Communication Between a VMware VM and an ECS Through NSX-Edge?

To enable this, perform the following operations:

1. On the VMware platform, provision logical routes and NSX Edge devices using NSX, configure OSPF, and connect an internal interface of the logical router to the logical switch of the VM.
2. Configure an SNAT rule on NSX Edge, and map the IP address and port of the VMware VM to the IP address and port of Edge uplink.
 - **Applicable To: uplink**
 - **Source IP Address/Range: 11.11.200.2**
 - **Protocol: Any**
 - **Source Port: Any**
 - **Converted Source IP Address/Range: 11.11.11.30**
 - **Converted Port/Range: Any**

What Configurations Are Required for Enabling a VMware VM Connected to NSX Edge to Provide Services Through an EIP?

To enable a VMware VM to provide services through an EIP, perform the following configurations:

1. Bind an EIP to the port reserved for the Edge uplink by adding a DNAT rule.

✕

Add DNAT Rule

i You do not need to add DNAT rules if your ECS has an EIP bound. [Learn more](#)

* Scenario VPC Direct Connect

* Port Type All ports Specific port

* Protocol All

* EIP

* Private IP Address

Assume that the IP address of the port reserved for Edge uplink is 11.11.11.30, and the EIP bound to the port is 188.xxx.xxx.xxx.

2. Configure a DNAT rule on Edge, and map the IP address and port of the VMware VM to the IP address and port of the Edge uplink.
 - **Applicable To: uplink**
 - **Protocol: tcp**
 - **Original Target IP Address/Range: 11.11.11.30**
 - **Source Port/Range: 2201**
 - **Converted Source IP Address/Range: 11.11.200.2**
 - **Converted Port/Range: 22**

In the preceding example, the configured DNAT rule will convert the packets of 11.11.11.30:2201 into those of 11.11.200.2:22.

After the configuration is complete, you can access port 2201 of EIP 188.xxx.xxx.xxx, which is converted into port 22 of a VMware VM whose IP address is 11.11.200.2 through DNAT.

1.3 XenServer on BMS

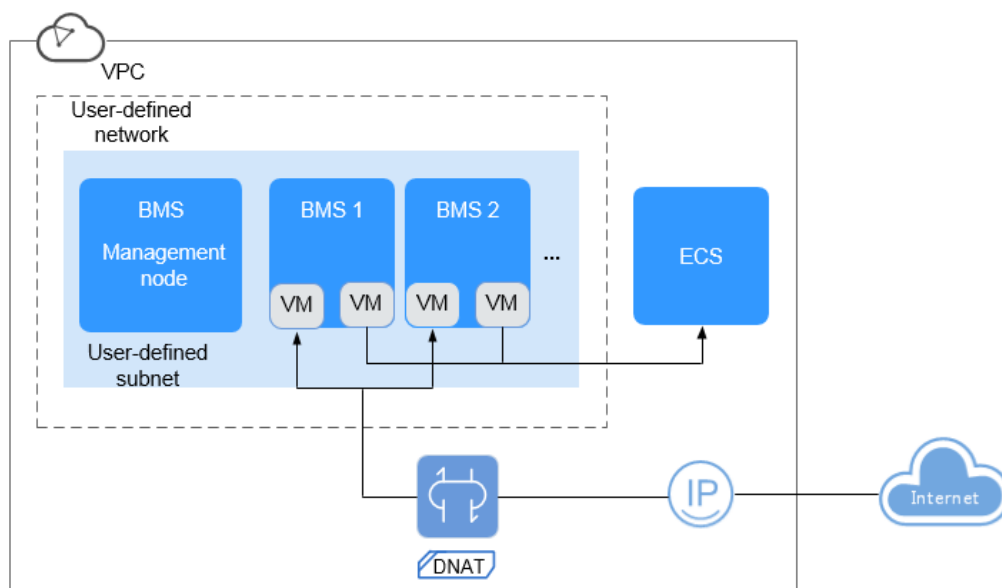
1.3.1 Solution Overview

The BMS service allows you to provision BMSs using the XenServer OS. You can build a XenServer private cloud based on the BMS cluster and seamlessly migrate your core applications to the cloud. Service VMs in the XenServer private cloud can communicate with Elastic Cloud Servers (ECSs) to provide you with a more flexible solution to migrate services to the cloud.

This document provides a reference solution for deploying a XenServer system using BMSs. You can also deploy the XenServer system in a different mode based on the capabilities provided by the public cloud.

Deployment Architecture

Figure 1-5 Deployment architecture



In XenServer on BMS, at least two BMSs are required. One is a multi-NIC Windows BMS with XenCenter installed that functions as the management node of the virtualization cluster. The other is a multi-NIC XenServer BMS that is used to provision service VMs. BMSs of this type can be added based on service requirements (single node and cluster).

If service VMs need to provide services to the Internet, an EIP must be bound to the BMS used to provision service VMs. If a cluster contains multiple BMSs, you can buy a NAT gateway and add a DNAT rule to it enable multiple BMSs to share an EIP. This reduces management cost.

If service VMs only communicate with ECSs in the same VPC, you do not need to buy an EIP.

In [Figure 1-5](#), the user-defined network is a network type provided by HUAWEI CLOUD for the XenServer on BMS solution. This network provides functions similar to those of the VPC for VMs.

1.3.2 Preparing for the Deployment

- Create a key pair.
To ensure system security, you are advised to use the key authentication mode to authorize the user who attempts to log in to a Linux BMS. Therefore, you must use an existing key pair or create a new one for remote login authentication.
Create a key pair by following the instructions in [Creating a Key Pair](#). If you already have a key pair, skip this step.
- Create a VPC.
BMSs use subnets and security groups provided by a VPC.
For how to create a VPC, see [Creating a VPC](#).
- (Optional) Assign an EIP.
If service VMs in the user-defined network need to provide services to the Internet, apply for an EIP. Otherwise, EIP is not required.
For how to apply for an EIP, see [Assigning an EIP](#).
- Register a private image.
Before creating a XenServer BMS, you need to register the XenServer image file as a private image.
For how to register a private image, see [Register a Private Image](#).

1.3.3 Buying a BMS

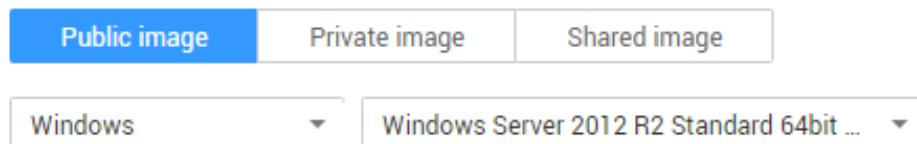
To deploy XenServer on BMS, you must create a multi-NIC Windows BMS and install XenCenter on it. The BMS functions as the management node of the virtualization cluster. Create a multi-NIC XenServer BMS (single-node scenario) or multiple multi-NIC XenServer BMSs (cluster scenario) used to provision service VMs.

Procedure

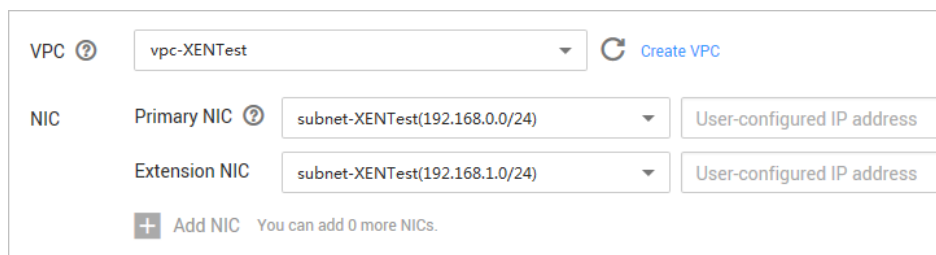
1. Log in to the management console.
2. Choose **Service List > Computing > Bare Metal Server**.
The BMS console is displayed.
3. Click **Buy BMS** in the upper right corner.
4. Configure the BMS specifications.
 - **Flavor:** Select a flavor of BMSs with multiple NICs, for example, **physical.s3.large**. This flavor provides two NICs. A 10GE NIC with two ports connects to the VPC and a 10GE extended NIC with two ports supports high-speed communication between BMSs.

Flavor name	CPU	Memory	Local Disk	Extended Configuration
<input type="radio"/> physical.d1.large	20 core Intel Xe...	128 GB DDR4	2*600G SAS System Disk RAID ...	2 x 2*10GE
<input checked="" type="radio"/> physical.s3.large	20 core Intel Xe...	128 GB DDR4	2*600G SAS System Disk RAID 1	2 x 2*10GE

- **Image:** Select a Windows public image, for example, **Windows Server 2012 R2 Standard 64bit Chinese for BareMetal**.



- **VPC and NIC:** Select the VPC created in section [Preparing for the Deployment](#) and add an extension NIC.



5. Click **Buy Now**.
After about 10 minutes, the BMS is created and its status changes to **Running**.
6. Log in to the BMS using the MSTSC password and install XenCenter.
For the software package and installation method, visit the official Citrix website.
7. Repeat steps **3** to **5** to create a XenServer BMS.
 - **Flavor:** Select a flavor of BMSs with multiple NICs.
 - **Image:** Select the XenServer private image registered in section [Preparing for the Deployment](#).
 - **VPC and NIC:** Select the VPC created in section [Preparing for the Deployment](#) and add an extension NIC.
 - **Login Mode:** If you select **Password**, enter and remember the password of user **root**. If you select the key pair mode, set the password of user **root** after your first login.
 - **Quantity:** Enter **1** for the single node scenario. Set this parameter as required for the cluster scenario.

1.3.4 Creating a User-defined Network

Prerequisites

Multi-NIC BMSs have been created.

Procedure

1. Log in to the management console.
2. Choose **Service List > Computing > Bare Metal Server**.
The BMS console is displayed.
3. On the **User-defined Networks** page, click **Create User-defined Network**.

4. Set **VPC, AZ, and User-defined Subnet**.
 - **Name:** Enter a name as needed, for example, **virtualnetwork-test**.
 - **VPC:** Select the VPC created in section [Preparing for the Deployment](#).
 - **AZ:** Select the AZ you have selected during BMS creation in [Buying a BMS](#).
 - **User-defined Subnet:** The subnet mask is from 16 to 29 bits. The gateway IP address ($x.x.x.1$) of the user-defined subnet cannot be used by tenants. In addition, the CIDR of the user-defined subnet cannot conflict with that of the VPC subnet. An example is **10.10.10.0/24**.

You also need to enter the VLAN value following the user-defined subnet. The value ranges from 1 to 4094. In addition, the VLAN of each user-defined subnet of the same user-defined network must be unique.
5. Click **OK**.

After the user-defined network is created successfully, it is displayed in the list of user-defined networks.
6. Click **virtualnetwork-test** to enter the page showing details of the user-defined network.
7. Click **10.10.10.0/24** to enter the page showing details of the user-defined subnet.

On the **BMSs** tab page, you can view the created multi-NIC BMSs.

IP Address		BMSs	
Name	ID		
bms-xencenter	9cf90234-4a8f-4980-869e-0d6a97f704a7		
bms-xenserver-1	3ade0981-76af-8902-987e-d9e339foa8		

1.3.5 Configuring a VPC NIC for the BMS

After the XenServer BMS is created, you need to bond network ports of the VPC NIC.

NOTE

The operations in this section are required only for the single-node scenario.

Procedure

1. Log in to the XenServer BMS.
2. Run the following command to configure the VPC NIC:
bash /opt/huawei/xenserver-bms-network-config.sh

1.3.6 Configuring the User-defined VLAN NIC for the BMS

After the XenServer BMS is created, you need to bond network ports of the user-defined VLAN NIC.

Single-Node Scenario

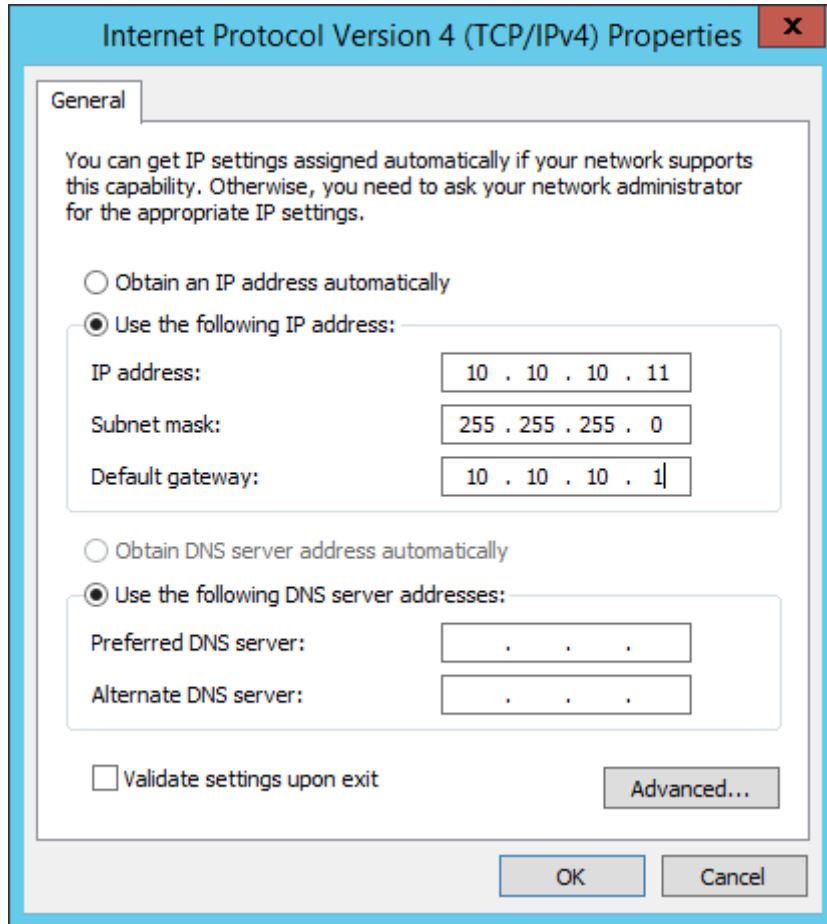
1. Log in to the XenServer BMS.
2. Run the **ip link** command to query the port information. The following figure shows an example command output.

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT
    qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth5: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master ovs-system s
    tate UP mode DEFAULT qlen 1000
    link/ether 04:25:c5:dd:a1:63 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master ovs-system s
    tate UP mode DORMANT qlen 1000
    link/ether a0:08:6f:a0:fb:59 brd ff:ff:ff:ff:ff:ff
4: eth6: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master ovs-system s
    tate UP mode DEFAULT qlen 1000
    link/ether 04:25:c5:dd:a1:64 brd ff:ff:ff:ff:ff:ff
5: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master ovs-system s
    tate UP mode DORMANT qlen 1000
    link/ether a0:08:6f:a0:fb:5a brd ff:ff:ff:ff:ff:ff
6: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master ovs-system s
    tate UP mode DORMANT qlen 1000
    link/ether a0:08:6f:7d:04:63 brd ff:ff:ff:ff:ff:ff
7: eth7: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master ovs-system s
    tate UP mode DEFAULT qlen 1000
    link/ether 04:25:c5:dd:a0:f1 brd ff:ff:ff:ff:ff:ff
8: eth4: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master ovs-system s
    tate UP mode DORMANT qlen 1000
    link/ether a0:08:6f:7d:04:64 brd ff:ff:ff:ff:ff:ff
9: eth8: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master ovs-system s
    tate UP mode DEFAULT qlen 1000
    link/ether 04:25:c5:dd:a0:f2 brd ff:ff:ff:ff:ff:ff
10: ovs-system: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAUL
    T qlen 1
10: ovs-system: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAUL
    T qlen 1
    link/ether 56:ec:f8:c6:6e:44 brd ff:ff:ff:ff:ff:ff
13: xenbr6: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNO
    WN mode DEFAULT qlen 1
    link/ether 04:25:c5:dd:a1:64 brd ff:ff:ff:ff:ff:ff
14: xenbr1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNO
    WN mode DEFAULT qlen 1
    link/ether a0:08:6f:a0:fb:59 brd ff:ff:ff:ff:ff:ff
15: xenbr3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNO
    WN mode DEFAULT qlen 1
    link/ether a0:08:6f:7d:04:63 brd ff:ff:ff:ff:ff:ff
18: xenbr2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNO
    WN mode DEFAULT qlen 1
    link/ether a0:08:6f:a0:fb:5a brd ff:ff:ff:ff:ff:ff
20: xenbr8: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNO
    WN mode DEFAULT qlen 1
    link/ether 04:25:c5:dd:a0:f2 brd ff:ff:ff:ff:ff:ff
21: xenbr7: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNO
    WN mode DEFAULT qlen 1
    link/ether 04:25:c5:dd:a0:f1 brd ff:ff:ff:ff:ff:ff
26: xenbr5: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNO
    WN mode DEFAULT qlen 1
    link/ether 04:25:c5:dd:a1:63 brd ff:ff:ff:ff:ff:ff
27: xenbr4: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNO
    WN mode DEFAULT qlen 1
    link/ether a0:08:6f:7d:04:64 brd ff:ff:ff:ff:ff:ff
```

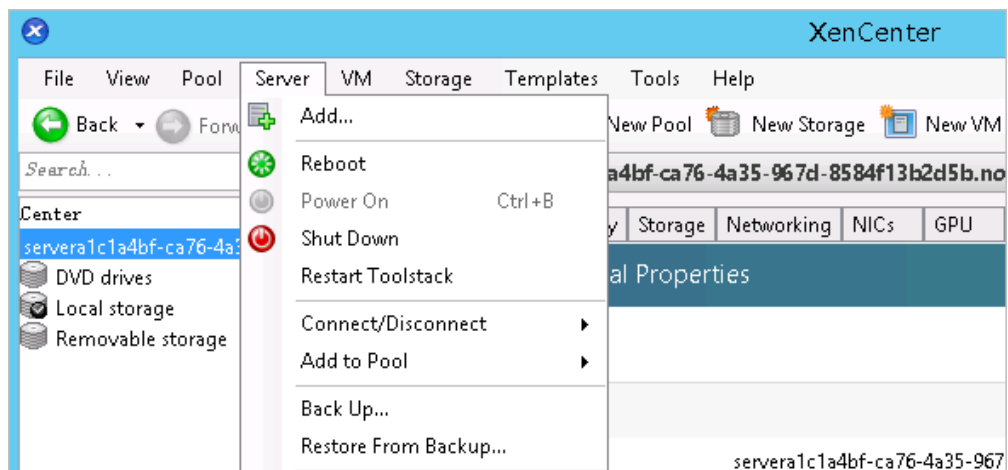
3. Select two network ports to be bonded from those displayed in the command output (for example, xenbr4 and xenbr5) and run the following command to configure a temporary IP address for one of the two ports, for example xenbr4 (the temporary IP address is an IP address in the user-defined subnet segment in section [Creating a User-defined Network](#) (except x.x.x.1)):

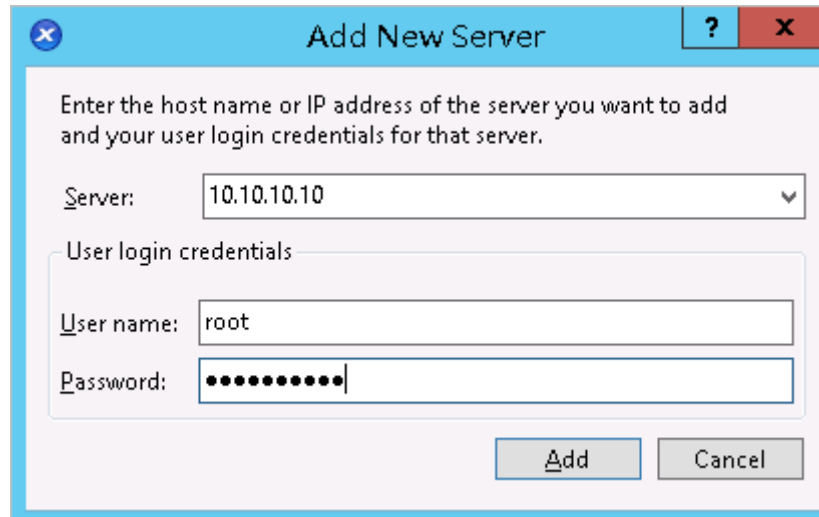
ip addr add 10.10.10.10/24 dev xenbr4

- 4. Log in to the XenCenter BMS, open the control panel, and set an IP address of one of the ports to enable the XenCenter BMS to communicate with the XenServer BMS. (The IP address must be in the same network segment as the IP address configured in step 3, for example 10.10.10.11.)



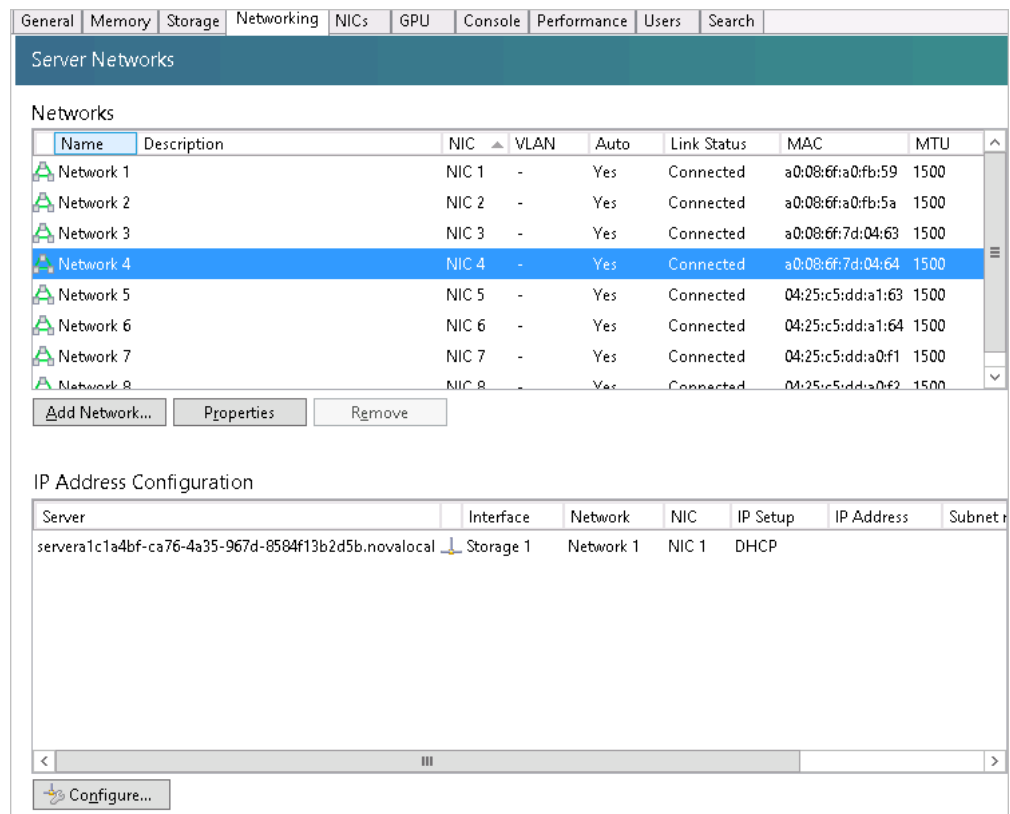
- 5. Start XenCenter, choose **Server > Add** and add the XenServer information.

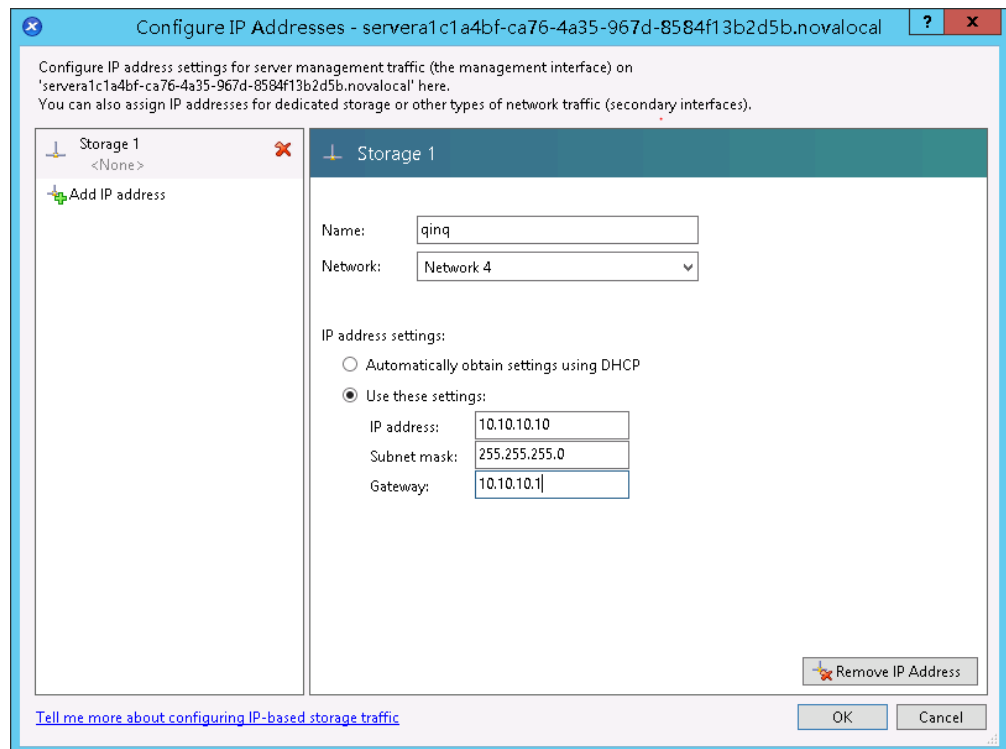




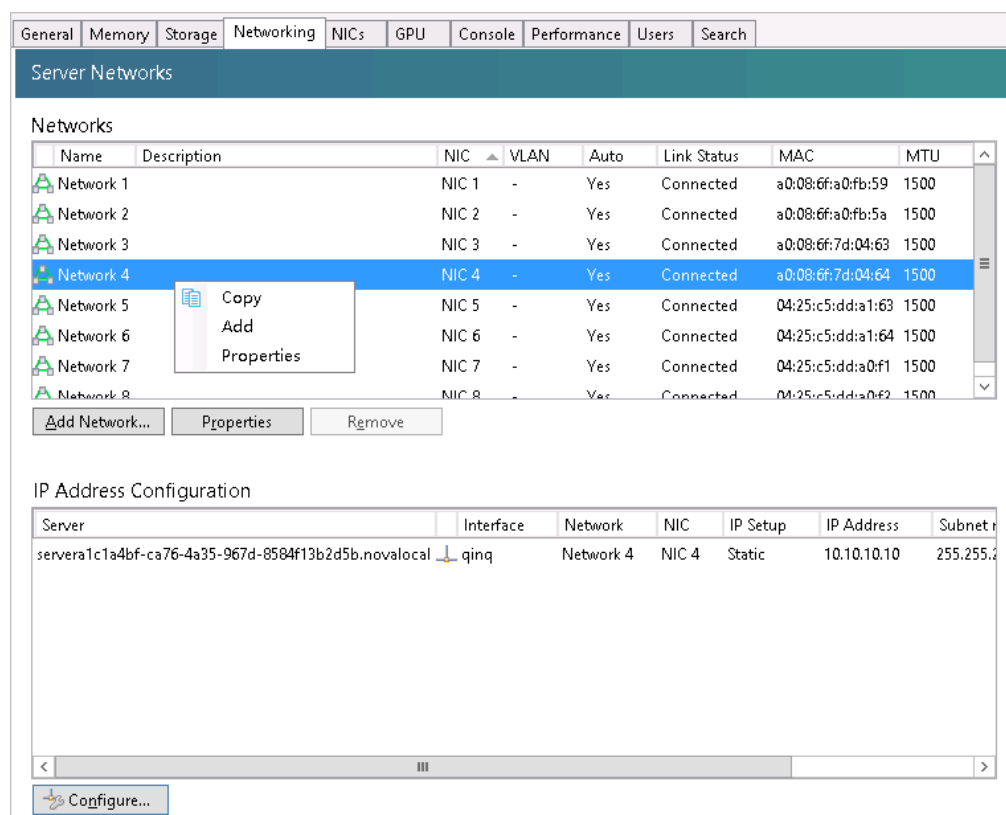
User name and **password** are user **root** of the XenServer OS and its password, respectively.

- After the XenCenter BMS is connected to XenServer, click the **Networking** tab, select network port NIC 4 (**xenbr4**), and click **Configure** in the lower left corner to configure a fixed IP address for the network port.





7. Bond network ports xenbr4 and xenbr5.
Right-click **NIC 4** and choose **Add**.



Select **Bonded Network** and click **Next**.

Select the type of new network you would like to create:

- External Network**
Create a network that passes traffic over one of your VLANs.
- Single-Server Private Network**
Create a network that does not leave each XenServer host.
This can be used as a private connection between VMs on the same host.
- Bonded Network**
Create a network that bonds together two or more of your NICs.
This will create a single higher performing channel.
- Cross-Server Private Network**
Create a network that does not leave the XenServer pool.
This can be used as a private connection between VMs in the pool.
This type of network requires the vSwitch Controller to be running.

i Cross-server private networks require the vSwitch Controller to be configured and running.

Select **NIC 4** and **NIC 5** and click **Finish**.

NIC	MAC	Link Status	Speed	Duplex	Vendor	Device
<input checked="" type="checkbox"/> NIC 4	a0:08:6f:7d:04:64	Connected	10000 Mbit/s	Full	Intel Corporation	82599ES
<input checked="" type="checkbox"/> NIC 5	04:25:e5:dd:a1:63	Connected	10000 Mbit/s	Full	Intel Corporation	Ethernet

Bond mode

- Active-active
- Active-passive
- LACP with load balancing based on IP and port of source and destination
- LACP with load balancing based on source MAC address

MTU: i Allowed MTU value: 1500

Automatically add this network to new virtual machines

8. After the network ports are bonded, a record of the bonded network ports numbered **xapi** is displayed on the XenServer.

```
xapi6: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
  inet 10.10.10.10 netmask 255.255.255.0 broadcast 10.10.10.255
  ether a0:08:6f:7d:04:64 txqueuelen 1 (Ethernet)
  RX packets 55 bytes 13628 (13.3 KiB)
  RX errors 0 dropped 0 overruns 0 frame 0
  TX packets 71 bytes 58081 (56.7 KiB)
  TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

Cluster Scenario

1. Log in to the XenServer BMS.
2. Run the following command to obtain the MAC addresses of the two physical network ports used by the VPC network:

```
bash /opt/huawei/GetVpcNicMac.sh
```

Information similar to the following is displayed:

```
a0:08:6f:a0:fb:59  
a0:08:6f:a0:fb:5a
```

3. Run the **ip link** command to query the port information.
4. In the displayed network ports, select two virtual network ports (for example xenbr4 and xenbr5) except for the network ports (xenbr1 and xenbr2) used by the VPC. Configure the user-defined VLAN NIC by following the instructions in steps 3 to 8 in [Single-Node Scenario](#).
5. Log in to other XenServer BMSs in the cluster one by one and perform the preceding operations to configure the user-defined VLAN NIC.

1.3.7 Configuring the vNIC for the Service VM

To enable service VMs created on the XenServer BMS to communicate with the ECSs in the same VPC or provide services to the Internet, configure vNICs for the service VMs.

NOTE

For how to create a VM on the XenServer BMS, see the official Citrix website or other channels. The vNIC selected during VM creation must be associated with the xapi NIC created in 8 of section [Configuring the User-defined VLAN NIC for the BMS](#).

Procedure

1. Log in to a VM.
2. Run the following command to enable the vNIC on the VM:

```
ip link set eth0 up
```
3. Run the following command to set the IP address of the vNIC to an IP address in the user-defined subnet segment in section [Creating a User-defined Network](#) (except *x.x.x.1*):

```
ip address add 10.10.10.2/24 dev eth0
```
4. Run the following command to configure a default route for the VM:

```
route add default gw 10.10.10.1
```

NOTE

The NIC name (**eth0**), IP address (**10.10.10.2/24**), and gateway address (**10.10.10.1**) in the preceding command are examples. Configure them based on the actual planning.

After the preceding operations are complete, service VMs can communicate with ECSs in the same VPC. You can run the **ping** command to verify the communication.

1.3.8 (Optional) Buying a NAT Gateway and Adding a DNAT Rule

Skip this section if service VMs in the user-defined network do not need to communicate with the Internet.

If service VMs in the user-defined network need to provide services to the Internet, buy a NAT gateway and add a DNAT rule so that BMSs can share an EIP to reduce management cost.

Procedure

1. Log in to the management console.
2. Choose **Service List > Network > NAT Gateway**.
The NAT Gateway console is displayed.
3. Click **Buy NAT Gateway** in the upper right corner.
4. Configure required parameters.

The screenshot shows the NAT Gateway configuration interface. It includes the following fields and options:

- Billing Mode:** Pay-per-use (selected)
- Region:** A dropdown menu with a warning note: "Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region."
- Name:** nat-5c15
- VPC:** vpc-XENTest (with a "View VPCs" link) and a note: "Only VPCs without NAT gateways and default routes are available for selection."
- Subnet:** subnet-XENTest (192.168...) (with a refresh icon) and a note: "Ensure that the subnet has at least one available IP address. One IP address in the subnet will be automatically assigned to the NAT gateway. This subnet is used in the background by the system and has no relationship with the subnet selected during SNAT and DNAT rule configuration."
- Type:** Small (selected), Medium, Large, Extra-large. A note below says: "Supports up to 1,000 new connections per second and up to 10,000 connections in total. [Learn more](#)"
- Description:** A text input field with a character count of 0/255.

- **Region:** Select the region to which the BMS created in section [Buying a BMS](#) belongs.
- **Name:** Enter a name as needed, for example, **nat-5c15**.
- **VPC:** Select the VPC created in section [Preparing for the Deployment](#).

NOTE

The BMS, user-defined network, and NAT gateway must be in the same VPC.

- **Subnet:** Select a subnet as needed.
 - **Type:** Select a type as needed.
 - **Description:** Enter description as needed.
5. Click **Buy Now**. Confirm specifications and click **Submit**.

After the NAT gateway is created successfully, it is displayed in the list of NAT gateways.

6. Click the NAT gateway name **nat-5c15** to enter the page showing details of the NAT gateway.
7. Click the **DNAT Rules** tab and then **Add DNAT Rule**, and specify the following information.
 - **Scenario**: Select **VPC**.
 - **Port Type**: Select **All ports**.
 - **Protocol**: Select **All**.
 - **EIP**: Select the EIP you have obtained in section [Preparing for the Deployment](#).
 - **Private IP Address**: Enter the vNIC IP address of the service VM configured in section [Configuring the vNIC for the Service VM](#).

 **NOTE**

You must add a DNAT rule for the virtual NIC IP address of each VM.

8. Click **OK**.

After the preceding operations are complete, service VMs can communicate with the Internet and the XenServer on BMS solution is deployed successfully.

1.4 Hyper-V on BMS

1.4.1 Solution Overview

Hyper-V is a virtualization product developed by Microsoft, and the company's first hypervisor technology similar to VMware and Xen. It enables users to deploy and use VMs on Windows OSs. Hyper-V is designed to provide cost-effective virtualization infrastructure software more familiar to a wide range of users. This reduces operating costs, improves hardware utilization, optimizes infrastructure, and improves server availability.

Bare Metal Servers (BMSs) have all the features and advantages of physical servers and support secondary virtualization. By provisioning Windows BMSs and deploying the Hyper-V role on them, you can build a private cloud. VMs in the private cloud can communicate with each other, the Internet, and Elastic Cloud Servers (ECSs). This document uses Windows Server 2012 Standard as an example to describe how to deploy Hyper-V on BMSs.

1.4.2 Preparing for the Deployment

- Create a key pair.

When creating a Windows BMS, you can only select the key pair login mode. Therefore, you need to create a key pair. When logging in to the Windows BMS, you need to use this key pair to obtain the password.

Create a key pair by following the instructions in [Creating a Key Pair](#). If you already have a key pair, skip this step.
- Create a VPC.

BMSs use subnets and security groups provided by a VPC.

For how to create a VPC, see [Creating a VPC](#).

- (Optional) Assign an EIP.
If service VMs in the user-defined network need to provide services to the Internet, apply for an EIP. Otherwise, EIP is not required.
For how to apply for an EIP, see [Assigning an EIP](#).

1.4.3 Purchasing a BMS

1. Log in to the management console.
2. Choose **Service List > Computing > Bare Metal Server**.
The BMS console is displayed.
3. In the upper right corner, click **Buy BMS**.
4. Configure the BMS parameters.
 - **Flavor:** Select a flavor of BMSs with multiple NICs. This flavor provides two NICs. A 10GE NIC with two ports connects to the VPC and a 10GE extended NIC with two ports supports high-speed communication between BMSs.

Flavor name	CPU	Memory	Local Disk	Extended Configuration
Sold physical.m2.medium	96 core 4*24Co...	32*b4 GB DIMM	2*600GB SAS System Disk RA...	2x2*10GE
<input type="radio"/> physical.s3.large	20 core Intel Xe...	128 GB DDR4	2*600G SAS System Disk RAI...	2 x 2*10GE
<input checked="" type="radio"/> physical.s4.3xlarge	44 core Intel Xe...	384 GB DDR4	NA	2 x 2*10GE
Sold physical.s4.large	20 core Intel Xe...	192 GB DDR4	NA	2 x 2*10GE
<input type="radio"/> physical.s4.medium	20 core Intel Xe...	128 GB DDR4	NA	2 x 2*10GE
<input type="radio"/> physical.s4.xlarge	28 core Intel Xe...	192 GB DDR4	NA	2 x 2*10GE

- **Image:** Select public image **Windows Server 2012 R2 Standard 64 bit for BareMetal**.

Public image
Private image
Shared image

Windows ▼

Windows Server 2012 R2 Standard 64bit ... ▼

- **VPC and NIC:** Select the VPC created in section [Preparing for the Deployment](#) and add an extension NIC.

VPC ? ↕ Create VPC

NIC ?

Primary NIC

Extension NIC

+ Add NIC You can add 0 more NICs.

5. Click **Buy Now**.
After about 10 minutes, the BMS is created and its status changes to **Running**.

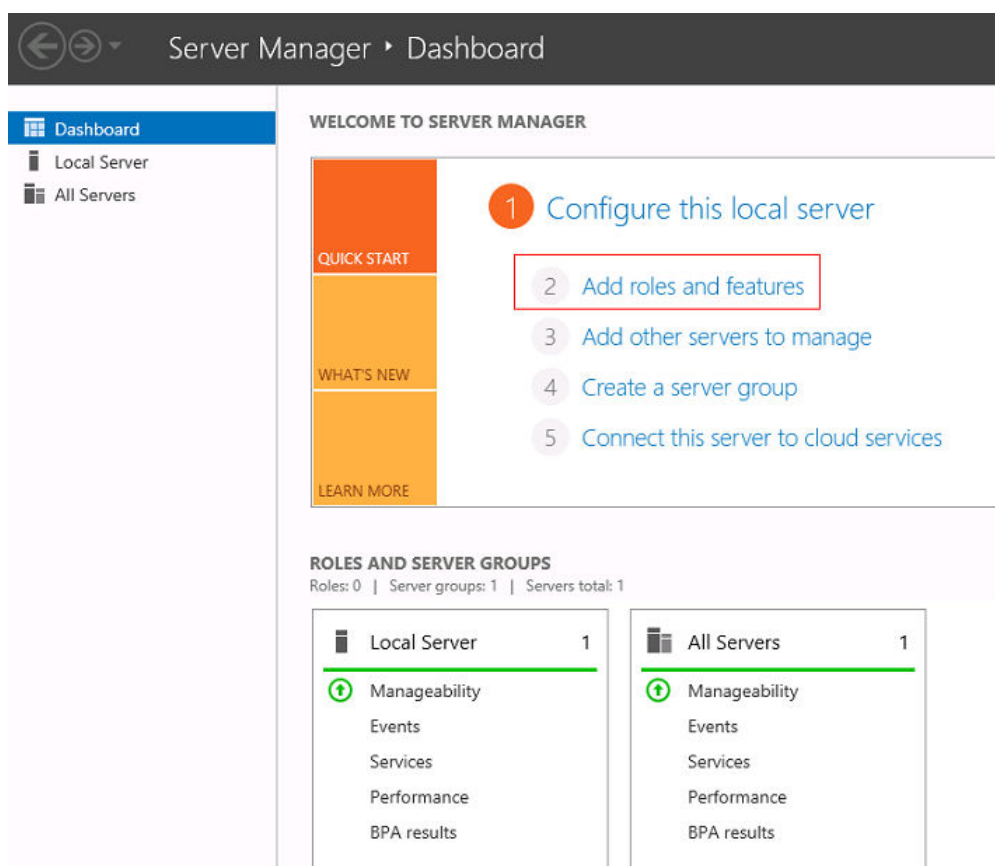
1.4.4 (Optional) Creating a User-defined Network and NAT Gateway and Adding a DNAT Rule

Create a user-defined network by following the instructions in [Creating a User-defined Network](#). To enable the BMS to communicate with the Internet, create a NAT gateway and add a DNAT rule by following the instructions in [\(Optional\) Buying a NAT Gateway and Adding a DNAT Rule](#).

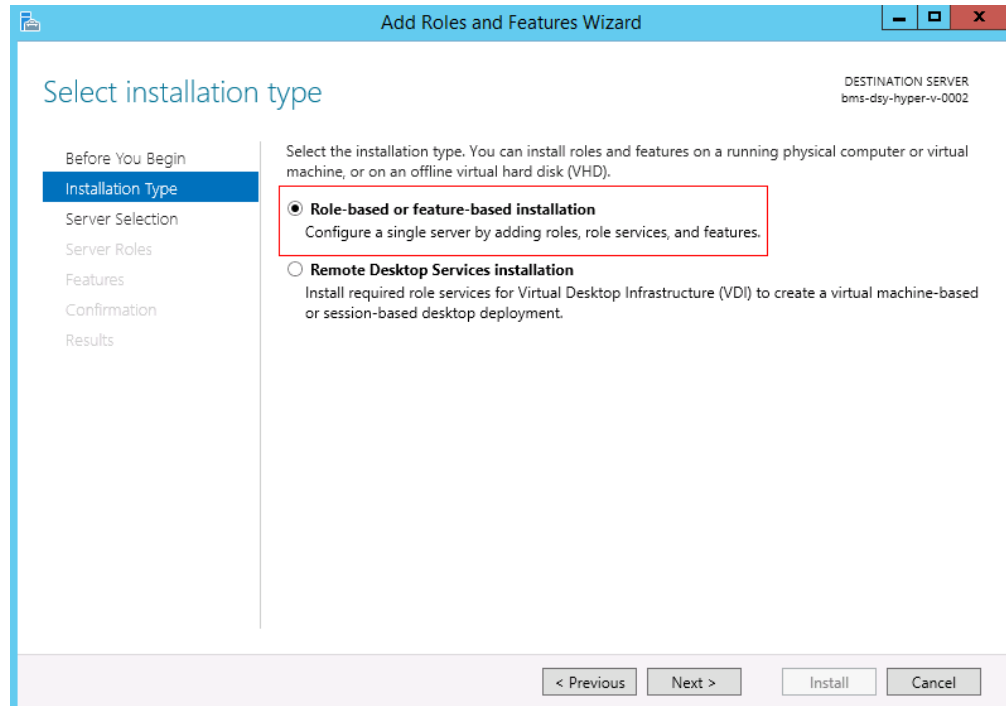
1.4.5 Deploying the Hyper-V Role

Windows Server 2012 R2 has integrated Hyper-V. You can add the Hyper-V role and install Hyper-V Manager to create, run, manage, and schedule VMs.

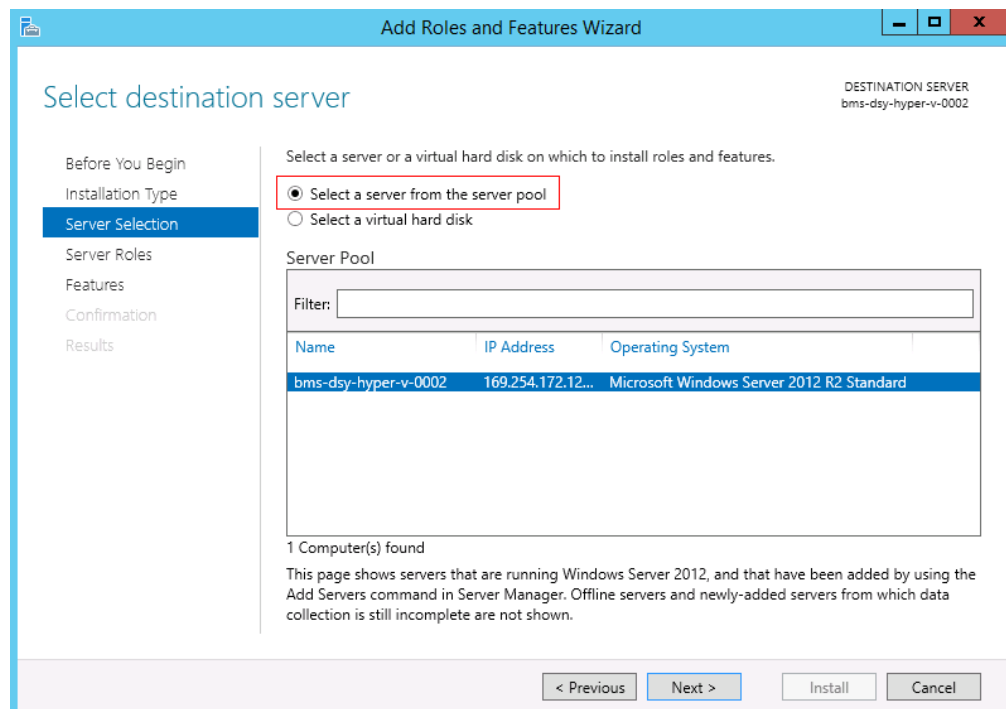
1. Log in to the Windows Server 2012 R2 BMS as the administrator and click **Add roles and features**.



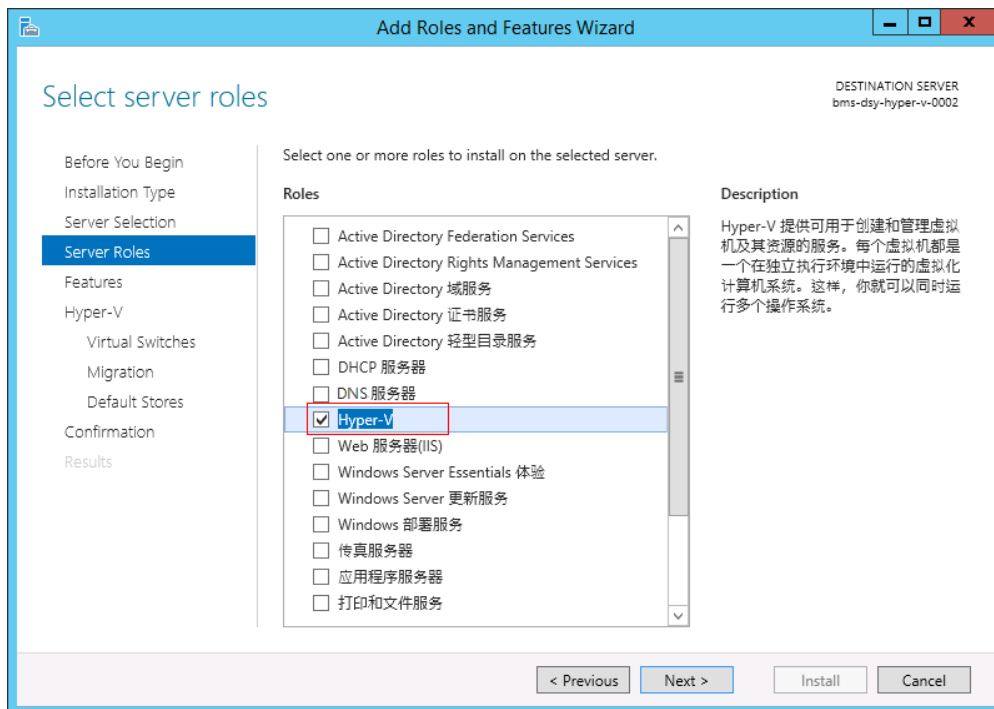
2. In the displayed **Select installation type** dialog box, select **Role-based or feature-based installation** (Hyper-V can only be deployed as a service).



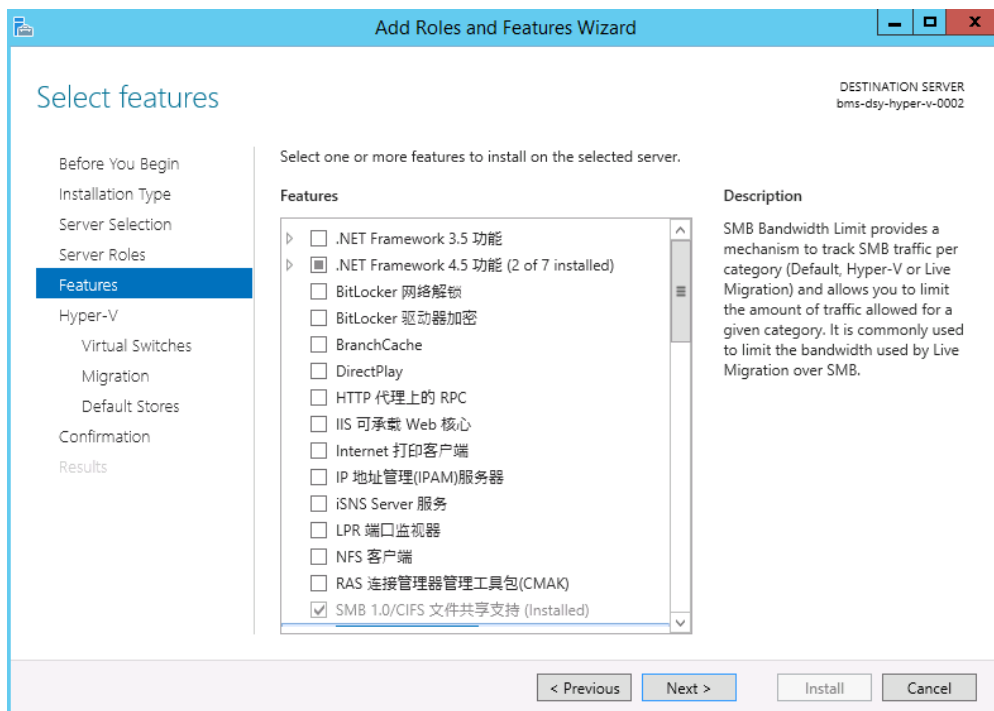
3. In the **Select destination server** dialog box, select **Select a server from the server pool**.



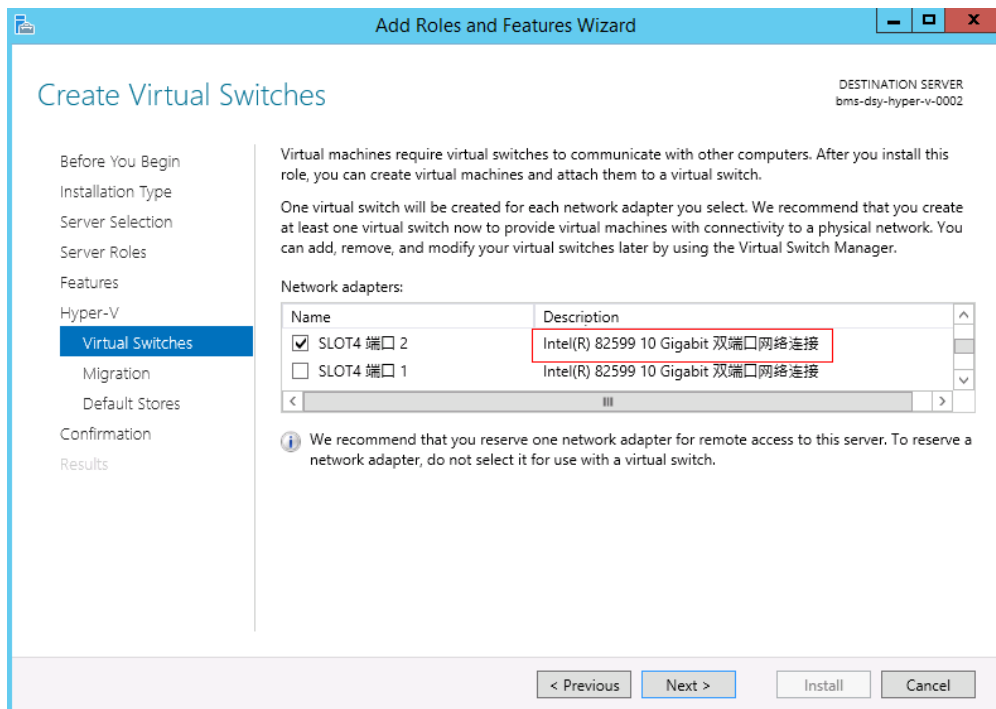
4. In the **Select server roles** dialog box, all available roles are displayed. Select **Hyper-V**. The **Add Roles and Features Wizard** dialog box is displayed. Select **Include management tools (if applicable)** and the **Hyper-V** role.



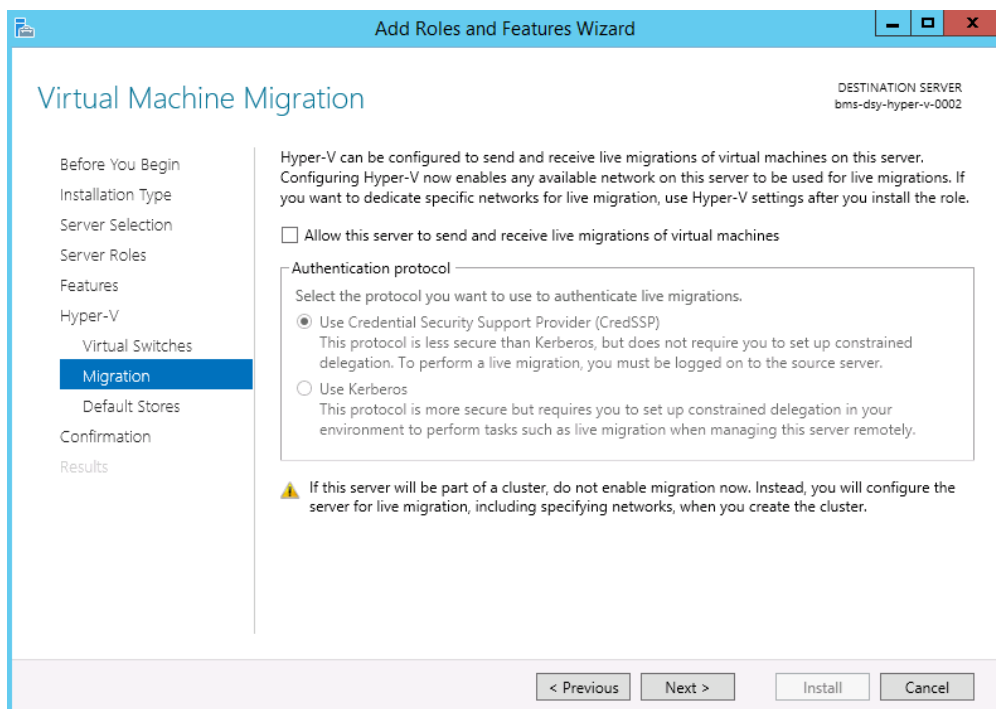
5. In the **Select features** dialog box, choose **Features** in the navigation pane and select the functions you need.



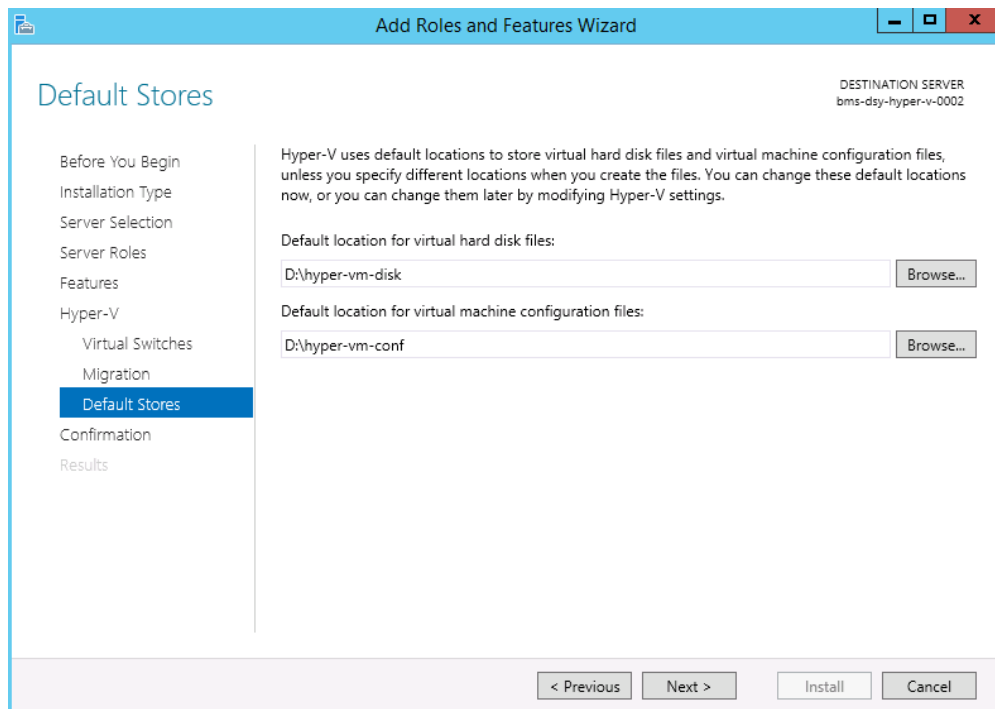
6. In the **Create Virtual Switches** dialog box, select the NIC used for creating the vSwitch. Generally, the 82599 10 Gbit/s extension network port is used for the NIC of the VM to communicate with external resources.



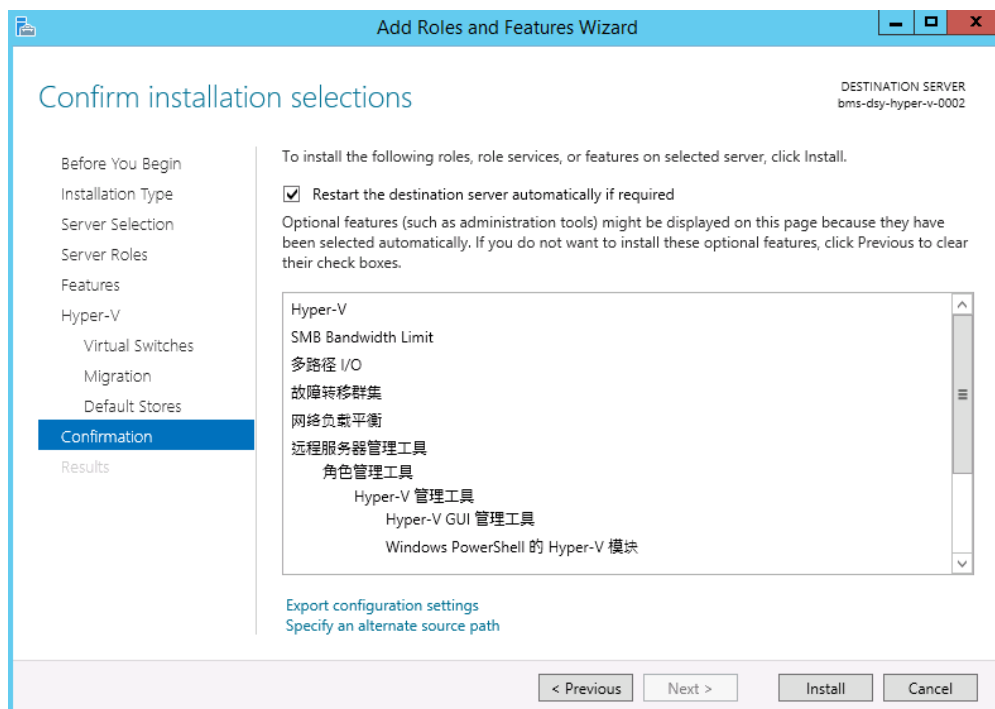
- In the **Virtual Machine Migration** dialog box, the VM migration function is disabled by default. To enable this function, select **Allow this server to send and receive live migrations of virtual machines** and set the authentication protocol. If VMs are deployed in a cluster, do not select this option.



- Set the storage path for the VM and the path for storing configuration files.



9. Click **Install**. After the installation is complete, restart the BMS for the Hyper-V role to take effect.

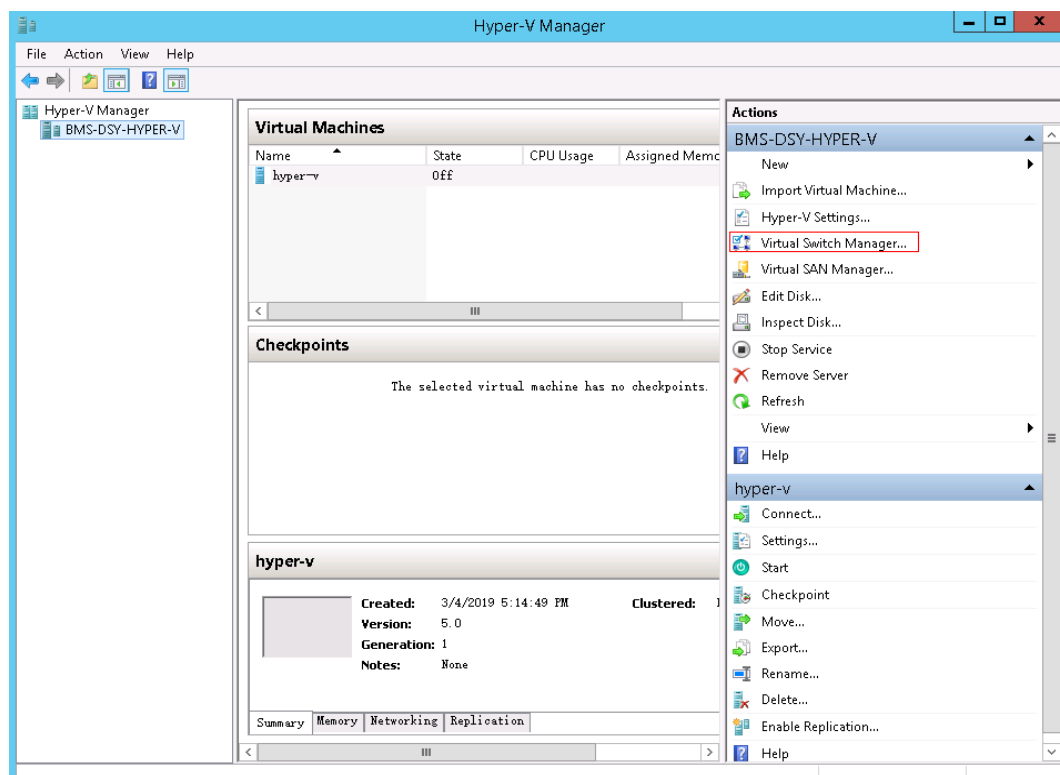


1.4.6 Creating a Hyper-V vSwitch

If you do not configure vSwitches during the Hyper-V deployment in [Deploying the Hyper-V Role](#), you can create a Hyper-V vSwitch by following the instructions in this section.

In the **Hyper-V Manager** window, click **Virtual Switch Manager**.

Figure 1-6 Hyper-V Manager

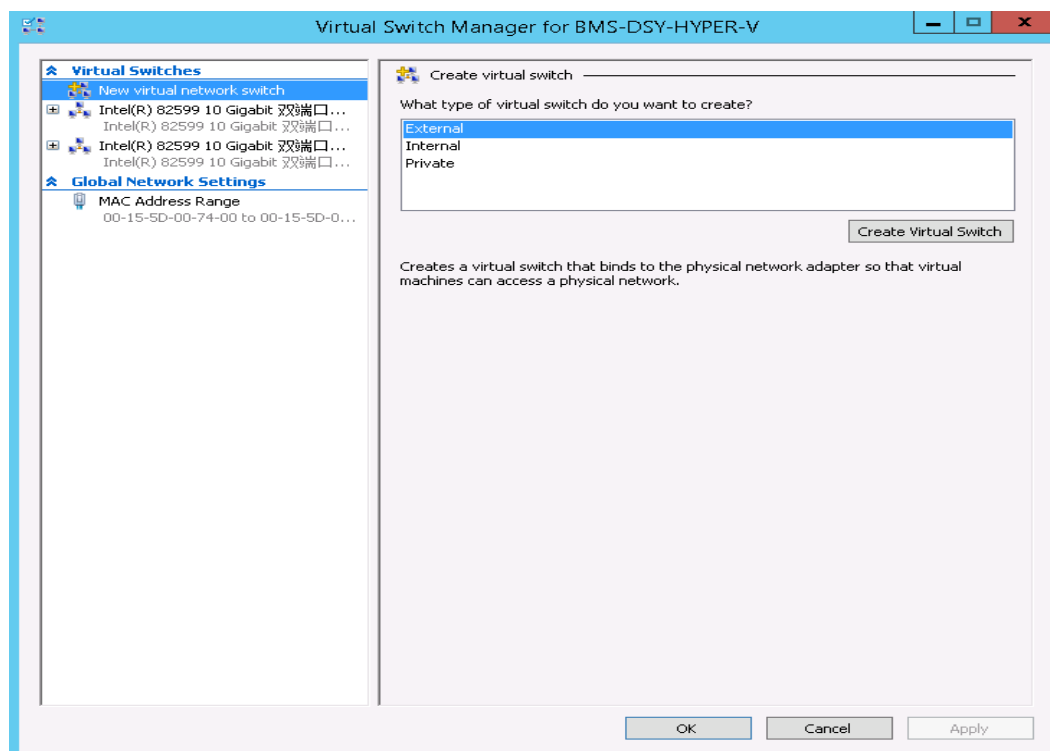


Hyper-V supports three types of vSwitches: external, internal, and dedicated.

- External: After the vSwitch is deployed, VMs and the physical server are connected to the same vSwitch. If you want to enable VMs to communicate with other servers in the LAN, choose this type.
- Internal: VMs on the physical server can communicate with each other and communicate with the physical server, but cannot communicate with other physical servers.
- Dedicated: VMs on the physical server can communicate with each other, but cannot communicate with any physical server.

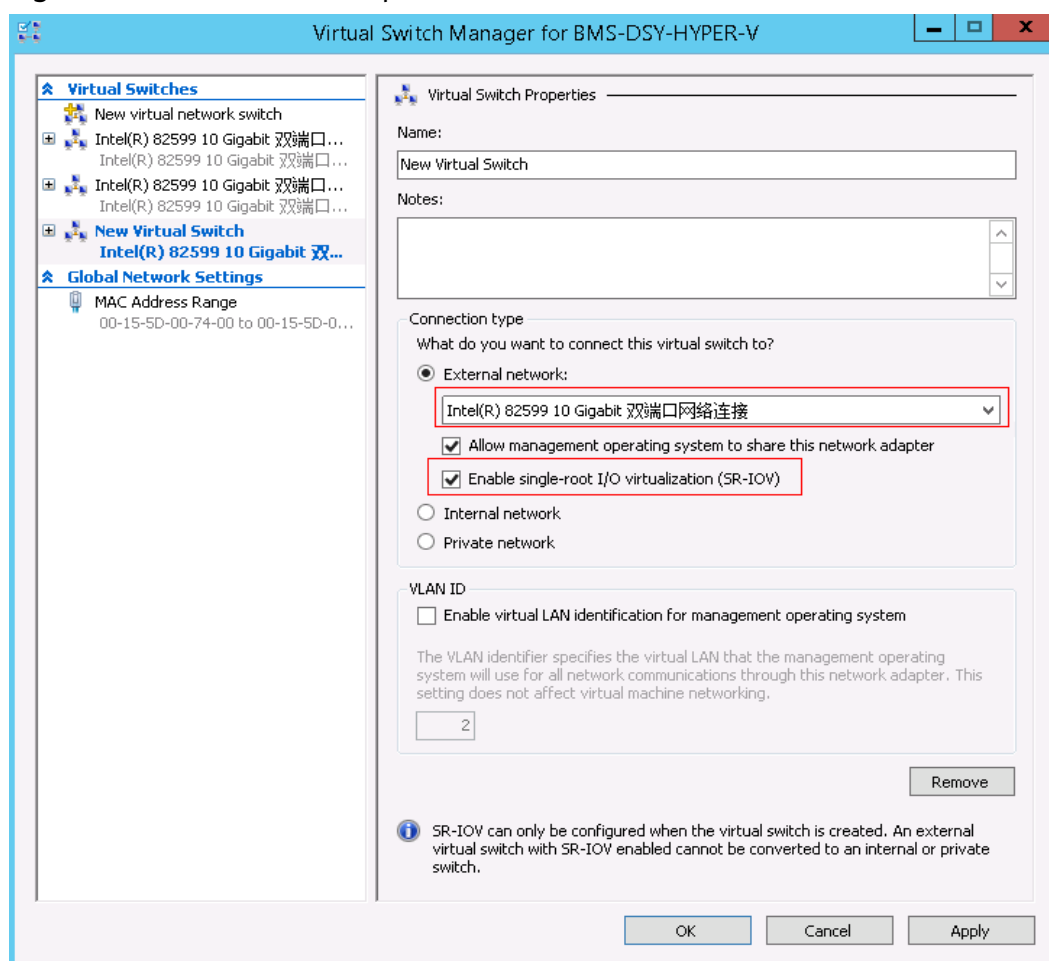
In this section, an external vSwitch is selected. Note that only one external vSwitch can be created for each NIC.

Figure 1-7 Creating a vSwitch



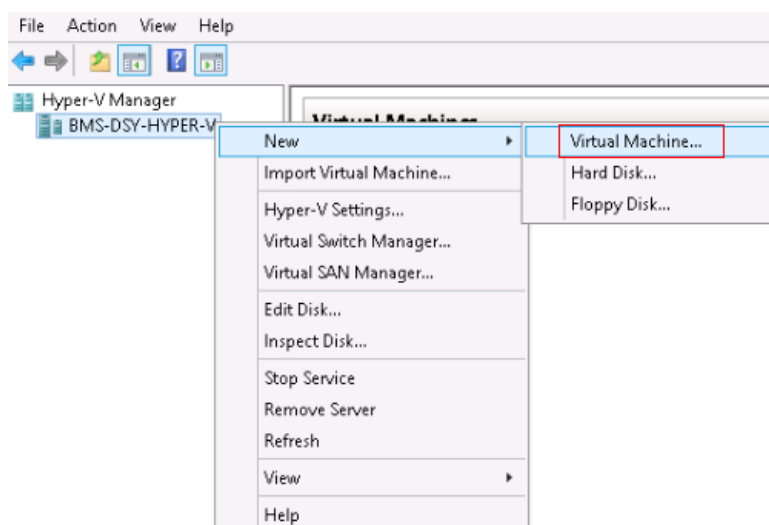
In the **Virtual Switch Properties** window, enter the vSwitch name and select the 82599 10GE extension NIC. When creating a vSwitch, you can select **Enable single-root I/O virtualization (SR-IOV)**.

Figure 1-8 Virtual Switch Properties

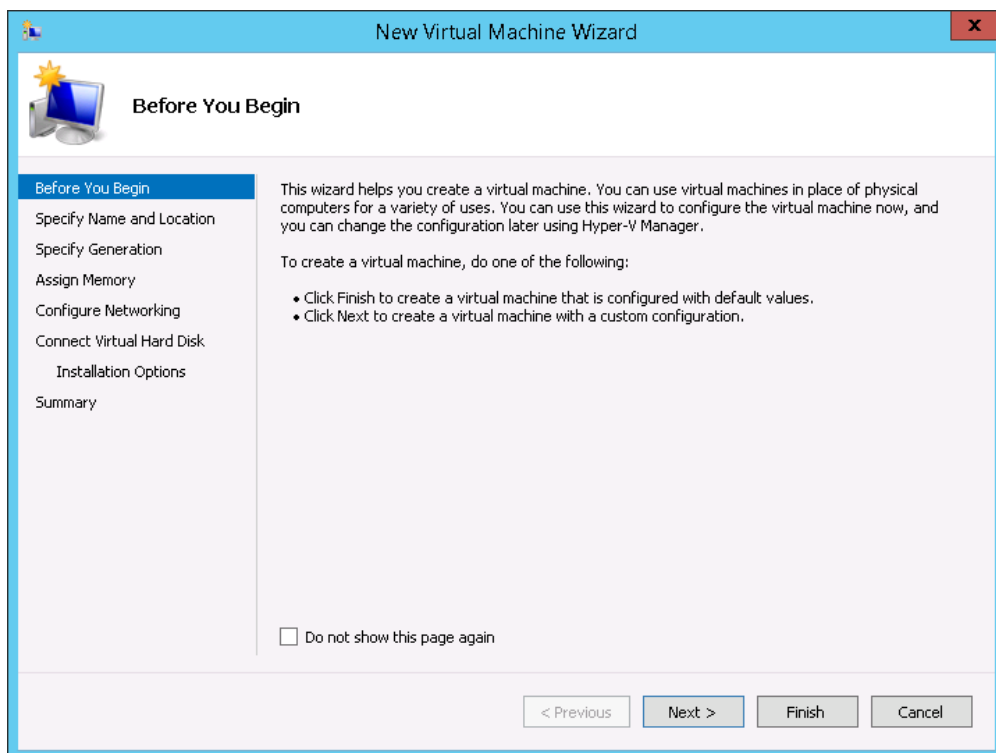


1.4.7 Creating a Hyper-V VM

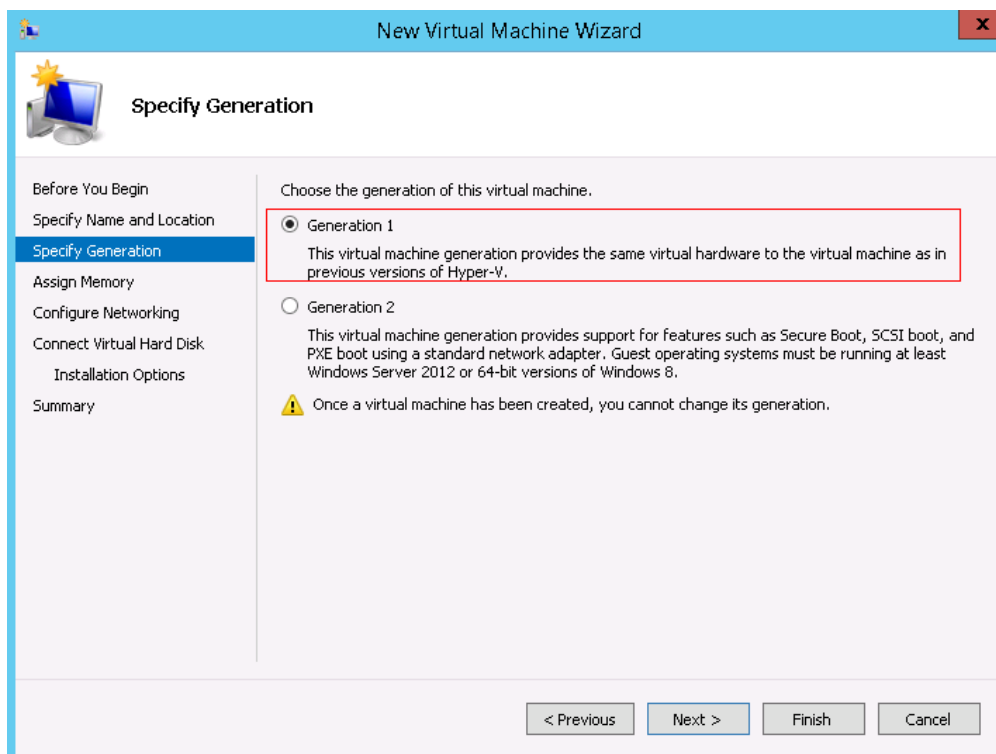
1. On the server management page, click **Hyper-V Manager**. Right-click the target Hyper-V host and choose **New > Virtual Machine**.



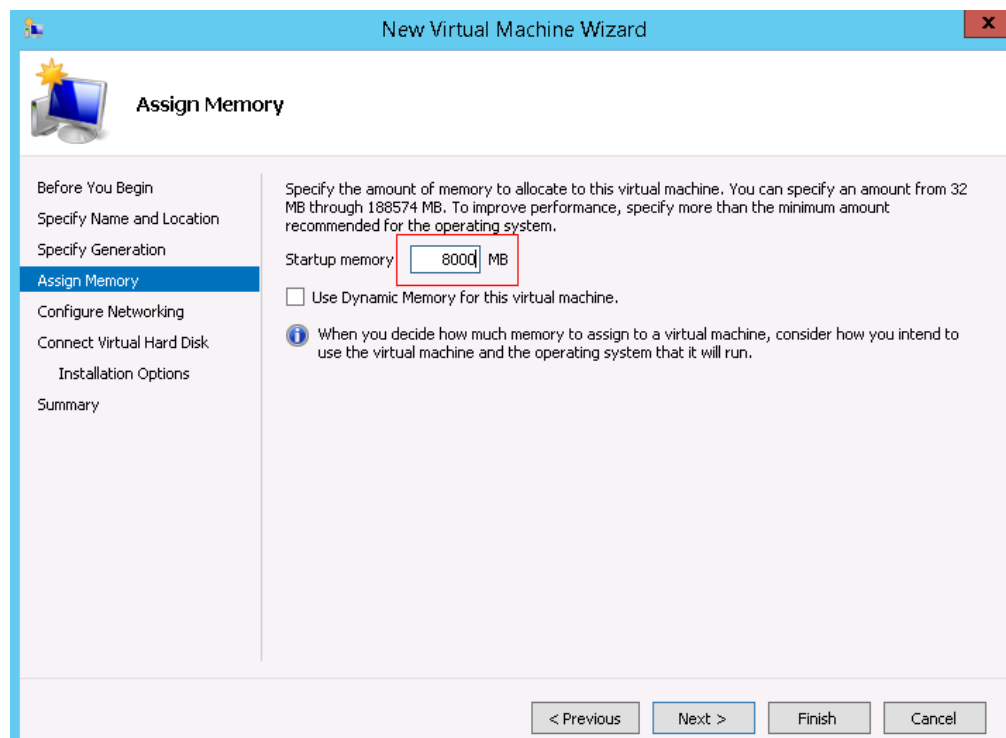
2. In the displayed **New Virtual Machine Wizard** dialog box, you can click **Finish** to create a VM with default configurations or click **Next** to create a VM with custom configurations. In most cases, the second method is used.



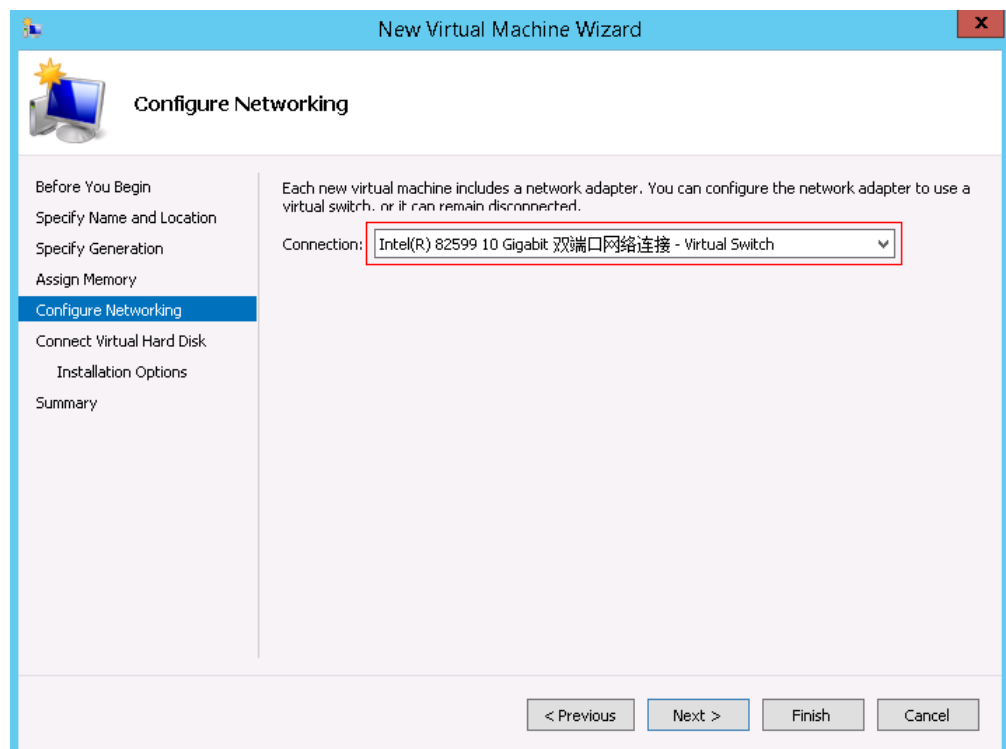
3. In the **Specify Generation** dialog box, select **Generation 1** (default). The first-generation VM inherits the virtual hardware of Windows Server 2012 and has good compatibility. It is also the most widely used VM.



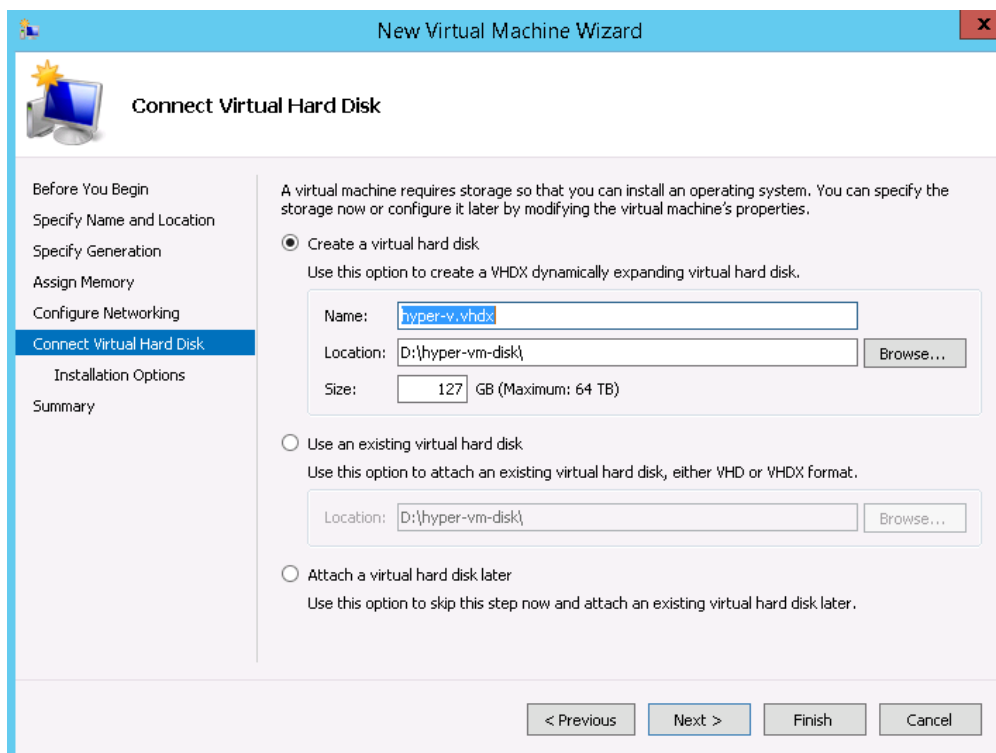
4. In the **Assign Memory** dialog box, select **Use dynamic memory for this VM**. If you do not select this, the value of this parameter is the static memory of the VM.



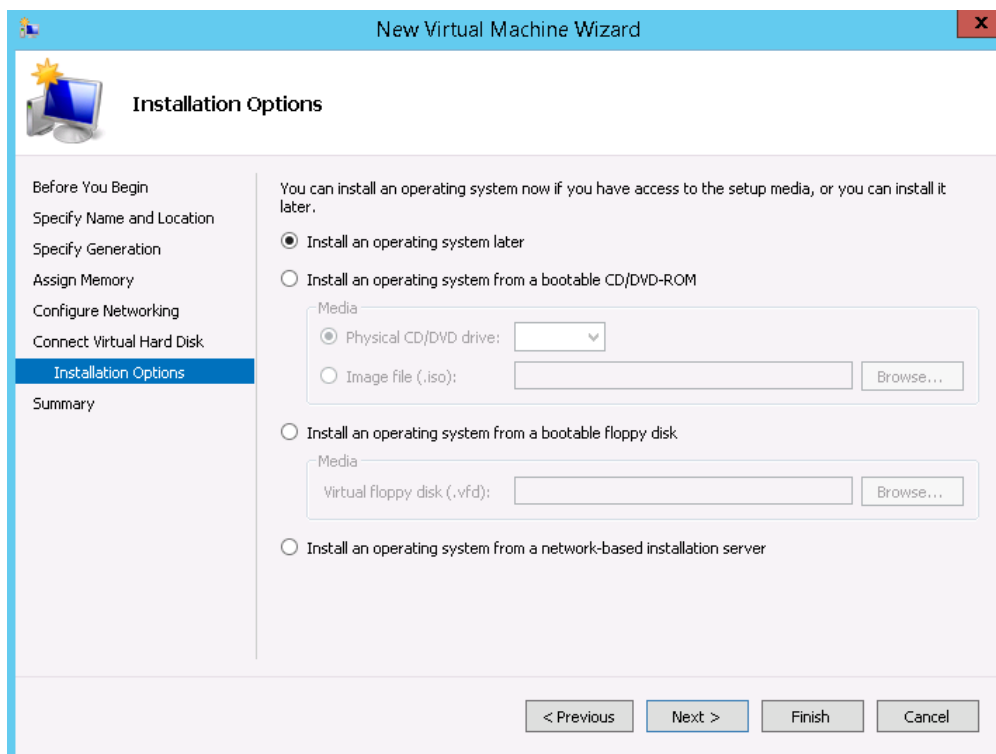
5. In the **Configure Networking** dialog box, select the Hyper-V vSwitch from the **Connection** drop-down list.



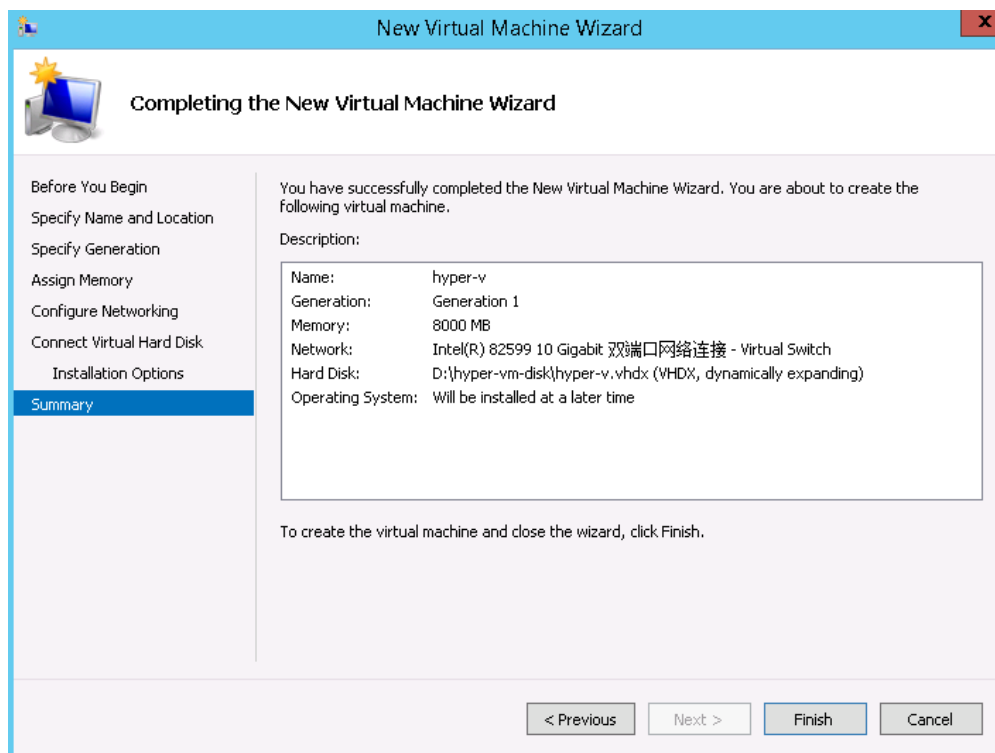
6. In the **Connect Virtual Hard Disk** dialog box, retain default option **Create a virtual hard disk** and create a virtual disk in VHDX format that can be dynamically expanded.



7. In the **Installation Options** dialog box, retain the default option **Install an operating system later**. You can select an ISO image that has been downloaded to install an OS for the VM.



- In the **Summary** dialog box, you can view the VM configurations. To change the configuration, click **Previous**. After confirming all configurations, click **Finish**.



1.4.8 Configuring the vNIC for the Service VM

This section uses CentOS 7 as an example to describe how to configure a network for the VM. Assume that the user-defined network segment is 172.168.0.0 and the IP address of the VM is 172.168.0.10. The procedure is as follows:

- Log in to the VM OS.
- Run the **ifconfig** command to query the NIC used by the current VM.
- Assume that the name of the NIC obtained in 2 is eth0. Run the following command to open the `/etc/sysconfig/network-scripts/ifcfg-eth0` file:

```
vim /etc/sysconfig/network-scripts/ifcfg-eth0
```

Edit the file as follows:

```
USERCTL=no  
ONBOOT=yes  
BOOTPROTO=static  
DEVICE=eth0  
TYPE=Ethernet  
IPADDR=172.168.0.10  
GATEWAY==172.168.0.1  
NETMASK=255.255.255.0  
MTU=8888
```

- Run the **systemctl restart network** command to restart the VM network to make the network configuration take effect.
- Run the following commands to disable the firewall:

```
systemctl disable firewalld  
systemctl stop firewalld
```

6. Edit the `/etc/resolv.conf` file and configure DNS.
`nameserver 114.114.114.114`

After the preceding operations are complete, the service VM can communicate with ECSs in the same VPC. You can run the **ping** command to verify the communication. If you have obtained a NAT gateway and bound an EIP to it, the VM can communicate with the Internet.

2 Monitoring

2.1 Overview

Solution Introduction

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

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

Constraints

- Agent can be installed only on BMSs running a 64-bit Linux OS.
- An agency must be configured for monitoring BMSs. For details, see [How Do I Create an Agency for Server Monitoring of the BMS?](#)
- Only CN North-Beijing1 (**cn-north-1**), CN South-Guangzhou (**cn-south-1**), and Asia Pacific-Hong Kong (**ap-southeast-1**) are supported now.
- Private images do not support this function.

[Table 2-1](#) lists the Linux images that support server monitoring.

Table 2-1 Linux images that support server monitoring

OS Type (64-bit)	Version
SUSE	Enterprise11 SP4, Enterprise12 SP1
CentOS	6.9 and 7.3
EulerOS	2.2
Debian	8.6

2.2 Installing and Configuring the Agent for an Existing BMS

2.2.1 Installing the Agent

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

1. **Adding the Resolved Domain Names:** Add the resolved domain names of regions to the `/etc/resolv.conf` file on the BMS.
2. **Configuring the Security Group:** Download the Telescope package, send metrics, and collect logs.
3. **Installing the Agent:** Manually install the Agent on the BMS.

Adding the Resolved Domain Names

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

Figure 2-1 Adding the resolved domain names

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


NOTE

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

- CN North-Beijing1: 100.125.1.250 and 100.125.21.250
- CN North-Beijing4: 100.125.1.250 and 100.125.129.250
- CN East-Shanghai1: 100.125.1.250,100.125.64.250
- CN South-Guangzhou: 100.125.1.250 and 100.125.136.29
- Asia Pacific-Hong Kong: 100.125.1.250 and 100.125.3.250
- Asia Pacific-Bangkok: 100.125.1.250,100.125.3.250
- LA-Santiago: 100.125.1.250

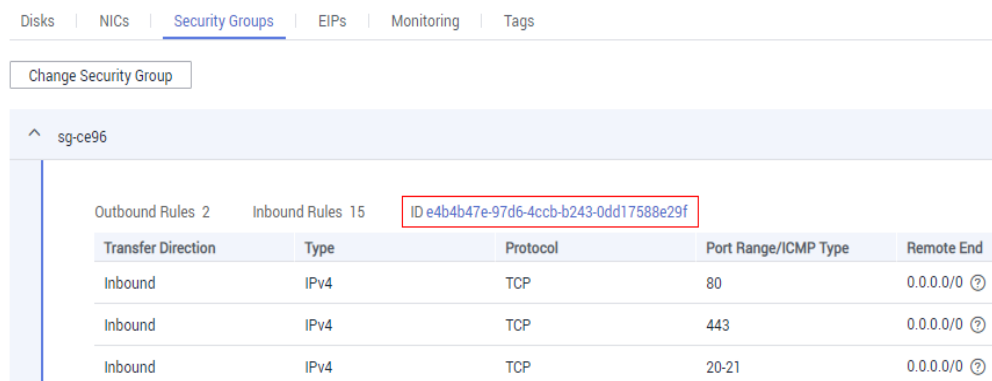
4. Press **Esc** and enter **:wq!** so save the configuration.

Configuring the Security Group

1. On the page showing the BMS details, click the **Security Groups** tab.
2. Click  to expand the security group details, showing the configured security group rules.

- In the upper right corner of the rule list, click the security group ID to go to the **Security Groups** page.

Figure 2-2 Security group rules



- In the **Operation** column, click **Manage Rule**. On the **Outbound Rules** tab page, click **Add Rule** to add a rule based on [Table 2-2](#).

Table 2-2 Security group rules

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

Installing the Agent

- Run the following command to install the Agent:

CN North-Beijing1 (x86):

```
cd /usr/local && curl -k -O https://obs.cn-north-1.myhuaweicloud.com/uniagent-cn-north-1/script/uniagent_install_amd64.sh && bash uniagent_install_amd64.sh
```

CN North-Beijing1 (Kunpeng):

```
cd /usr/local && curl -k -O https://obs.cn-north-1.myhuaweicloud.com/uniagent-cn-north-1/script/uniagent_install_arm64.sh && bash uniagent_install_arm64.sh
```

CN North-Beijing4 (x86):

```
cd /usr/local && curl -k -O https://obs.cn-north-4.myhuaweicloud.com/uniagent-cn-north-4/script/uniagent_install_amd64.sh && bash uniagent_install_amd64.sh
```

CN North-Beijing4 (Kunpeng):

```
cd /usr/local && curl -k -O https://obs.cn-north-4.myhuaweicloud.com/uniagent-cn-north-4/script/uniagent_install_arm64.sh && bash uniagent_install_arm64.sh
```

CN South-Guangzhou (x86):

```
cd /usr/local && curl -k -O https://obs.cn-south-1.myhuaweicloud.com/uniagent-cn-south-1/script/uniagent_install_amd64.sh && bash uniagent_install_amd64.sh
```

CN South-Guangzhou (Kunpeng):

```
cd /usr/local && curl -k -O https://obs.cn-south-1.myhuaweicloud.com/uniagent-cn-south-1/script/uniagent_install_arm64.sh && bash uniagent_install_arm64.sh
```

CN East-Shanghai1 (x86):

```
cd /usr/local && curl -k -O https://obs.cn-east-3.myhuaweicloud.com/uniagent-cn-east-3/script/uniagent_install_amd64.sh && bash uniagent_install_amd64.sh
```

CN East-Shanghai1 (Kunpeng):

```
cd /usr/local && curl -k -O https://obs.cn-east-3.myhuaweicloud.com/uniagent-cn-east-3/script/uniagent_install_arm64.sh && bash uniagent_install_arm64.sh
```

CN East-Shanghai2 (x86):

```
cd /usr/local && curl -k -O https://obs.cn-east-2.myhuaweicloud.com/uniagent-cn-east-2/script/uniagent_install_amd64.sh && bash uniagent_install_amd64.sh
```

CN East-Shanghai2 (Kunpeng):

```
cd /usr/local && curl -k -O https://obs.cn-east-2.myhuaweicloud.com/uniagent-cn-east-2/script/uniagent_install_arm64.sh && bash uniagent_install_arm64.sh
```

CN Southwest-Guiyang1:

```
cd /usr/local && wget https://telescope-cn-southwest-2.obs.cn-southwest-2.myhuaweicloud.com/scripts/agentInstall.sh && chmod 755 agentInstall.sh && ./agentInstall.sh
```

AP-Hong Kong:

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

AP-Bangkok:

```
cd /usr/local && wget https://telescope-ap-southeast-2.obs.ap-southeast-2.myhuaweicloud.com/scripts/agentInstall.sh && chmod 755 agentInstall.sh && ./agentInstall.sh
```

AP-Singapore:

```
cd /usr/local && wget https://telescope-ap-southeast-3.obs.ap-southeast-3.myhuaweicloud.com/scripts/agentInstall.sh && chmod 755 agentInstall.sh && ./agentInstall.sh
```

AF-Johannesburg:

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

LA-Santiago:

```
cd /usr/local && wget https://telescope-la-south-2.obs.la-south-2.myhuaweicloud.com/scripts/agentInstall.sh && chmod 755 agentInstall.sh && ./agentInstall.sh
```

LA-Sao Paulo1:

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

The Agent is installed successfully if the command output similar to [Figure 2-3](#) is displayed.

Figure 2-3 Successful installation

```
telescope_linux_amd64/  
telescope_linux_amd64/uninstall.sh  
telescope_linux_amd64/install.sh  
telescope_linux_amd64/bin/  
telescope_linux_amd64/bin/conf.json  
telescope_linux_amd64/bin/telescope  
telescope_linux_amd64/bin/conf_ces.json  
telescope_linux_amd64/bin/conf_lts.json  
telescope_linux_amd64/bin/record.json  
telescope_linux_amd64/bin/logs_config.xml  
telescope_linux_amd64/bin/agent  
telescope_linux_amd64/telescoped  
telescope_linux_amd64/telescope-1.0.12-release.json  
Current user is root.  
Current linux release version : CENTOS  
Start to install telescope...  
In chkconfig  
Success to install telescope to dir: /usr/local/telescope.  
Starting telescope...  
Telescope process starts successfully.  
[root@ecs-74e5-7 local]#
```

2. After the installation is complete, configure the Agent as instructed in [Manually Configuring the Agent for Linux](#).

2.2.2 (Optional) Managing the Agent

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

NOTE

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

Checking the Agent Status

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

```
service telescoped status
```

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

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

Starting the Agent

Run the following command to start the Agent:

```
/usr/local/telescope/telescoped start
```

Restarting the Agent

Run the following command to restart the Agent:

```
/usr/local/telescope/telescoped restart
```

Stopping the Agent

Run the following command to stop Agent:

```
service telescoped stop
```

NOTE

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

```
/usr/local/telescope/telescoped stop
```

Uninstalling the Agent

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

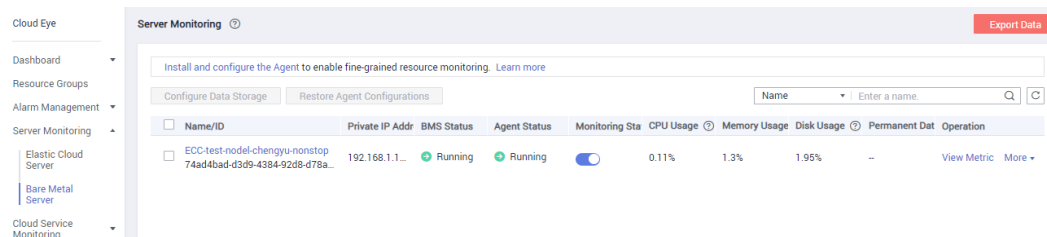
Run the following command to uninstall the Agent:

```
/usr/local/telescope/uninstall.sh
```

2.3 Monitoring Data

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

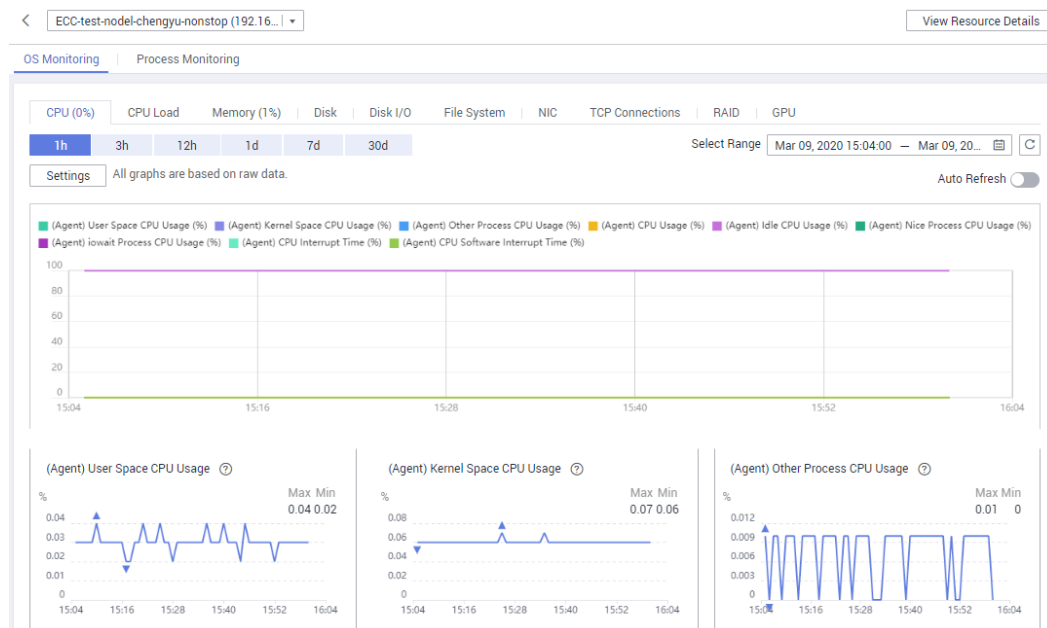
Figure 2-4 Server monitoring



Name/ID	Private IP Addr	BMS Status	Agent Status	Monitoring Sta	CPU Usage	Memory Usage	Disk Usage	Permanent Dat	Operation
ECC-test-nodel-chengyu-nonstop 74ad4bad-d3d9-4384-92d8-d78a...	192.168.1.1...	Running	Running	<input checked="" type="checkbox"/>	0.11%	1.3%	1.95%	--	View Metric More

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

Figure 2-5 Visualized monitoring graph



2.4 Supported Monitoring Metrics (with Agent Installed)

Description

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

NOTE

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

In the CN East-Shanghai1, CN East-Shanghai2, CN North-Beijing1, CN North-Beijing4, and CN South-Guangzhou regions, the Agent of the latest version is used for server monitoring with simplified monitoring metrics. For details, see [Supported Monitoring Metrics \(with Agent Installed and Simplified Metrics\)](#).

Namespace

SERVICE.BMS

Metrics

Supported BMS **OS Monitoring** metrics include CPU metrics listed in [Table 2-3](#), CPU load metrics listed in [Table 2-4](#), memory metrics listed in [Table 2-5](#), disk metrics listed in [Table 2-6](#), disk I/O metrics listed in [Table 2-7](#), file system metrics listed in [Table 2-8](#), NIC metrics listed in [Table 2-9](#), software RAID metrics listed in [Table 2-10](#), and process metrics in [Table 2-11](#).

 NOTE

To monitor software RAID metrics, Agent 1.0.5 or later is required.
Currently, BMSs running the Windows OS cannot be monitored.

Table 2-3 CPU metrics

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

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

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

Table 2-4 CPU load metrics

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
load_average_1	(Agent) 1-Minute Load Average	CPU load averaged from the last 1 minute Obtain its value by dividing the load1/ value in /proc/loadavg by the number of logical CPUs. Run the top command to check the load1 value.	≥ 0	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
load_average_5	(Agent) 5-Minute Load Average	CPU load averaged from the last 5 minutes Obtain its value by dividing the load5/ value in /proc/loadavg by the number of logical CPUs. Run the top command to check the load5 value in the /proc/loadavg file.	≥ 0	BMS	1 minute
load_average_15	(Agent) 15-Minute Load Average	CPU load averaged from the last 15 minutes Obtain its value by dividing the load15/ value in /proc/loadavg by the number of logical CPUs. Run the top command to check the load15 value in the /proc/loadavg file.	≥ 0	BMS	1 minute

Table 2-5 Memory metrics

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mem_available	(Agent) Available Memory	Available memory size of the monitored object Obtain the MemAvailable value by checking the file /proc/meminfo . If it is not displayed in the file, MemAvailable = MemFree + Buffers + Cached Unit: GB	≥ 0 GB	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mem_usedPercent	(Agent) Memory Usage	Memory usage of the monitored object Obtain its value by checking the file /proc/meminfo . Memory Usage = (MemTotal - MemAvailable) / MemTotal Unit: percent	0-100%	BMS	1 minute
mem_free	(Agent) Idle Memory	Amount of memory that is not being used Obtain its value by checking the file /proc/meminfo . Unit: GB	≥ 0 GB	BMS	1 minute
mem_buffers	(Agent) Buffer	Memory that is being used for buffers Obtain its value by checking the file /proc/meminfo . Run the top command to check the KiB Mem:buffers value. Unit: GB	≥ 0 GB	BMS	1 minute
mem_cached	(Agent) Cache	Memory that is being used for file caches Obtain its value by checking the file /proc/meminfo . Run the top command to check the KiB Swap:cached Mem value. Unit: GB	≥ 0 GB	BMS	1 minute

Table 2-6 Disk metrics

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mountPointPrefix_disk_free	(Agent) Available Disk Space	<p>Available disk space of the monitored object</p> <p>Run the df -h command to check the data in the Avail column.</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: GB</p>	≥ 0 GB	BMS	1 minute
mountPointPrefix_disk_total	(Agent) Disk Storage Capacity	<p>Disk storage capacity of the monitored object</p> <p>Run the df -h command to check the data in the Size column.</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: GB</p>	≥ 0 GB	BMS	1 minute
mountPointPrefix_disk_used	(Agent) Used Disk Space	<p>Used disk space of the monitored object</p> <p>Run the df -h command to check the data in the Used column.</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: GB</p>	≥ 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: Disk Usage = Used Disk Space/Disk Storage Capacity.</p> <p>Disk Usage = Used Disk Space/Disk Storage Capacity</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

Table 2-7 Disk I/O metrics

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mountPointPrefix_disk_agt_read_bytes_rate	(Agent) Disks Read Rate	<p>Volume of data read from the monitored object per second</p> <p>The disk read rate is calculated by checking data changes in the sixth column of the corresponding device in the /proc/diskstats file in a collection period.</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: byte/s</p>	≥ 0 bytes/s	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mountPoint Prefix _disk_agt_read_requests_rate	(Agent) Disks Read Requests	<p>Number of read requests sent to the monitored object per second</p> <p>The disks read requests are calculated by checking data changes in the fourth column of the corresponding device in the /proc/diskstats file in a collection period.</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: request/s</p>	≥ 0	BMS	1 minute
mountPoint Prefix _disk_agt_write_bytes_rate	(Agent) Disks Write Rate	<p>Volume of data written to the monitored object per second</p> <p>The disks write rate is calculated by checking data changes in the tenth column of the corresponding device in the /proc/diskstats file in a collection period.</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: byte/s</p>	≥ 0 bytes/s	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mountPointPrefix_disk_agt_write_requests_rate	(Agent) Disks Write Requests	<p>Number of write requests sent to the monitored object per second</p> <p>The disks write requests are calculated by checking data changes in the eighth column of the corresponding device in the /proc/diskstats file in a collection period.</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: request/s</p>	≥ 0	BMS	1 minute
disk_readTime	(Agent) Average Read Request Time	<p>Average amount of time that read requests have waited on the disks</p> <p>The average read request time is calculated by checking data changes in the seventh column of the corresponding device in the /proc/diskstats file in a collection period.</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: ms/count</p>	≥ 0 ms/Count	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
disk_writeTime	(Agent) Average Write Request Time	<p>Average amount of time that write requests have waited on the disks</p> <p>The average write request time is calculated by checking data changes in the eleventh column of the corresponding device in the /proc/diskstats file in a collection period.</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: ms/count</p>	≥ 0 ms/Count	BMS	1 minute
disk_ioUtils	(Agent) Disk I/O Usage	<p>Disk I/O usage of the monitored object</p> <p>Check the data changes in the thirteenth column of the corresponding device in the /proc/diskstats 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)
disk_queue_length	(Agent) Disk Queue Length	<p>Average number of read or write requests to be processed for the monitored disk in the monitoring period</p> <p>The average disk queue length is calculated by checking data changes in the fourteenth column of the corresponding device in the /proc/diskstats file in a collection period.</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: count</p>	≥ 0	BMS	1 minute
disk_write_bytes_per_operation	(Agent) Average Disk Write Size	<p>Average number of bytes in an I/O write for the monitored disk in the monitoring period</p> <p>The average disk write size is calculated by dividing the data changes in the tenth column of the corresponding device by that of the eighth column in the /proc/diskstats file in a collection period.</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: KB/op</p>	≥ 0 KB/operation	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
disk_read_bytes_per_operation	(Agent) Average Disk Read Size	<p>Average number of bytes in an I/O read for the monitored disk in the monitoring period</p> <p>The average disk read size is calculated by dividing the data changes in the sixth column of the corresponding device by that of the fourth column in the /proc/diskstats file in a collection period.</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: KB/op</p>	≥ 0 KB/op	BMS	1 minute
disk_io_svc_time	(Agent) Disk I/O Service Time	<p>Average time in an I/O read or write for the monitored disk in the monitoring period</p> <p>The average disk I/O service time is calculated by dividing the data changes in the thirteenth column of the corresponding device by the sum of data changes in the fourth and eighth columns in the /proc/diskstats file in a collection period.</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: ms/op</p>	≥ 0 ms/op	BMS	1 minute

Table 2-8 File system metrics

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

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
disk_inodesUsedPercent	(Agent) Percentage of Total Inode Used	<p>Percentage of used inodes on the disk</p> <p>Run the df -i command to check data in the IUse% column.</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

Table 2-9 NIC metrics

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

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

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

Table 2-10 Software RAID metrics

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

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

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

Table 2-11 Process metrics

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

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
proc_pHashId_mem	Memory Usage	<p>Memory consumed by a process. pHashId (process name and process ID) is the value of md5.</p> <p>Memory Usage = $RSS \times PAGESIZE / MemTotal$</p> <ul style="list-style-type: none"> Obtain the RSS value by checking the second column of the file /proc/pid/statm. Obtain the PAGESIZE value by running the getconf PAGESIZE command. Obtain the MemTotal value by checking the file /proc/meminfo. <p>Unit: percent</p>	0-100 %	BMS	1 minute
proc_pHashId_file	Opened Files	<p>Number of files opened by a process. pHashId (process name and process ID) is the value of md5.</p> <p>Run the ls -l /proc/pid/fd command to view the number of opened files.</p>	≥ 0	BMS	1 minute
proc_running_count	(Agent) Running Processes	<p>Number of running processes</p> <p>You can obtain the status of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.</p>	≥ 0	BMS	1 minute
proc_idle_count	(Agent) Idle Processes	<p>Number of idle processes</p> <p>You can obtain the status of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.</p>	≥ 0	BMS	1 minute

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

2.5 Supported Monitoring Metrics (with Agent Installed and Simplified Metrics)

Description

This section describes the monitoring metrics reported by BMS in the CN East-Shanghai1, CN East-Shanghai2, CN North-Beijing1, CN North-Beijing4, and CN South-Guangzhou regions. In these regions, the Agent of the latest version is used for server monitoring with simplified monitoring metrics.

 **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 2-12](#) lists the metrics supported by BMS.

Table 2-12 Metrics

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
cpu_usage	(Agent) CPU Usage	<p>CPU usage of the monitored object</p> <p>Obtain its value by checking metric value changes in the /proc/stat file in a collection period.</p> <p>Run the top command to check the %Cpu(s) value.</p> <p>Unit: percent</p>	0-100 %	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
load_aver age5	(Agent) 5- Minute Load Average	CPU load averaged from the last 5 minutes Obtain its value by dividing the load5/ value in /proc/loadavg by the number of logical CPUs. Run the top command to check the load5 value in the /proc/loadavg file.	≥ 0	BMS	1 minute
mem_ usedP ercent	(Agent) Memory Usage	Memory usage of the monitored object Obtain its value by checking the file /proc/meminfo . Memory Usage = (MemTotal - MemAvailable)/MemTotal Unit: percent	0-100 %	BMS	1 minute
moun tPoint Prefix _disk_ free	(Agent) Available Disk Space	Available disk space of the monitored object Run the df -h command to check the data in the Avail column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Unit: GB	≥ 0 GB	BMS	1 minute

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: Disk Usage = Used Disk Space/ Disk Storage Capacity.</p> <p>Disk Usage = Used Disk Space/Disk Storage Capacity</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
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 /proc/diskstats 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)
disk_inodesUsedPercent	(Agent) Percentage of Total Inode Used	Percentage of used inodes on the disk Run the df -i command to check data in the IUse% column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Unit: percent	0-100 %	BMS	1 minute
net_bitSent	(Agent) Inbound Bandwidth	Number of bits sent by this NIC per second Check metric value changes in the /proc/net/dev file in a collection period. Unit: bit/s	≥ 0 bit/s	BMS	1 minute
net_bitRecv	(Agent) Outbound Bandwidth	Number of bits received by this NIC per second Check metric value changes in the /proc/net/dev file in a collection period. Unit: bit/s	≥ 0 bit/s	BMS	1 minute
net_packetRecv	(Agent) NIC Packet Receive Rate	Number of packets received by this NIC per second Check metric value changes in the /proc/net/dev file in a collection period. Unit: count/s	≥ 0 counts /s	BMS	1 minute
net_packetSent	(Agent) NIC Packet Send Rate	Number of packets sent by this NIC per second Check metric value changes in the /proc/net/dev file in a collection period. Unit: count/s	≥ 0 counts /s	BMS	1 minute

Metric ID	Metric	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
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

2.6 FAQs

2.6.1 Why Does Not the Cloud Eye Console Display Monitoring Data or Why Is There a Delay in Data Display After Agent Is Installed and Configured?

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

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

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

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

If monitoring data is still not displayed, check the configuration as instructed in [Manually Configuring the Agent for Linux](#).

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

1. On the management console homepage, choose **Service List > Management & Deployment > Identity and Access Management**.
2. 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 or the sub-project of the region 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.

3. Click **OK**.

The operations to create an agency for server monitoring of the BMS are complete.

3 Backup

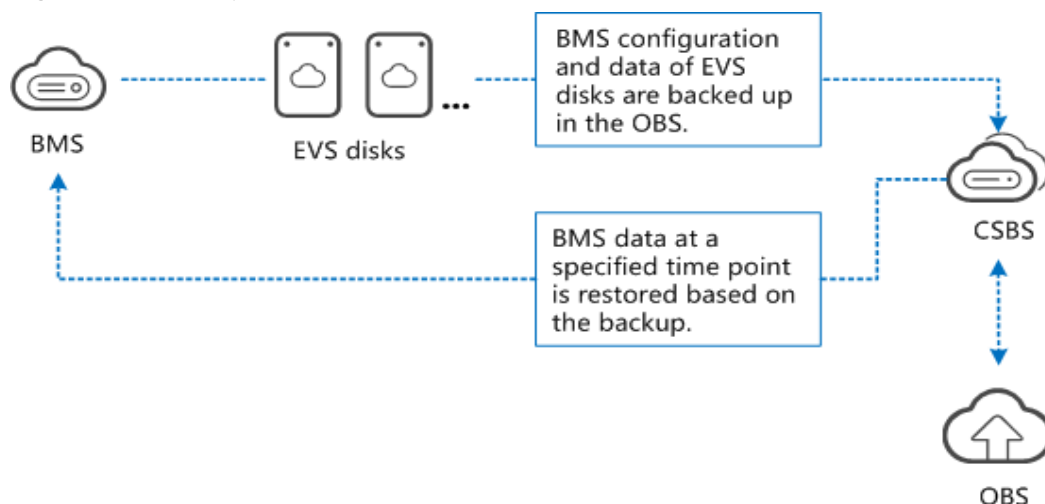
3.1 Overview

Solution Introduction

When a system deployed on a BMS encounters an external virus intrusion, misoperation, or service software bug, you can use the Cloud Server Backup Service (CSBS) which backs up your BMS periodically and automatically (the minimum backup period is one hour). When the preceding fault occurs, you can use the backup to quickly restore your BMS to avoid data loss.

CSBS backs up the BMS configuration and the data of EVS disks of the BMS in the highly reliable Object Storage Service (OBS) to ensure the security of your data.

Figure 3-1 Backup and restoration



Application Scenarios

In scenarios that require high data security, such as enterprise core databases and government and enterprise finance, you are advised to use this scheme to back up BMSs to improve data reliability.

Advantages

- Convenient: You can customize policies for online backup. The backup settings are simple and easy to operate.
- Efficient: Incremental backup and restoration are supported. The RTO is minute-level.
- Reliable: All EVS disks are backed up simultaneously, avoiding data inconsistency caused by the different backup creation time. The OBS durability is 99.999999999%.

Restrictions and Limitations

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

Prerequisites

1. Create a key pair.

To ensure system security, you are recommended to use the key authentication mode to authorize the user who attempts to log in to a BMS. Therefore, you must use an existing key pair or create a new one for remote login authentication.

Create a key pair by following the instructions in [Creating a Key Pair](#). If you already have a key pair, skip this step.

2. Create a VPC.

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

For how to create a VPC, see [Creating a VPC](#).

Procedure

The BMS backup and restoration procedure is as follows:

1. [Creating a Backup Policy](#): Set the backup time, backup period, and retention rules for automatically backup of BMS data.
2. [Purchasing and Configuring a BMS](#): Purchase a BMS that supports quick provisioning and bind a backup policy to it. Initialize EVS disks and deploy required application software.
3. [Creating a BMS Backup](#): The BMS is automatically backed up at fixed time points based on the preset backup policy. You can also manually back up the BMS at any time.
4. [Viewing Backups and Restoring Data](#): View the backups created automatically or manually on the management console. Restore the BMS data to a specified time point as you need.

3.2 Creating a Backup Policy

1. Log in to the HUAWEI CLOUD management console.
2. Choose **Service List > Storage > Cloud Server Backup Service**.
The CSBS console is displayed.
3. Click the **Policies** tab.
4. Click **Create Policy** and configure the following parameters:

Create Policy

Backups generated according to the backup policy are not free.

Basic Information

* Name:

Backup Policy



Enable:

* Execution Time:
The backup interval cannot be shorter than one hour.

* Backup Period: Weekly
 Daily Every days

* Retention Rule: Time Period
 Backup Quantity
 Permanent

Table 3-1 Parameter description

Parameter	Description	Remarks
Name	Backup policy name. It is a string of 1 to 255 characters that can contain only digits, letters, underscores (_), and hyphens (-).	backup_policy
Enable	Whether to enable the backup policy. <ul style="list-style-type: none"> • Enable:  • Disable:  	Only after the backup policy is enabled, the system automatically backs up servers associated with the backup policy and deletes expired backups.

Parameter	Description	Remarks
Execution Time	<p>Execution time</p> <p>A maximum of 24 backup times can be set in a day. The backup interval must not be shorter than one hour. If backup jobs are executed in two consecutive days, the interval between the execution times of the last backup of the former day and the first backup of the latter day must not be shorter than one hour.</p>	<p>00:00, 02:00</p> <p>It is recommended that backup jobs be executed during off-peak hours or when there are no services running.</p>
Backup Period	<p>Dates for executing the backup task.</p> <ul style="list-style-type: none"> • Weekly Specifies on which days of each week the backup task will be executed. You can select multiple days. • Daily Specifies the interval (every 1 to 30 days) for executing the backup task. 	<p>Every 1 day</p> <p>If you select Daily, the first backup time is supposed to be on the day when the backup policy was created. If the creation time of the backup policy is later than the latest execution time, the initial backup will be performed in the next backup cycle.</p> <p>It is recommended that backup jobs be executed during off-peak hours or when there are no services running.</p>

Parameter	Description	Remarks
Retention Rule	<p>Rule that specifies how backups will be retained.</p> <ul style="list-style-type: none"> • Time Period You can choose to retain replicas for one month, three months, six months, or one year, or for any desired number (1 to 99999) of days. • Backup Quantity Specifies the maximum allowed number of backups for a single ECS. The value ranges from 1 to 99999. • Permanent <p>NOTE</p> <ul style="list-style-type: none"> • When the number of backups to be retained has exceeded the value of Backup Quantity, the system automatically deletes the earliest backups. When the retention time of a backup has exceeded the value of Time Period, the system automatically deletes all expired backups. The system will automatically clear backups every other day at 00:00. Each automatic or manual job scheduling also triggers clearing. After a backup is deleted, the other backups can still be used for restoration. • This parameter applies only to backups automatically scheduled by a backup policy. Those backups generated by a manually executed backup policy are not affected by this parameter and are not automatically deleted. You can manually delete them from the backup list. • A maximum of 10 backups are retained for failed periodic backup jobs. They are retained for one month and can be manually deleted. 	Three months

 **NOTE**

More frequent backup creates more backups or retains backups for a longer time, protecting data with a higher level but occupying more storage space. Set an appropriate backup period as required.

- Click **OK** to complete the policy creation.

The created backup policy is displayed in the backup policy list.

Policy Name	Policy Status	Execution Time	Backup Period	Retention Rule	Operation
▼ backup_policy_102726	➔ Enabled	14:30	Every day	90 days	Edit Associate Server More ▼

3.3 Purchasing a BMS

- In the navigation pane on the left, choose **Bare Metal Server** and click **Buy BMS** in the upper right corner.
- Configure the BMS specifications.
 - Flavor:** Select a BMS flavor, for example, **physical.s4.xlarge**.

Flavor name	CPU	Memory	Local Disk	Extended Configuration
Sold physical.m2.medium	96 core 4*24Co...	32*64 GB DIMM	2*600GB SAS System Disk RA...	2x2*10GE
<input type="radio"/> physical.s3.large	20 core Intel Xe...	128 GB DDR4	2*600G SAS System Disk RAL...	2 x 2*10GE
<input type="radio"/> physical.s4.3xlarge	44 core Intel Xe...	384 GB DDR4	NA	2 x 2*10GE
Sold physical.s4.large	20 core Intel Xe...	192 GB DDR4	NA	2 x 2*10GE
<input type="radio"/> physical.s4.medium	20 core Intel Xe...	128 GB DDR4	NA	2 x 2*10GE
<input checked="" type="radio"/> physical.s4.xlarge	28 core Intel Xe...	192 GB DDR4	NA	2 x 2*10GE

- Disk:** Set system disk parameters and add a data disk.

EVS

If you select a Linux image, change the disk identifier in the fstab file to **UUID** after the BMS is created. Otherwise, the BMS OS or service will be unavailable due to a failure to find the disk **UUID** after the BMS is restarted.

System Disk GB | IOPS limit 800, IOPS burst limit 2,200

Common I/O - 150 +

Data Disk GB | IOPS limit 700, IOPS burst limit 2,200 [Delete](#)

Common I/O - 100 +

Share [?](#)

- Auto Backup:** Select **Enable automatic backup** and select the backup policy created in section **Creating a Backup Policy** from the drop-down list.

Auto Backup [?](#) Enable auto backup Recommended It is good practice to back up BMS data. Standard charges apply for successfully generated BMS backups. 1 GB of backup data costs ¥0.268. (To obtain more cost-effective services, [purchase a package](#)).

Backup Policy [Manage Backup Policy](#)

backup_policy_102726... ▼

Note that DSS, DESS, and shared disks cannot be attached to the BMS if auto backup is enabled.

- Click **Buy Now** to create the BMS.

After about five minutes, the BMS is created and its status changes to **Running**.

<input type="checkbox"/> bms-da7e_test 011cacf5-e2c8-...	➔ Running	CPU: Intel Xeon... Memory: 384GB... Local Disk: NA Extended Config...	CentOS 7.4 64bi...	192.168.11.240	-	AZ3	Yearly/Monthly 30 days remaining until expira	Remote Login More ▼
---	-----------	--	--------------------	----------------	---	-----	--	---

- Log in to the BMS remotely or using the key pair, initialize the data disk, and deploy your applications.

3.4 Creating a BMS Backup

Automatic Backup

After a backup policy is bound to a BMS, the BMS will be periodically backed up based on the backup policy. You can focus on service running rather than whether backups are created.

Manual Backup

If you want to back up BMS data at any time, you can manually create BMS backups.

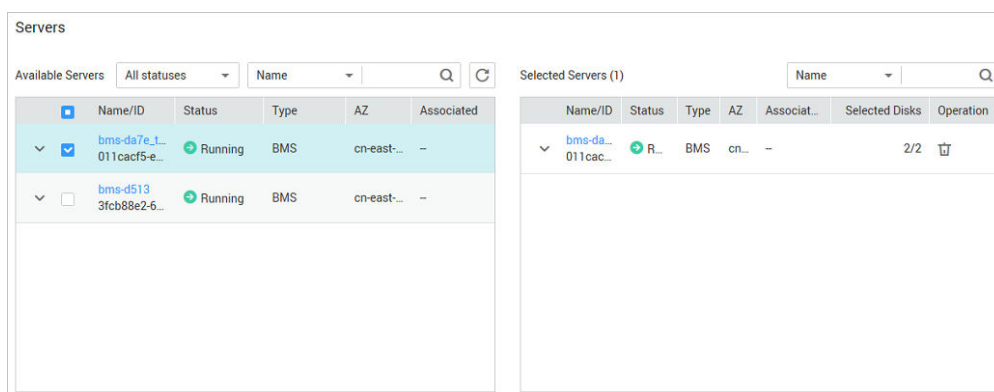
1. On the CSBS console, click **Create CSBS Backup** in the upper right corner.
2. Select a BMS from the server list and select the disks to be backed up.

NOTICE

To ensure post-restoration data consistency, you are advised to back up the entire BMS.

If you want to back up specific disks to reduce costs, ensure that data on these disks is independent of the rest. Otherwise, data inconsistency may occur.

For example, data of the Oracle application is scattered in different disks. If only some of the disks are backed up, restoration restores only the data of the disks that have been backed up, with the rest data unchanged. As a result, data inconsistency occurs and even the application cannot be started.



3. Select **Backup Now**, enter the name and description, and select **Enable** for **Full Backup**.

An auto backup policy has been bound to the BMS. To replace the policy, select **Auto Backup** and select a new policy from the drop-down list. The original policy will be unbound from the BMS and the new policy will be bound to the BMS.

4. Click **Apply Now**. On the **Details** page, confirm the specifications and click **Submit**.

On the **Backups** tab page, if the generated backups are in the **Available** state, the one-off manual backup task is successful.

3.5 Viewing Backups and Restoring Data

View Backups Created Automatically and Manually

You can view all backups on the **Backups** page.

Names of backups created automatically start with **auto** and names of backups created manually start with **manual**.

Figure 3-2 Backup list

<input type="checkbox"/>	Backup Name	Backup Type	Backup Status	Created	Server Name	Server Type	Operation
▼ <input type="checkbox"/>	autobk_83ff	Enhanced backup	Available	Sep 18, 2018 17:30:22 GMT...	bms-da7e_test	BMS	Create Image Restore Delete
▼ <input type="checkbox"/>	manualbk_3acb	Enhanced backup	Available	Sep 18, 2018 17:15:33 GMT...	bms-da7e_test	BMS	Create Image Restore Delete

Click ▼ in the row that contains a backup to expand the details, including the disk backups of the BMS.

Figure 3-3 Backup details

Backup		Server	
Backup Name	autobk_83ff	Server Name	bms-da7e_test
Backup ID	5140e4cd-9603-4e04-8775-cd382895a4ac	Server ID	011cacf5-e2c8-4712-81e6-c89a6e161cc1
Backup Status	Available	AZ	AZ3
Created	Sep 18, 2018 17:30:22 GMT+08:00		
Description	--		

Backup Name	Backup Status	Disk Name	Function	Disk Capacity (GB)	Shared
autobk_63a3_bms-da7e_test...	Available	bms-da7e_test-volume-0000	System Disk	150	No
autobk_63a3_bms-da7e_test...	Available	bms-da7e_test-volume-0001	Data Disk	100	No

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

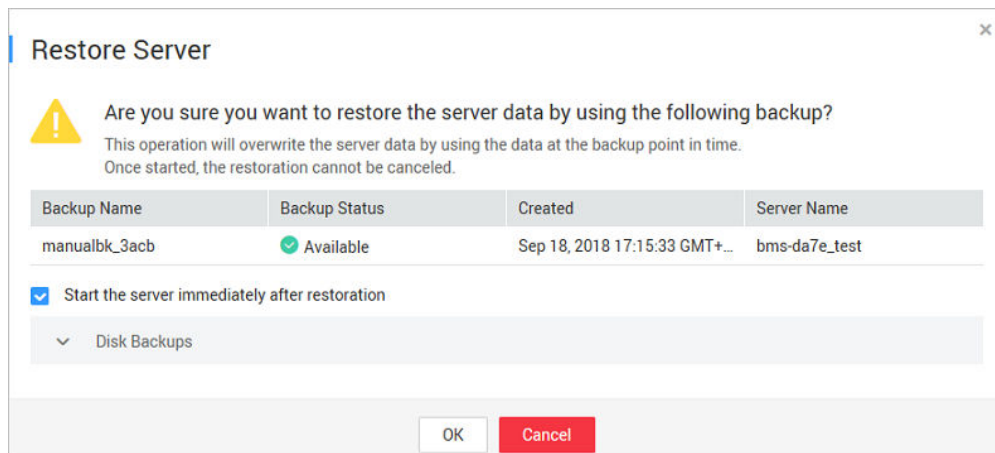
Figure 3-4 Backup Tree

bms-da7e_test-volume-0000 System Disk 150GB		Detach	
ID	c123f704-27d7-4823-b18f-d6b2b2029384	Name	bms-da7e_test-volume-0000
Description		Bootable	Yes
Created	Sep 18, 2018 16:59:30 GMT+08:00	Device Name	/dev/sda
Image	CentOS 7.4 64bit for BareMetal	Capacity (GB)	150
Order	CS1809181658GGNF1	AZ	AZ3
Backups	2	Billing Mode	Yearly/Monthly
Type	Common I/O	Expire At	Oct 18, 2018 23:59:59 GMT+08:00
Disk Sharing	No	Max. Throughput	90MB/s
Max. IOPS	800/2,200 2 IOPS per GB with an IOPS limit of 800 and an IOPS burst limit of 2,200.	Backup Tree	<ul style="list-style-type: none"> manualbk_feab_bms-da7e_test-volume-0000 Sep 18, 2018 17:16:02 GMT+08:00 autobk_63a3_bms-da7e_test-volume-0000 Sep 18, 2018 17:30:45 GMT+08:00
Device Type	SCSI		
Device Identifier	688860300008f7aafa162b5e26007430		

Restore Data

When an external virus intrusion, misoperation, or service software bug occurs, you need to restore BMS configuration and EVS disk data to those at a specified time point to minimize loss.

1. On the CSBS console, locate the row that contains a backup created at a specified time point and click **Restore** in the **Operation** column.
The **Restore Server** dialog box is displayed.

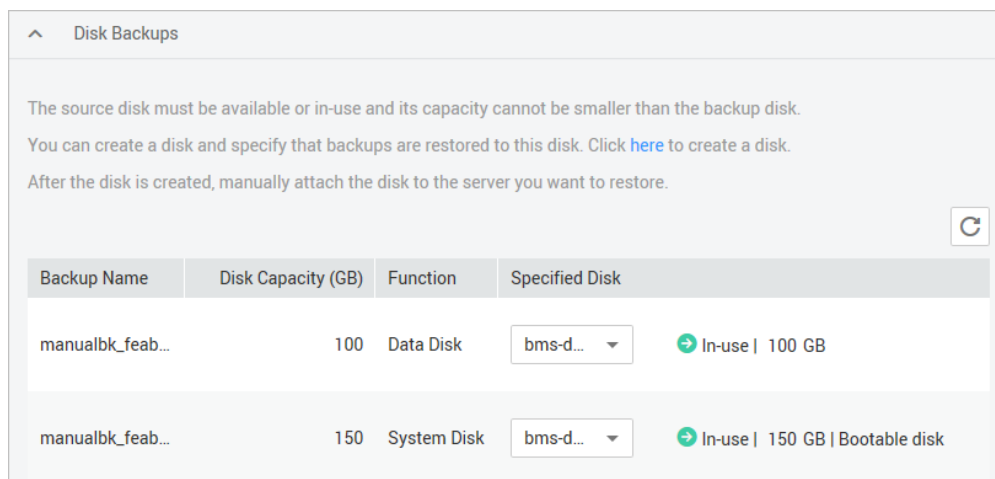


2. Select **Start the server immediately after restoration**.
If you deselect **Start the server immediately after restoration**, manually start the server after the restoration is complete.

NOTICE

Servers are shut down during restoration. Therefore, perform restoration jobs during off-peak hours.

3. In the **Specified Disk** drop-down list, select the target EVS disk to which the backup will be restored.



NOTE

- If the server has only one disk, the backup is restored to the only disk by default.
 - If the server has multiple disks, the backup will be restored to the original disk. Alternatively, you can select another disk for the restoration. The specified EVS disk must have an equal capacity to or a larger capacity than the original EVS disk.
 - Data on data disks cannot be restored to system disks.
4. Click **OK** and confirm the restoration is successful.
You can view the restoration status in the backup list. The restoration is successful if **Backup Status** of the backup becomes **Available**. Then check whether your services are restored.

A Change History

Release On	Description
2020-01-13	This issue is the fourth official release. Optimized descriptions in Overview .
2019-04-25	This issue is the third official release. Added the VMware on BMS solution.
2019-03-18	This issue is the second official release. Added the Hyper-V on BMS solution.
2018-04-15	This issue is the first official release.