

### 9.2.1

This book is the hardware maintenance guide for Hitachi Content Platform (HCP) systems. It contains instructions for adding new nodes and storage, replacing failed components, and performing other HCP hardware maintenance procedures.

© 2007, 2020 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively “Hitachi”). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. “Materials” mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at <https://support.hitachivantara.com/en-us/contact-us.html>.

**Notice:** Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

**Notice on Export Controls.** The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.

---

# Contents

<b>Preface.....</b>	<b>6</b>
Intended audience.....	6
Product version.....	6
Release notes.....	6
Related documents.....	7
Accessing product documentation.....	8
Getting help.....	8
Comments.....	8
<b>Chapter 1: HCP system overview.....</b>	<b>9</b>
Introduction to Hitachi Content Platform.....	9
HCP VM system components and architecture.....	9
Host platform.....	10
Compute.....	10
Storage.....	10
Front-end network.....	11
Back-end network.....	11
Management port network.....	11
Dedicated database volume.....	11
Hardware monitoring and alerting.....	12
HCP software.....	12
Storage licensing.....	12
<b>Chapter 2: Configuration guidelines for the HCP VM environment.....</b>	<b>13</b>
Supported KVM versions.....	13
HCP VM hardware requirements.....	13
HCP VM system limits.....	14
HCP VM availability considerations.....	15
<b>Chapter 3: Installing KVM.....</b>	<b>16</b>
Configuring the BIOS for KVM.....	16
Logging in to the KVM host.....	16
Making sure virtualization technology is enabled.....	16
(Optional) Connecting a monitor and keyboard to your KVM host.....	17

(Optional) Entering the BIOS.....	17
(Optional) Enabling virtualization technology.....	17
Installing the KVM packages.....	17
Logging in to the KVM host.....	18
Installing and initiating the KVM packages.....	18
Installing Virtual Machine Manager on your local computer.....	18
<b>Chapter 4: Configuring KVM networking.....</b>	<b>20</b>
Configuring the KVM host for networking.....	20
(Optional) Connecting a monitor and keyboard to your KVM host.....	20
Backing up the interface configuration files.....	20
Enabling the bonding kernel module.....	21
Creating the bridge network files.....	21
Creating the bond network files.....	23
Enslaving the front-end interface configuration files.....	24
Enslaving the back-end interface configuration files.....	26
Restarting networking on the KVM host network services.....	27
Restarting the individual KVM host.....	27
Verifying that the KVM host network has been updated.....	28
<b>Chapter 5: Maintaining the HCP VM system.....</b>	<b>30</b>
Adding logical volumes.....	30
Moving storage node databases to optimal volumes.....	35
Deleting databases from older database volumes.....	37
Adding HCP VM nodes.....	39
Recovering storage nodes.....	43
Recovering storage nodes and preserving volumes.....	43
Recovering storage nodes and clearing volumes.....	45
<b>Chapter 6: Deploying the HCP VM system.....</b>	<b>50</b>
Prerequisites.....	50
Deploying the HCP VMsystem.....	50
Downloading the .iso installation files.....	50
Copying the .iso file and sending the Zip file to the KVM host.....	51
Adding an HCP VM node connection.....	51
Creating the HCP VM node.....	52
Customizing the HCP VM node for HCP configuration.....	59
Performing the OS installation.....	64
Changing the install user password.....	66
Installing the HCP software.....	67
Identifying the nodes in the HCP VM system.....	69
Configuring the HCP system.....	70

Running the HCP installation.....	75
Verifying the HCP software installation.....	80
Monitoring and alerting.....	81
Software monitoring.....	82
HCP VM resource monitoring.....	82
HCP VM diagnostic menu.....	83
<b>Chapter 7: Configuring HCP monitoring with Hitachi Remote Ops.....</b>	<b>84</b>
Enabling SNMP in HCP.....	84
Configuring Hitachi Remote Ops.....	85
Log in to Hitachi Remote Ops.....	85
Set the base configuration.....	85
(Optional) Configure transport agents.....	86
Identify the HCP system.....	87
<b>Appendix A: Configuring SAN storage for the KVM host.....</b>	<b>89</b>
Logging in to the KVM host.....	89
Configuring multipathing.....	89
Creating the physical volume on the multipath disk.....	91
Mounting the file system.....	91
Adding the new storage pool to the KVM host.....	92
<b>Appendix B: Creating an HCP VM node using the command line...</b>	<b>95</b>
Logging in to the KVM host.....	95
Performing the command line initialization.....	95

---

## Preface

This book is the setup guide for Hitachi Content Platform Virtual Machine (HCP VM) systems. This book provides the information you need to deploy a virtualized HCP system in your Kernel-based Virtual Machine (KVM) environment.

### Intended audience

This book is intended for people who are responsible for deploying an HCP VM system in a KVM environment. This book assumes you have experience with computer networking and creating virtual machines. This book also assumes you have familiarity with KVM concepts, and that you have a basic understanding of HCP systems.

### Product version

This book applies to release 9.2.1 or later of Hitachi Content Platform.

### Release notes

Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>.

## Related documents

The following documents contain additional information about Hitachi Content Platform:

- *HCP System Management Help*

This Help system is a comprehensive guide to administering and using an HCP system. The Help contains complete instructions for configuring, managing, and maintaining HCP system-level and tenant-level features and functionality. The Help also describes the properties of objects stored in HCP namespaces and explains how to access those objects.

- *HCP Tenant Management Help*

This Help system contains complete instructions for configuring, managing, and maintaining HCP namespaces. The Help also describes the properties of objects stored in HCP namespaces and explains how to access those objects.

- *Managing the Default Tenant and Namespace*

This book contains complete information for managing the default tenant and namespace in an HCP system. The book provides instructions for changing tenant and namespace settings, configuring the protocols that allow access to the namespace, managing search and indexing, and downloading the installation files for HCP Data Migrator. The book also explains how to work with retention classes and the privileged delete functionality.

- *Using the Default Namespace*

This book describes the file system HCP uses to present the contents of the default namespace. This book provides instructions for using HCP-supported protocols to store, retrieve, and deleting objects, as well as changing object metadata such as retention and shred settings.

- *Using HCP Data Migrator*

This book contains the information you need to install and use HCP Data Migrator (HCP-DM), a utility that works with HCP. This utility enables you to copy data between local file systems, namespaces in HCP, and earlier HCAP archives. It also supports bulk delete operations and bulk operations to change object metadata. Additionally, it supports associating custom metadata and ACLs with individual objects. The book describes both the interactive window-based interface and the set of command-line tools included in HCP-DM.

- *Installing an HCP System*

This book provides the information you need to install the software for a new HCP system. It explains what you need to know to successfully configure the system and contains step-by-step instructions for the installation procedure.

- *Deploying an HCP VM System on ESXi*

This book contains all the information you need to install and configure an HCP VM system. The book also includes requirements and guidelines for configuring the VMware® environment in which the system is installed.

- Installing an HCP SAIN System - Final On-site Setup

This book contains instructions for deploying an assembled and configured single-rack HCP SAIN system at a customer site. It explains how to make the necessary physical connections and reconfigure the system for the customer computing environment. It also contains instructions for configuring Hitachi Remote Ops to monitor the nodes in an HCP system.

- Installing an HCP RAIN System - Final On-site Setup

This book contains instructions for deploying an assembled and configured single-rack HCP RAIN system at a customer site. It explains how to make the necessary physical connections and reconfigure the system for the customer computing environment. It also contains instructions for configuring Hitachi Remote Ops to monitor the nodes in an HCP system.

## Accessing product documentation

Product user documentation is available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

## Getting help

[Hitachi Vantara Support Connect](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: [https://support.hitachivantara.com/en\\_us/contact-us.html](https://support.hitachivantara.com/en_us/contact-us.html).

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to [community.hitachivantara.com](https://community.hitachivantara.com), register, and complete your profile.

## Comments

Please send us your comments on this document to [doc.comments@hitachivantara.com](mailto:doc.comments@hitachivantara.com). Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

**Thank you!**



---

# Chapter 1: HCP system overview

This chapter introduces Hitachi Content Platform. It describes the architecture of an HCP VM system installed in a KVM environment.

## Introduction to Hitachi Content Platform

Hitachi Content Platform (HCP) is a distributed storage system designed to support large, growing repositories of fixed-content data. An HCP system consists of both hardware (physical or virtual) and software.

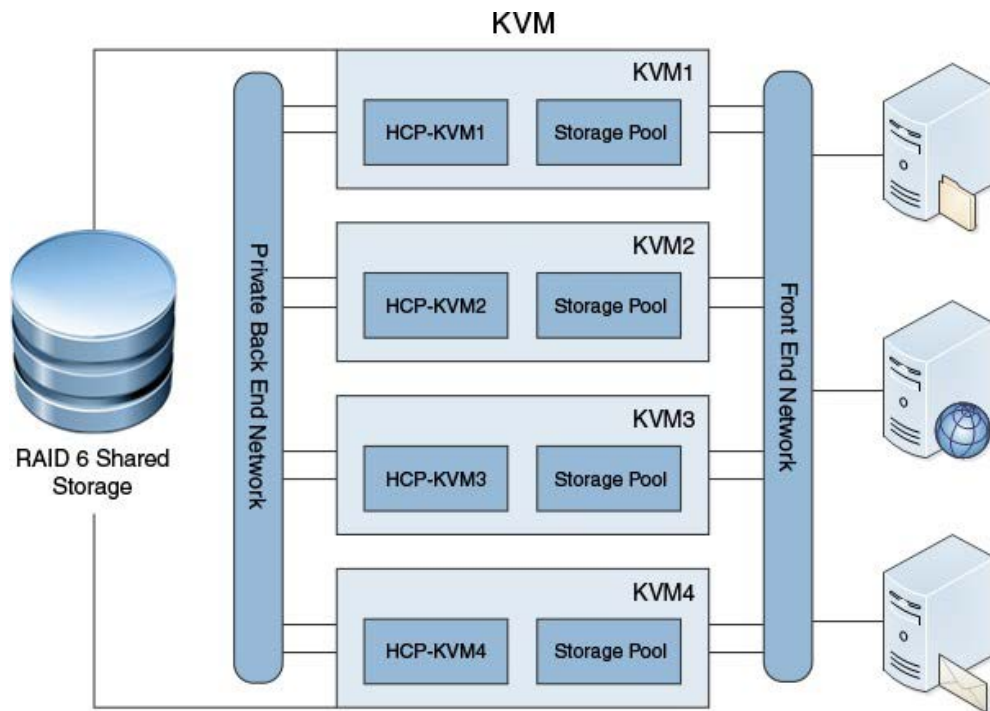
HCP stores objects that include data and the metadata that describes the data. HCP distributes these objects across the storage space. HCP represents objects either as URLs or as files in a standard file system.

An HCP *repository* is partitioned into namespaces. Each *namespace* consists of a distinct logical grouping of objects with its own folder structure. Namespaces are owned and managed by tenants.

HCP gives access to objects through a variety of industry-standard protocols and various HCP interfaces.

## HCP VM system components and architecture

This section describes the components and architecture of an HCP VM system. The following figure shows the architecture of an HCP VM system running on KVM infrastructure.



## Host platform

In an HCP VM system, each HCP VM node runs in a virtual machine on a KVM host, but only one HCP VM node should be present on a single KVM host. The KVM host, however, can have additional virtual machines running other applications.

## Compute

An HCP VM node must have at least eight virtual CPUs and 32 GB of allocated RAM.

The minimum processing requirements ensure that HCP VM system performance is not slowed by multiple client logins, and that activities such as encryption, scheduled services, and routine database maintenance continue running.

If you are deploying an HCP VM small-instance configuration, each HCP VM node must have at least four virtual CPUs and 16 GB of allocated RAM.

## Storage

The KVM hosts that run HCP VM nodes use shared storage. This shared storage should be backed by SAN storage with RAID 6 protection, or by storage that is internal to the KVM hosts. With SAN storage, each KVM host is connected to the storage array through two Fibre Channel switches, which helps make sure data is available if a single host fails.

SAN storage is presented to each KVM host as logical units (represented by LUNs). On each KVM host, the logical units are formatted into an ext4 file system. Using Virtual Machine Manager, you create a storage pool from the file system on each host. You allocate virtual machine disks, which HCP VM reads as internal drives, out of the storage pools. The HCP operating system and software are installed on the virtual machine disks.

With storage that is internal to the KVM hosts, each LUN allocated from the internal storage corresponds to a storage pool. Like SAN storage, the physical storage underlying the internal storage must be RAID protected.

## Front-end network

The HCP front-end network is used for client and management access. For HCP front-end networks, the KVM host has a bonded interface of two physical network interface cards (*pNICs*). Having two pNICs dedicated to HCP ensures redundancy and consistent performance.HCP

## Back-end network

The HCP private back-end network is used for internode communication and data transfer. The KVM host has a bonded interface that maps two pNICs to the KVM server.

The physical NICs dedicated to the back-end network must be connected to two physical switches on an isolated network. On all KVM hosts, pNIC-1 must connect to the same physical switch (switch1), and pNIC-2 on all KVM hosts must connect to the same second physical switch (switch2). The physical switches must be cabled for an inter-switch connection. To guarantee data security and HCP reliability, back-end switches must be configured with spanning tree disabled and multicast traffic enabled. The back-end switches must be at least 1 GbE and must be dedicated to HCP.

To support HCP VM inter-node communication, the back-end network must have multicast enabled. In most cases, enabling multicast on the switch is not sufficient to allow for multicast traffic. Most switches need additional configuration parameters. To allow multicast traffic between the HCP VM nodes, follow the switch-vendor documentation to configure the network.



**Note:** The HCP VM system can be deployed without multicast enabled on the switches. However, if the switches are not configured for multicast, the HCP VM nodes cannot communicate.

If the HCP VM back-end network is on a public network, the HCP VM system should reside on its own VLAN.

## Management port network

The HCP management port is a separate network that can be used to isolate management access from client access. For the management port network, a single virtual Network Interface Card (*NIC*) must be created on the ESXi host on a single physical NIC.

## Dedicated database volume

A separate volume can be created on each HCP virtual machine to separate the storage of user data and metadata from the HCP database. During the installation, you are asked if you want a dedicated database volume if each virtual machine is configured with three or more data disks, at least one of which is greater than 50 GB.

## **Hardware monitoring and alerting**

HCP hardware has built-in redundancy, monitoring, alerting, and failover behavior that cannot be used in a virtualized environment.

To maintain performance and data integrity on an HCP VM system, the system must be connected to Hitachi Remote Ops.

Hitachi Remote Ops supports Hitachi hardware only. To monitor hardware supplied by other vendors, you must use third-party monitoring tools.

## **HCP software**

An HCP VM system uses the same HCP operating system and software as HCP RAIN and HCP RAIN systems. Data is RAID protected, and HCP policies and services ensure data integrity, data security, and storage optimization. HCP VM management and data access interfaces are the same as that of HCP HCP RAIN and HCP RAIN systems.

Because HCP VM software is not bound to hardware, the software does not support zero-copy failover, and hardware cannot be monitored in the System Management Console.

## **Storage licensing**

HCP VM systems come with a basic storage license that includes two terabytes of active and HCP S Series storage. The basic storage license also includes two terabytes of extended storage. If you need additional storage, contact your Hitachi Vantara sales representative.

---

## Chapter 2: Configuration guidelines for the HCP VM environment

This chapter describes the requirements and guidelines for installing and using an HCP VM system.

### Supported KVM versions

HCP VM supports multiple versions of KVM.

For a list of supported KVM versions, see the release notes on the Hitachi Content Platform:

[https://knowledge.hitachivantara.com/Documents/Storage/Content\\_Platform](https://knowledge.hitachivantara.com/Documents/Storage/Content_Platform)

### HCP VM hardware requirements

HCP VM systems can be configured in two ways: standard and small instance.

#### Standard HCP VM configuration

The following are needed to deploy a standard configuration:

- A shared SAN storage using a RAID 6 system
- A minimum of four 1.2 TB LUNs



**Note:** Due to the overhead associated with disk formatting and database storage, the estimated usable storage available with this minimum is approximately 3.66 TB.

- A minimum of four HCP VM nodes
- A minimum of two 500 GB VMDKs on each HCP VM node
- A minimum of eight virtual CPUs on each HCP VM node
- A minimum of 32 GB of RAM on each HCP VM node.



**Note:** To avoid the possibility of slowing system performance, do not commit more than 256 GB of RAM for an HCP VM node.

- Two physical NICs on each KVM host dedicated to the HCP back-end network
- Two physical NICs for the KVM management network and HCP VM front-end network

- Two-port Fibre Channel HBA cards for shared SAN storage connectivity (when applicable)
- A minimum of 2 GB of physical RAM for KVM host management

### Small instance HCP VM configuration

A small-instance HCP VM system has the same requirements as a standard configuration with the following exceptions:

- A minimum of 4 virtual CPUs on each HCP VM node
- A minimum of 16 GB of RAM on each HCP VM node

A small-instance deployment can support:

- Five tenants
- 25 namespaces
- A single active/passive replication link
- An ingest duty cycle of 12 hours per day, 5 days per week

Other factors can affect whether the small-instance deployment meets your performance requirements, such as heavy metadata query engine (MQE) querying or object and folder counts above published maximums.

## HCP VM system limits

There are system limits for an HCP VM system.

For a standard HCP VM system configuration, the following limits are supported:

- HCP VM nodes: 40
- Data LUNs: Limited by the number of device slots available for LUNs in the VirtIO-blk para-virtualized storage back-end, which depends on the number of other devices configured for the guest OS that also use the VirtIO-blk back-end. In a typical HCP configuration, 17 slots are available.
- Virtual machine disks: 15.90 TB

For an HCP VM system small-instance configuration, the following limits are supported:

- HCP VM nodes: 16
- Data LUNs: Limited by the number of device slots available for LUNs in the VirtIO-blk para-virtualized storage back-end, which depends on the number of other devices configured for the guest OS that also use the VirtIO-blk back-end. In a typical HCP configuration, 17 slots are available.
- Virtual machine disks: 15.90 TB

For more information about the supported limits for the file system partition used for HCP storage pools on a KVM host, see the Linux documentation.

For more information about HCP supported limits, see the *HCP Release Notes*.

## **HCP VM availability considerations**

An HCP cluster is considered in a state of continuous availability if there is one HCP VM node per KVM host, and if over half of the HCP VM nodes are healthy and running. When the HCP cluster is in this state, the system can survive a single KVM host failure without affecting HCP functionality.

If your HCP cluster is not in a state of continuous availability because you have multiple HCP VM nodes per KVM host and one of your KVM hosts fails, the HCP VM system enters a state of metadata unavailability. Metadata unavailability prohibits HCP namespaces from accepting write requests. The data stored in the affected nodes becomes inaccessible until the HCP system repairs itself. The repair process can take between one and five minutes.

HCP VM systems do not support zero-copy failover. If a namespace has a data protection level of one, the loss of a single HCP VM node causes the node to enter a state of data unavailability until the node is restored.

Oversubscribing the CPU, RAM, or disk of KVM hosts can cause HCP system instability.

---

## Chapter 3: Installing KVM

This section covers how to configure the KVM host BIOS so that it can run KVM. This section also covers how to install the necessary KVM packages on your KVM host and local Linux machine. The packages listed in this section are required to deploy an HCP VM system according to the guidelines in this manual.

### Configuring the BIOS for KVM

On some KVM hosts, virtualization technology is disabled by default in the BIOS. You need to enable virtualization technology for the KVM packages to run.

### Logging in to the KVM host

#### Procedure

1. Open a new terminal.
2. Use SSH to log into the KVM host.

### Making sure virtualization technology is enabled

#### Procedure

1. Enter the following command to verify that virtualization is enabled on the KVM host:

```
lsmod | grep kvm
```

The output should contain one of the following lines of text:

```
kvm_intel
```

```
kvm_amd
```

2. If the output does not contain `kvm_intel` or `kvm_amd`, enter the following command to determine whether KVM is disabled by the BIOS:

```
dmesg | grep -i kvm
```

If the output contains the following line of text, virtualization technology must be enabled on the BIOS:

```
kvm: disabled by bios
```



## (Optional) Connecting a monitor and keyboard to your KVM host

### Before you begin

To configure networking for a KVM host, you will need the following items:

- USB keyboard
- VGA monitor

Connect your monitor and keyboard to the KVM host.

## (Optional) Entering the BIOS

To enter the BIOS, you must restart the KVM host and access the BIOS as the KVM host powers on:

### Procedure

1. Press and hold the power button until the KVM host shuts down.
2. Press the power button again and wait for the KVM host to restart.
3. On the startup window, press the button that accesses the BIOS.

## (Optional) Enabling virtualization technology

### Procedure

1. After you are in the BIOS, navigate to **Processor Settings**, then press **Enter**
2. Navigate to **Virtualization Technology**, then press **Enter** to change **Virtualization Technology** to **Enabled**.
3. Press **Esc** until the **Exit** window appears.
4. Navigate to **Save changes and exit**, then press **Enter**.

## Installing the KVM packages

To deploy an HCP VM node on your KVM host, you must install several KVM packages on your host. You also must install a separate instance of Virtual Machine Manager on your local machine.

These are the packages to install:

- **libvirt**: a virtualization API and toolkit that manages virtualization hosts; libvirt brings together every server-side RPM required by libvirt.
- **libvirt-daemon**: a server-side daemon that is required to manage the KVM hypervisor.
- **libvirt-daemon-kvm**: brings together the server-side daemon, drivers, and the KVM binaries required for hardware-accelerated virtualization.
- **qemu-kvm**: installs all KVM-specific libraries.
- **virt-manager**: a user interface for performing administrator tasks on virtual machines.

- `guestfs-browser`: a graphic interface for browsing the virtual machine file system and disk images.
- `libguestfs-tools`: a set of tools for accessing and modifying virtual machine disk images.
- `python-libguestfs`: a `libguestfs` tools Python library.
- `virt-top`: monitors a KVM guest virtual machine CPU, memory, and performance.
- `virt-install`: CLI command support for creating guest virtual machines for KVM.
- `bridge-utils`: needed to create and manage bridge devices; `bridge-utils` is used to set up networks for a hosted virtual machine.
- `virt-viewer`: displays the graphic console of a virtual machine.

## Logging in to the KVM host

### Procedure

1. Open a new terminal.
2. Use SSH to log into the KVM host.

## Installing and initiating the KVM packages

### Procedure

1. Enter the following command to download and install the packages:
 

```
sudo dnf install libvirtlibvirt-daemon-kvm qemu-kvm virt-manager
guestfs-browser libguestfs-tools python-libguestfs virt-top virt-
install bridge-utils virt-viewer
```
2. After the packages are installed on the node, enter the following command to start `virtlogd`:
 

```
sudo systemctl start virtlogd
```
3. Enter the following command to start `libvirtd`:
 

```
sudo systemctl start libvirtd
```
4. Enter the following command to enable `libvirtd`:
 

```
sudo systemctl enable libvirtd
```

## Installing Virtual Machine Manager on your local computer

To install Virtual Machine Manager on your local computer, you must download the library, enable the necessary functions, and download the application on your local computer. To start Virtual Machine Manager:

### Procedure

1. Open a new terminal that is not using SSH to connect to your KVM host and enter the following command:
 

```
sudo dnf install virt-manager
```

2. After the packages are installed, enter the following command to start virtlogd:

```
sudo systemctl start virtlogd
```

3. Enter the following command to start libvirtd:

```
sudo systemctl start libvirtd
```

4. Enter the following command to enable libvirtd:

```
sudo systemctl enable libvirtd
```

---

## Chapter 4: Configuring KVM networking

This chapter covers configuring the KVM host network so that it is ready for HCP VM system deployment. The steps in this section must be performed individually on each KVM host.

### Configuring the KVM host for networking

To configure the KVM host for networking, you must create bond files that bond the two physical front-end ports together and bond the two physical back-end ports together. Then you must create two bridge files that make the bonded physical front-end ports and back-end ports accessible to HCP VM node virtual NICs. The procedure makes the HCP VM system accessible to outside networks.

For the examples in this chapters, the four KVM host, physical network ports are named eno1, eno2, eno3, and eno4. The example KVM host front-end IP address is 192.168.210.16. The example KVM host gateway address is 192.168.210.254. The example KVM host netmask address is 255.255.254.0.



**Note:** (Optional) Connect a monitor and keyboard to your KVM host.

### (Optional) Connecting a monitor and keyboard to your KVM host

#### Before you begin

To configure networking for a KVM host, you will need the following items:

- USB keyboard
- VGA monitor

Connect your monitor and keyboard to the KVM host.

### Backing up the interface configuration files

Before you make changes to the network, you need to back up the original KVM host interface configuration files. This example assumes that the interface configuration files are named `ifcfg-eno1`, `ifcfg-eno2`, `ifcfg-eno3`, and `ifcfg-eno4`.

#### Procedure

1. Enter the following command to open the network-scripts folder:

```
cd /etc/sysconfig/network-scripts/
```

2. Enter the following command to look inside the folder:

```
ls
```

3. Verify that the folder contains the interface configuration files.
4. Enter the following command to create a backup of the interface configuration files in your root folder:

```
cp ifcfg-eno* /root/
```

5. Enter the following command to list the contents of the root folder:

```
ls /root
```

6. Verify that the root folder contains the four copied interface configuration files.

## Enabling the bonding kernel module

Configure the bonding kernel module so that it is enabled every time the KVM host is started:

### Procedure

1. Enter the following command to enable bonding mode:

```
modprobe --first-time bonding
```

2. Enter the following command to create a configuration file in the `modules-load.d` folder:

```
vi /etc/modules-load.d/bonding.conf
```

3. Press **I** to edit the file.

4. Enter the following text in the file:

```
#Load the bonding kernel module at boot bonding
```

5. Press **Esc**.

6. Enter the following command to save and exit the file:

```
:wq
```

7. Restart the KVM host.

8. Enter the following command to verify that the file is working:

```
lsmod | grep bonding
```

The output should contain the following text:

```
bonding
```

## Creating the bridge network files

You must create two interface configuration bridge files, one for the front end and one for the back end, to connect the virtual HCP VM network to the KVM host network. In this example, the front-end bridge network file is named `ifcfg-front-end`, and the back-end bridge network file is named `ifcfg-back-end`.

### Procedure

1. Enter the following command to open the network-scripts folder:

```
cd /etc/sysconfig/network-scripts/
```

2. Enter the following command to create the front-end bond network file:

```
vi ifcfg-front-end
```

3. Press **I** to edit the file.

4. Enter the following text:

```
DEVICE="Device-Name"
ONBOOT="yes"
TYPE="Bridge"
BOOTPROTO="none"
IPADDR="Node-IP-Address"
NETMASK="Netmask-IP-Address"
GATEWAY="Gateway-IP-Address"
IPV6INIT="yes"
IPV6_AUTOCONF="no"
DHCPV6C="no"
STP="on"
DELAY="0.0"
```

Here is an example of the completed file:

```
DEVICE="front-end"
ONBOOT="yes"
TYPE="Bridge"
BOOTPROTO="none"
IPADDR="192.168.210.16"
NETMASK="255.255.254.0"
GATEWAY="192.168.210.254"
IPV6INIT="yes"
IPV6_AUTOCONF="no"
DHCPV6C="no"
STP="on"
DELAY="0.0"
```

5. Press **Esc**.
6. Enter the following command to save and exit the file:  
:wq
7. Enter the following command to create the back-end bond network file:  
vi ifcfg-bond1
8. Press **I** to edit the file.
9. Enter the following text:

```
DEVICE="Device-Name"
STP="on"
TYPE="Bridge"
BOOTPROTO="none"
IPV4_FAILURE_FATAL="no"
NAME="Bridge-Name"
```

```
ONBOOT="yes"
DELAY="0.0"
```

Here is an example of the completed file:

```
DEVICE="Back-End"
STP="on"
TYPE="Bridge"
BOOTPROTO="none"
IPV4_FAILURE_FATAL="no"
NAME="back-end"
ONBOOT="yes"
DELAY="0.0"
```

10. Press **Esc**.
11. Enter the following command to save and exit the file:  
:wq

## Creating the bond network files

You must create two interface configuration bond files to bond the two KVM host front-end physical NIC ports together and bond the two KVM host back-end physical NIC ports together. In this example, the front-end bond network file is named `ifcfg-bond0`, and the back-end bond network file is named `ifcfg-bond1`.

### Procedure

1. Enter the following command to open the network-scripts folder:  
`cd /etc/sysconfig/network-scripts/`
2. Enter the following command to create the front-end bond network file:  
`vi ifcfg-bond0`
3. Press **I** to edit the file.
4. Enter the following text:

```
DEVICE=Device-Name
NAME=Front-End-Bond-File-Name
TYPE=Bond
BONDING_MASTER=yes
BOOTPROTO=none
ONBOOT=yes
BONDING_OPTS="primary=Interface-Configuration-File1 \
mode=active-backup miimon=100 updelay=3000 downdelay=500"
BRIDGE="Front-End-Bridge-Device"
```

Here is an example of the completed file:

```
DEVICE=bond0
NAME=bond0
TYPE=Bond
```

```
BONDING_MASTER=yes
BOOTPROTO=none
ONBOOT=yes
BONDING_OPTS="primary=enol mode=active-backup miimon=100
updelay=3000 downdelay=500"
BRIDGE="front-end"
```

5. Press **Esc**.
6. Enter the following command to save and exit the file:  
:wq
7. Enter the following command to create the back-end bond network file:  
vi ifcfg-bond1
8. Press **I** to edit the file.
9. Enter the following text:

```
DEVICE=Device-Name
NAME=Back-End-Bond-File-Name
TYPE=Bond
BONDING_MASTER=yes
BOOTPROTO=none
ONBOOT=yes
BONDING_OPTS="primary=Interface-Configuration-File1 \
mode=active-backup miimon=100 updelay=3000 downdelay=500"
BRIDGE="Back-End-Bridge-Device"
```

Here is an example of the completed file:

```
DEVICE=bond1
NAME=bond1
TYPE=Bond
BONDING_MASTER=yes
BOOTPROTO=none
ONBOOT=yes
BONDING_OPTS="primary=enol mode=active-backup miimon=100
updelay=3000 downdelay=500"
BRIDGE="back-end"
```

10. Press **Esc**.
11. Enter the following command to save and exit the file:  
:wq

## Enslaving the front-end interface configuration files

To bridge the front-end network, you must enslave the front-end interface configuration files. In this example, the first front-end network configuration file is named `ifcfg-enol`, and the second front-end network configuration file is named `ifcfg-enol3`.





**Note:** Do not delete information that is specific to your system network configuration.

### Procedure

1. Enter the following command to open the network-scripts folder:  
`cd /etc/sysconfig/network-scripts/`
2. Enter the following command to access the first front-end network configuration file:  
`vi ifcfg-eno1`
3. Press **I** to edit the file.
4. Replace the existing contents with the following text:

```
NAME=Slave-Name
DEVICE=Device-Name
ONBOOT=yes
MASTER=Front-End-Bond-File
SLAVE=yes
```



**Note:** Do not delete information that is particular to your system network configuration.

Here is an example of the completed file:

```
NAME=bond0-slave1
DEVICE=eno1
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

5. Press **Esc**.
6. Enter the following command to save and exit the file:  
`:wq`
7. Enter the following command to access the second front-end network configuration file:  
`vi ifcfg-eno3`
8. Press **I** to edit the file.
9. Replace the existing contents with the following text:

```
NAME=Slave-Name
DEVICE=Device-Name
ONBOOT=yes
MASTER=Front-End-Bond-File
SLAVE=yes
```



**Note:** Do not delete information that is particular to your system network configuration.

Here is an example of the completed file:

```
NAME=bond0-slave2
DEVICE=eno3
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

10. Press **Esc**.
11. Enter the following command to save and exit the file:  
:wq

## Enslaving the back-end interface configuration files

To bridge the back-end network, you must enslave the back-end interface configuration files. In this example, the first back-end network configuration file is named `ifcfg-eno2`, and the second back-end network configuration file is named `ifcfg-eno4`.

To enslave the interface network configuration files:

### Procedure

1. Enter the following command to open the network-scripts folder:  
`cd /etc/sysconfig/network-scripts/`
2. Enter the following command to access the first back-end network configuration file:  
`vi ifcfg-eno2`
3. Press **I** to edit the file.
4. Replace the existing contents with the following text:

```
NAME=Slave-File-Name
DEVICE=Device-Name
ONBOOT=yes
MASTER=Back-End-Bond-File
SLAVE=yes
ETHTOOL_OPTS=""
```



**Note:** Do not delete information that is particular to your system network configuration.

Here is an example of the completed file:

```
NAME=bond1-slave1
DEVICE=eno2
ONBOOT=yes
MASTER=bond1
SLAVE=yes
ETHTOOL_OPTS=""
```

5. Press **Esc**.
6. Enter the following command to save and exit the file:  
:wq
7. Enter the following command to access the second back-end network configuration file:  
vi ifcfg-eno4
8. Press **I** to edit the file.
9. Replace the existing contents with the following text:

```
NAME=Slave-File-Name
DEVICE=Device-Name
ONBOOT=yes
MASTER=back-end-Bond-File
SLAVE=yes
ETHTOOL_OPTS=""
```



**Note:** Do not delete information that is particular to your system network configuration.

Here is an example of the completed file:

```
NAME=bond1-slave2
DEVICE=eno4
ONBOOT=yes
MASTER=bond1
SLAVE=yes
ETHTOOL_OPTS=""
```

10. Press **Esc**.
11. Enter the following command to save and exit the file:  
:wq

## Restarting networking on the KVM host network services

After the network files are created and edited, restart the KVM host network services by issuing the following command:

```
sudo systemctl restart network
```



**Note:** If you used SSH to perform the network configuration, you might experience connection issues after the network restarts.

## Restarting the individual KVM host

Restarting a KVM host causes it to restart. While the KVM host is in the process of restarting, you have no access to it.

**Procedure**

1. Press and hold the power button until the KVM host shuts down.
2. Press the power button again and wait for the KVM host to restart.  
The KVM host restarts. When the KVM host has finished restarting, it is available for access.



**Note:** If the KVM host remains unavailable after the restart, restart the KVM host network services by issuing the following command:

```
sudo systemctl restart network
```

## Verifying that the KVM host network has been updated

Once the KVM host has restarted and is operational, you need to verify that the network has been updated.

**Procedure**

1. Enter the following command to verify that the bridge and interface configuration files are working:

```
IP link
```

The output should contain the following text:

```
Interface-Configuration-File1: <BROADCAST, MULTICAST, SLAVE, UP,
LOWER_UP>
Interface-Configuration-File2: <BROADCAST, MULTICAST, SLAVE, UP,
LOWER_UP>
Interface-Configuration-File3: <BROADCAST, MULTICAST, SLAVE, UP,
LOWER_UP>
Interface-Configuration-File4: <BROADCAST, MULTICAST, SLAVE, UP,
LOWER_UP>
Front-End-Bridge-File: <BROADCAST, MULTICAST, UP, LOWER_UP>
Back-End-Bridge-File: <BROADCAST, MULTICAST, UP, LOWER_UP>
```

2. Enter the following command to verify that the bond0 file is operational:

```
cat /proc/net/bonding/bond0
```

The output should contain the following text:

```
Primary Slave: eno1 (primary_reselect always)
Currently Active Slave: Interface-Configuration-File1
MII Status: up
MII Polling Interval (ms): 100 Up Delay (ms): 3000
Down Delay (ms): 500
Slave Interface: Interface-Configuration-File3
MII Status: up Speed: 1000 Mbps Duplex: full
Link Failure Count: 0 Slave queue ID: 0
Slave Interface: Interface-Configuration-File1
MII Status: up Speed: 1000 Mbps Duplex: full
Link Failure Count: 0 Slave queue ID: 0
```

3. Enter the following command to verify that the bond1 file is operational:

```
cat /proc/net/bonding/bond1
```

The output should contain the following text:

```
Bonding Mode: fault-tolerance (active-backup)
Primary Slave: Interface-Configuration-File2 (primary_reselect always)
Currently Active Slave: Interface-Configuration-File2
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 3000
Down Delay (ms): 500

Slave Interface: Interface-Configuration-File2
MII Status: up
Speed: 1000
Mbps Duplex: full
Link Failure Count: 0
Slave queue ID: 0

Slave Interface: Interface-Configuration-File4
MII Status: up
Speed: 1000
Mbps Duplex: full
Link Failure Count: 0
Slave queue ID: 0
```

---

## Chapter 5: Maintaining the HCP VM system

This chapter describes how to keep your HCP VM system running at an optimal performance level.

### Adding logical volumes

#### Procedure

1. In the Virtual Machine Manager, open the virtual machine console for the highest numbered storage node.
2. Log in as the install user.  
The **HCP Configuration Menu** is displayed. The following window is included for illustrative purposes. Your screen will specify the HCP version that you are running.

```
HCP 9.0 Configuration Menu
=====

[1] Get HCP Setup Files
[2] Install an HCP System
[3] Upgrade an HCP System
[4] Add a Node to an HCP System
[5] Perform Checks for Offline Upgrade
[6] Perform Checks for Online Upgrade
[v] Add Logical Volumes to an HCP System
[s] Perform a Service Procedure
[q] Log Out

Currently installed version:  9.0
Version on CD/DVD:           None
Extracted version:           9.0

Enter a selection:
```

3. Enter `v` to add logical volumes to an HCP system.

```
Add New Storage
=====

[1] Add Storage to the HCP System while It Is Online

[q] Return to Configuration Menu
```

```
Enter your choice
[Default: 1]:
```

4. Enter 1 to add storage to the HCP system while it is online.

```
Add Storage to the HCP System while it is Online
=====
```

```
This option adds storage to a node. It should be used only by trained
and
Qualified personnel.
```

```
FOR SAIN SYSTEMS: Before you begin, make sure that someone present is
qualified and authorized to use the storage management application for
your
storage array.
```

```
WARNING (SAIN SYSTEMS): Be sure the new storage is properly configured
at the
storage tier before continuing this procedure. Trying to add improperly
configured storage can result in data loss or can cause the system to
become
inoperable.
```

```
Are you sure you still want to continue?
[Default: no]: yes
```

```
You chose "yes", is that correct?
[Default: yes]:
```

5. Enter `yes` to continue the procedure.
6. Press **Enter** to confirm.

After you have confirmed that you want to add storage, HCP Setup performs a set of installation prechecks.

```
Verifying correct menu
Verifying system name
Verifying run location
Verifying running as install
Verifying node connections
Verifying SSH keys
Verifying SSH
Verifying systemwide SSH
Verifying all network links
Verifying software versions
Verifying all nodes available
Verifying upgrade state
Verifying 64-bit hardware platform
Verifying drive size
```

```

Verifying disk space
Verifying nobody using /fcfs_*
Verifying nobody using /fs/*
Verifying storage tiering service is disabled
Searching for new storage volumes
Verifying multicast enabled
Found these new volumes:
    Node 001:
1.    /dev/sdd at 2:0:0:3 (500GB)
2.    /dev/sde at 2:0:0:4 (1TB)

    Node 002:
1.    /dev/sdd at 2:0:0:3 (500GB)
2.    /dev/sde at 2:0:0:4 (1TB)
    Node 03:
1.    /dev/sdd at 2:0:0:3 (500GB)
    Node 04:
1.    /dev/sdd at 2:0:0:3 (500GB)
Is this correct? [y/n]: y

```

7. Enter `y` to verify the new volumes that were found.

Typically this shows storage added to all nodes.

(Optional) To configure a dedicated database volume, the volume size must be at least 50 GB and at least 1.5 times the size of the existing database for each node.

All dedicated database volumes must be the same size. If you already have a dedicated database volume, any newly-added dedicated database volume needs to be larger than the current one.

8. If HCP Setup asks whether you want to select a dedicated volume for the database, perform one of the following:

- a. If you do not want to select a dedicated volume for the database, enter `no`. HCP Setup formats and adds the new volumes.

During this process, HCP Setup reports on its progress.

```

Syncing install password to all nodes.
Updating EULA
Syncing date to all nodes.
Syncing HCP package to all nodes
Starting to poll nodes for progress
Fri Mar  6 11:51:59 2020 Current status:
    node 150: 53% Complete (7/13): Running formatDrives
Fri Mar  6 11:52:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 27% complete
(0 / 1 volumes)
Fri Mar  6 11:53:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 35% complete
(0 / 1 volumes)
Fri Mar  6 11:54:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 42% complete
(0 / 1 volumes)

```



```

Fri Mar  6 11:55:15 2020 Current status:
      node 150: 53% Complete (7/13): Formatting 99% complete
(0 / 1 volumes)
Fri Mar  6 11:56:15 2020 Current status:
      node 150: 53% Complete (7/13): Formatting 100% complete
(1 / 1 volumes)
Fri Mar  6 11:56:25 2020 Current status:
      node 150: 76% Complete (10/13): Running
create_volume_config
Fri Mar  6 11:56:30 2020 Current status:
      node 150: 84% Complete (11/13): Running start_new_volumes
Fri Mar  6 11:56:46 2020 Current status:
      node 150: 92% Complete (12/13): Running
sync_new_local_volumes
Fri Mar  6 11:57:01 2020 Current status:
      node 150: 100% Complete: Storage addition complete
Fri Mar  6 11:57:22 2020 Current status:
      node 150: 100% Complete: Starting new volumes
Fri Mar  6 11:57:27 2020 Current status:
      node 150: 100% Complete: Starting new volumes

>>> HCP Logical Volume Addition completed successfully
Press ENTER to continue:

```

- b. To select a dedicated volume for the database, enter *yes*.

Then complete the next substeps:

- i. To select a dedicated database volume for the first node, enter *yes*.
  - ii. If you entered *yes*, select the dedicated database volume for the first node.
  - iii. Press **Enter** to confirm your selection.
  - iv. Repeat the three previous substeps for each node in the system. 9.
9. After you have selected the dedicated database volumes for each node, HCP Setup confirms your selections then asks if you want to continue the procedure. Enter *yes* to continue the procedure.

```

Do you want to select a dedicated volume for database? [Default: no]:
yes
Do you want to select a new dedicated PG LUN for node 001? [Default:
no]: yes
Enter a selection for node 001 [1, 2]: 2
You chose: "2. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]: yes
Do you want to select a new dedicated PG LUN for node 002? [Default:
no]: yes
Enter a selection for node 002: 2
You chose: "2. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]: yes
Do you want to select a new dedicated PG LUN for node 003? [Default:

```

```

no]: yes
Do you want to select a new dedicated PG LUN for node 004? [Default:
no]: yes
This will add the new volumes and move database to the following
dedicated
volumes. Do you want to continue?
    node 001: 2:0:0:4 (1TB)
    node 002: 2:0:0:4 (1TB)
    node 003: 2:0:0:3 (500GB)
    node 004: 2:0:0:3 (500GB)
[Default: no]: yes

```

HCP Setup formats and adds the new volumes. During this process, HCP Setup reports on its progress.

```

Syncing install password to all nodes.
Updating EULA
Syncing date to all nodes.
Syncing HCP package to all nodes
Starting to poll nodes for progress
Fri Mar  6 11:51:59 2020 Current status:
    node 150: 53% Complete (7/13): Running formatDrives
Fri Mar  6 11:52:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 27% complete (0 / 1
volumes)
Fri Mar  6 11:53:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 35% complete (0 / 1
volumes)
Fri Mar  6 11:54:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 42% complete (0 / 1
volumes)
Fri Mar  6 11:55:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 99% complete (0 / 1
volumes)
Fri Mar  6 11:56:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 100% complete (1 / 1
volumes)
Fri Mar  6 11:56:25 2020 Current status:
    node 150: 76% Complete (10/13): Running create_volume_config
Fri Mar  6 11:56:30 2020 Current status:
    node 150: 84% Complete (11/13): Running start_new_volumes
Fri Mar  6 11:56:46 2020 Current status:
    node 150: 92% Complete (12/13): Running sync_new_local_volumes
Fri Mar  6 11:57:01 2020 Current status:
    node 150: 100% Complete: Storage addition complete
Fri Mar  6 11:57:22 2020 Current status:
    node 150: 100% Complete: Starting new volumes
Fri Mar  6 11:57:27 2020 Current status:
    node 150: 100% Complete: Starting new volumes

```

```
>>> HCP Logical Volume Addition completed successfully
Press ENTER to continue:
```

10. 10. When the formatting is complete, press **Enter** to continue.
11. Log in to the HCP System Management Console to verify the newly added volumes.

## Moving storage node databases to optimal volumes

### Procedure

1. From the **HCP Configuration Menu**, enter `s` to display the HCP Service menu.
2. In response to the confirming prompt, enter `y` or `yes` to confirm your entry or `n` or `no` to try again.

When you enter `y` or `yes`, the **HCP Service Menu** appears.

```
HCP Service Menu
=====

[1] Recover an HCP Node
[2] Display the Current System Status
[3] Shut down the HCP System
[4] Free Space on Nodes
[5] Install a Hotfix
[6] Finalize Migration
[7] Change Region Count
[8] Configure Routine Database Maintenance
[9] Configure Time Settings
[10] Migrate Metadata Regions
[11] Perform Node Swap
[s] Perform SSD Maintenance
[m] Manage Database Volumes

[q] Return to Configuration Menu

Enter your choice.
[Default: 1]: m
```

3. From the **HCP Service Menu**, enter `m`.
4. In response to the confirming prompt, enter `y` or `yes` to confirm your entry or `n` or `no` to try again.

When you enter `y` or `yes`, the Setup displays the **Manage Database Volumes menu**.

```
Manage Database Volumes
=====

[1] Move Database
```

```
[d] Delete Old Database

[q] Return to Configuration Menu

[Enter your choice.
[Default: 1]:
```

5. From the Manage Database Volumes menu, enter 1.
6. In response to the confirming prompt, enter `y` or `yes` to confirm your entry or `n` or `no` to try again.

When you enter `y` or `yes`, HCP Setup displays the **Move Database** menu.

```
Move Database
=====

Volumes to move:

Node          | Type | Current volume | New volume
              |      | (Available/Total) | (Available/Total)
-----
172.20.59.125 | pgdata | /RIS/archive33 | /RIS/archive94
              | pgidx | (450M/1006)    | (1806/2006)
172.20.59.125 | pgxlog | /RIS/archive34 | /RIS/archive95
              |      | (450M/1006)    | (1806/2006)
172.20.59.129 | pgxlog | /RIS/archive33 | /RIS/archive94
              |      | (350M/1006)    | (1506/2006)

Executing this procedure will move the database from the current
volume to the new volume. This process cannot be undone after it is
complete.
Please review the changes before proceeding.

Do you want to move the database?
[Default: no]: yes

You chose: "yes", is this correct?
[Default: yes]: yes
```

From the **Move Database** menu, the HCP setup asks you to review your database configuration and warns you that the process cannot be undone after the database move is complete.

7. When you have reviewed the configuration, enter `y` or `yes` to confirm your entry or `n` or `no` to try again.
8. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

After the procedure begins, the progress of the database move appears and details the current status of the **HCP Service** procedure.

```
Starting to poll nodes for progress
Thu Jan 16 09:51:15 2020 Current status:
    node 042: 40% Complete (2/5): Running arcShutdown
Thu Jan 16 09:52:13 2020 Current status:
    node 042: 60% Complete (3/5): Running mountDisks
Thu Jan 16 09:52:48 2020 Current status:
    node 042: 80% Complete (4/5): Running move_pgdata
Thu Jan 16 09:52:55 2020 Current status:
    node 042: 100% Complete: Deploy complete
Thu Jan 16 09:54:07 2020 Current status:
    node 042: 100% Complete: Rebooting node
Thu Jan 16 09:54:12 2020 Current status:
    node 042: 100% Complete: Waiting for node to become available.
Thu Jan 16 09:56:12 2020 Current status:
    node 042: 100% Complete: Waiting for node to become available.
Thu Jan 16 09:58:13 2020 Current status:
    node 042: 100% Complete: Waiting for node to become available.
Thu Jan 16 10:00:03 2020 Current status:
    node 042: 100% Complete: All nodes available
Thu Jan 16 10:00:38 2020 Current status:
    node 042: 100% Complete: All nodes available and metadata is
balanced

>>> HCP Service Procedure successful
Press ENTER to continue:
```

When the procedure is complete, press **Enter** to return to the **HCP Service** menu. You can now delete the database from the older database volume.

## Deleting databases from older database volumes

### Procedure

1. From the **HCP Configuration Menu**, enter **s** to display the **HCP Service Menu**.
2. In response to the confirming prompt, enter **y** or **yes** to confirm the move or **n** or **no** to try again.

When you enter **y** or **yes**, the **HCP Service Menu** is displayed.

```
HCP Service Menu
=====

[1] Recover an HCP Node
[2] Display the Current System Status
[3] Shut down the HCP System
```

```

[4] Free Space on Nodes
[5] Install a Hotfix
[6] Finalize Migration
[7] Change Region Count
[8] Configure Routine Database Maintenance
[9] Configure Time Settings
[10] Migrate Metadata Regions
[11] Perform Node Swap
[s] Perform SSD Maintenance
[m] Manage Database Volumes

[q] Return to Configuration Menu

Enter your choice.
[Default: 1]: m

```

3. From the **HCP Service Menu**, enter m.
4. In response to the confirming prompt, enter y or yes to confirm the move or n or no to try again.  
When you enter y or yes, HCP Setup displays the **Manage Database Volumes** menu.

```

Manage Database Volumes
=====

[1] Delete Old Database

[q] Return to Configuration Menu

[Enter your choice.
[Default: d]:

```

5. From the **Manage Database Volumes** menu, enter d.  
You can delete the database from the older database volume only if you have completed the database move procedure.
6. In response to the confirming prompt, enter y or yes to confirm the move or n or no to try again.  
When you enter y or yes, HCP Setup displays the **Delete Old Database** menu.

```

Delete Old Database
=====

WARNING: This procedure deletes the old HCP database from its original
storage volumes.
The database on the optimal storage volumes will be preserved.

Do you want to continue? Yes or No.
[Default: no]: yes

Deleting the old database: #

```

```
The old database has been deleted. Press ENTER to continue:
```

7. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

After the procedure is complete, press **Enter** to return to the HCP **Service Menu**.

## Adding HCP VM nodes

### Before you begin

1. Add new KVM hosts or find existing KVM hosts that can support an HCP node. See *Installing KVM*??.
2. Unpack and upload the ISO files to the selected KVM hosts. See *Deploying the HCP VM system*??.
3. Create the new virtual machine. See *Deploying the HCP VM system*??.
4. Configure the HCP VM network on the newly deployed HCP VM nodes. See *Performing the OS installation*??.
5. From the highest active HCP VM node, run the Add Node service procedure. For more information, see the manual *Installing and Maintaining an HCP System*.

### Procedure

1. From the **Configuration Menu**, enter 4 to run the HCP Setup wizard.
2. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

When you enter `y` or `yes`, the **Membership Update Menu** is displayed.

```
HCP Setup: Membership Update Menu
```

```
=====
```

```
[1] Add Storage Nodes to the System (no updates)
```

```
[v] Review Updated Configuration (disabled, no updates)
```

```
[x] Add Nodes to an Existing HCP System (disabled, no updates)
```

```
[q] Return to Configuration menu
```

```
Enter your choice.
```

```
[Default: 1]:
```

3. From the **Membership Update Menu**, enter `x` to perform the node addition. The wizard displays an explanation of the node addition procedure.

```
Add Nodes to an Existing HCP System
```

```
=====
```

This option will erase all data on the new nodes, install the HCP software on those nodes, and add the nodes to the system configuration.

Note: Control-C cancels input.

Enter yes or no.

[Default: no]:

4. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

The wizard prompts again for confirmation.

5. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

After you confirm that you want to add nodes, the wizard downloads and displays the current system configuration.

Configuration confirmation.

=====

```
DNS Server (s) = 172.18.4.4S
Allow Data at Rest Encryption = No
Customer Support Contact Information = United States:
(800) 446-0744. Outside the United states: (858) 547-4526
Storage configuration = internal
Gateway Router IPv4 Address = 172.20.59.254
Encrypt Data at Rest on Primary Storage = No
Time Settings Compliance Mode = No
HCP System Serial Number = 00001
MOE Index-only volumes = No
Enable DNS = Yes
Chassis = None
Enable Replication on This System = Yes
Multicast Network = 238.172.59.42
Time Zone = America/New_York
Current Date and Time = None
Domain Name for the System = hcp.example.com
Allow Data in Flight Encryption / SSL = Yes
Blade Servers = No
Distributor/OEM Key Access = Arizona
Time Server(s) = internal
Gateway Router Secondary IPv6 Address = None
Gateway Router IPv6 Address = None
Spindown Volumes = No
HCP Storage Nodes: 1
172.59.42.5
Configure Dedicated Database Volumes = Yes

Use SHIFT+PGUP to review the Configuration.
```



```
IS this Configuration Correct?
(Default: no) . yes
```

6. Review the configuration and take one of the following actions.
  - If the configuration is incorrect:
    - a. Enter `n` or `no`.
    - b. In response to the confirming prompt, enter `y` or `yes`.
    - c. Exit the wizard and contact your HCP support center for help.
  - If the configuration is correct:
    - a. Enter `y` or `yes`.
    - b. In response to the confirming prompt, enter `y` or `yes`. HCP Setup performs a set of prechecks.

```
You chose "yes", is this correct?
[Default: yes]:
Verifying system name
Verifying run location
Verifying running as install
Verifying node connections
Verifying SSH keys
Verifying SSH
Verifying systemwide SSH
Verifying total memory > 32GB
Verifying all network links
Verifying software versions
Verifying 64-bit hardware platform
Verifying drive size
Verifying disk space
Verifying nobody using /fcfs_*
Verifying nobody using /fs/*
Verifying multicast enabled
Syncing install password to all nodes.
Updating EULA
Syncing timezone to all nodes
Syncing date to all nodes.
Generating auth keys
Generating system UUID
Syncing HCP package to all nodes
Checking to see if we need to run update schemaupgrade scripts
False
Updating schema scripts for upgrade.
```

If your current HCP system has dedicated database volumes, each newly added node must have at least three volumes. Also, the dedicated database volume size needs to be at least 50 GB. All dedicated database volumes need to be the same size.

7. Select dedicated database volumes.

- a. When prompted, select the dedicated database volume for the first node.
- b. Press **Enter** to confirm your selection.
- c. Repeat the previous two substeps for each node in the system.

After you select the dedicated database volumes for each node, HCP Setup confirms your selections. Then the wizard asks if you want to continue the node addition.

```
Select dedicated volume for each node.
Found these volumes:
    node 001:
        1. /dev/sdd at 2:0:0:1 (500GB)
        2. /dev/sde at 2:0:0:2 (500GB)
        3. /dev/sdd at 2:0:0:3 (500GB)
        4. /dev/sde at 2:0:0:4 (1TB)

    node 002:
        1. /dev/sdd at 2:0:0:1 (500GB)
        2. /dev/sde at 2:0:0:2 (500GB)
        3. /dev/sdd at 2:0:0:3 (500GB)
        4. /dev/sde at 2:0:0:4 (1TB)

    node 003:
        1. /dev/sdd at 2:0:0:1 (500GB)
        2. /dev/sde at 2:0:0:2 (500GB)
        3. /dev/sdd at 2:0:0:3 (500GB)
        4. /dev/sde at 2:0:0:4 (1TB)

    node 004:
        1. /dev/sdd at 2:0:0:1 (500GB)
        2. /dev/sde at 2:0:0:2 (500GB)
        3. /dev/sdd at 2:0:0:3 (500GB)
        4. /dev/sde at 2:0:0:4 (1TB)

Select dedicated database volume for node 001: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]:
Select dedicated database volume for node 002: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]:
Select dedicated database volume for node 003: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]:
Select dedicated database volume for node 004: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]:
Following volumes will be configured as dedicated database volumes:
    node 001: 4. /dev/sde at 2:0:0:4 (1TB)
    node 002: 4. /dev/sde at 2:0:0:4 (1TB)
    node 003: 4. /dev/sde at 2:0:0:4 (1TB)
    node 004: 4. /dev/sde at 2:0:0:4 (1TB)
Do you want to continue? [Default: yes]?
```

8. Press **Enter** to continue the procedure.
  - If the prechecks are successful, the HCP software is installed on the new nodes. For RAIN and VM systems, the software is installed on four nodes at a time. For SAIN systems, the software is installed on one cross-mapped pair of nodes at a time. After the software installation is complete, the nodes are restarted. When the node addition is complete, the **HCP Configuration Menu** redisplay.
  - If any of the prechecks fail, HCP Setup exits. In this case, fix the problem and then start the node addition procedure again.
  - If HCP Setup exits at any time before the node addition processing is complete, contact your HCP support center for help.
9. From the **HCP Configuration Menu**, enter `q` to log out of the install shell.
10. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

## Recovering storage nodes

The following sections describe how to recover storage nodes when preserving storage volumes, and how to recover storage nodes when clearing storage volumes. For more information about these and other procedures, see the manual *Installing and Maintaining an HCP System*.

## Recovering storage nodes and preserving volumes

### Procedure

1. From the **HCP Configuration Menu**, enter `s` to display the **HCP Service Menu**.
2. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

When you enter `y` or `yes`, the **HCP Service Menu** appears.

```
HCP Service Menu
=====

[1] Recover an HCP Node
[2] Display the Current System Status
[3] Shut down the HCP System
[4] Free Space on Nodes
[5] Install a Hotfix
[6] Finalize Migration
[7] Change Region Count
[8] Configure Routine Database Maintenance
[9] Configure Time Settings
[10] Migrate Metadata Regions
[11] Perform Node Swap
[s] Perform SSD Maintenance
[m] Manage Database Volumes
```

```
[q] Return to Configuration Menu
```

```
Enter your choice.
```

```
[Default: 1]: m
```

3. From the **HCP Service Menu**, enter 1 for recovery.
4. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

When you enter `y` or `yes`, HCP Setup displays the **Node Recovery Menu**.

```
HCP Setup: Node Recovery Menu
```

```
=====
```

```
[1] Recover storage nodes
[2] Recover all storage nodes
[3] Reinitialize internal database
```

```
[b] Go Back to the Previous Menu
[q] Return to the Configuration Menu
```

```
Enter your choice.
```

```
[Default: 1]:
```

5. From the **Node Recovery Menu**, take one of these actions:
  - To recover selected storage nodes, enter 1. Then follow the on-screen instructions to identify the nodes you want to recover. Make sure to use the back-end IP address to identify each node.



**Note:** If you choose to use a range of IP addresses to identify the nodes, ensure that the range you specify includes only the nodes you want to recover.

- If you identify fewer than half of the nodes in the HCP system, HCP Setup asks whether you want to delete and then try to rebuild the database on those nodes.

```
Do you want to delete the database and have HCP try to rebuild it?
Enter yes only if you know that the database is unrecoverable. If
you
are unsure, enter no.
[Default: no]:
```

- a. Enter `y` or `yes` to delete the database while recovering the OS or `n` or `no` to recover the OS without deleting the database.
- b. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.
- c. (Optional) If you are performing the OS recovery on an HCP system with

dedicated database volumes, HCP Setup displays the next prompt.

```
Do you want HCP to format the internal database drive before
rebuilding it?
If you enter yes, only the internal database drive is formatted.
Enter yes only if you know that the internal database drive
needs formatting.
If you are unsure, enter no.
[Default: no]:
```

- To format only the internal database drive, enter `y` or `yes`.
- To keep the internal database drive in its original state, enter `n` or `no`.
- If you identify half or more of the nodes in the HCP system, HCP Setup displays a unique key and prompts you to enter it.
- To recover all storage nodes, enter `2`.

HCP Setup displays a unique key and prompts you to enter it back.

6. Enter the unique key exactly as it is shown.

HCP Setup performs a series of prechecks and, if they are successful, recovers the OS on the selected nodes or all nodes, as applicable.

If any of the prechecks fail, HCP Setup exits.

If that happens, fix the problem and then start the OS recovery procedure again.

When the node recovery is complete, HCP restarts all the nodes that it recovered and displays this message:

```
>>> HCP Service Procedure successful Press ENTER to continue:
```

7. If the node that you are logged in to is one of the recovered nodes, the SSH or console session ends automatically when HCP restarts the node.

If this is not the case, in response to the prompt to continue, press **Enter**.

The HCP Service Menu is displayed.



**Important:** If HCP Setup exits at any time before the OS recovery processing is complete, contact your HCP support center. Do not try to recover the OS again.

8. From the **HCP Service Menu**, enter `q` to return to the **HCP Configuration Menu**.
9. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.
10. From the **HCP Configuration Menu**, enter `q` to log out of the install shell.
11. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

## Recovering storage nodes and clearing volumes

### Procedure

1. From the **HCP Configuration Menu**, enter `s` to display the **HCP Service Menu**.

2. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

If you enter `y` or `yes`, the **HCP Service Menu** is displayed.

```
HCP Service Menu
=====

[1] Recover an HCP Node
[2] Display the Current System Status
[3] Shut down the HCP System
[4] Free Space on Nodes
[5] Install a Hotfix
[6] Finalize Migration
[7] Change Region Count
[8] Configure Routine Database Maintenance
[9] Configure Time Settings
[10] Migrate Metadata Regions
[11] Perform Node Swap
[s] Perform SSD Maintenance
[m] Manage Database Volumes

[q] Return to Configuration Menu

Enter your choice.
[Default: 1]: m
```

3. Enter `1` for recovery.
4. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

If you enter `y` or `yes`, HCP Setup displays the **Node Recovery Menu**.

```
HCP Setup: Node Recovery Menu
=====

[1] Recover storage nodes
[2] Recover all storage nodes
[3] Reinitialize internal database

[b] Go Back to the Previous Menu
[q] Return to the Configuration Menu

Enter your choice.
[Default: 1]:
```

5. From the **Node Recovery Menu**, enter `1` to recover selected storage nodes. Then follow the instructions to identify the nodes you want to recover. Make sure to use the back-end IP address to identify each node.



**Note:** If you choose to use a range of IP addresses to identify the nodes, make sure that the range you specify includes only the nodes you want to recover.

6. (Optional) If you set **FORCE\_FORMAT** to 1, when you enter **y or yes**, HCP Setup displays the `FORCE_FORMAT` prompt.

```
Enabling FORCE FORMAT will format all disks. Are you sure you want to
do
this?
[Default: no]:
```



**Important:** If you receive this prompt, continuing with this procedure formats all disks and erases all data on the targeted node or nodes. The data cannot be recovered. Perform this action only if you are sure the data can be deleted.

7. Enter `y` or `yes` to allow the system to format all disks.  
Then follow the on-screen instructions to identify the node containing the logical volumes you want to recover.  
HCP Setup displays a unique key and prompts you to enter it back.
8. Enter the unique key exactly as it is shown.  
After you have entered the unique key, HCP Setup performs a set of prechecks.

```
You chose "yes", is this correct?
[Default: yes]:
Verifying system name
Verifying run location
Verifying running as install
Verifying node connections
Verifying SSH keys
Verifying SSH
Verifying systemwide SSH
Verifying total memory > 32GB
Verifying all network links
Verifying software versions
Verifying 64-bit hardware platform
Verifying drive size
Verifying disk space
Verifying nobody using /fcfs_*
Verifying nobody using /fs/*
Verifying multicast enabled
Syncing install password to all nodes.
Updating EULA
Syncing timezone to all nodes
Syncing date to all nodes.
Generating auth keys
Generating system UUID
Syncing HCP package to all nodes
```

```
Checking to see if we need to run update schemaupgrade scripts False
Updating schema scripts for upgrade.
```

9. (Optional) Complete the next substeps only if your current HCP system has dedicated database volumes.
  - a. When prompted, select the dedicated database volume for the first node.
  - b. Press **Enter** to confirm your selection.
  - c. Repeat the two previous substeps for each node in the system.

After you have selected the dedicated database volumes for each node, HCP Setup confirms your selections then asks if you want to continue the node recovery.

```
Select dedicated volume for each node.
Found these volumes:
    node 001:
        1. /dev/sdd at 2:0:0:1 (500GB)
        2. /dev/sde at 2:0:0:2 (500GB)
        3. /dev/sdd at 2:0:0:3 (500GB)
        4. /dev/sde at 2:0:0:4 (1TB)

    node 002:
        1. /dev/sdd at 2:0:0:1 (500GB)
        2. /dev/sde at 2:0:0:2 (500GB)
        3. /dev/sdd at 2:0:0:3 (500GB)
        4. /dev/sde at 2:0:0:4 (1TB)

    node 003:
        1. /dev/sdd at 2:0:0:1 (500GB)
        2. /dev/sde at 2:0:0:2 (500GB)
        3. /dev/sdd at 2:0:0:3 (500GB)
        4. /dev/sde at 2:0:0:4 (1TB)

    node 004:
        1. /dev/sdd at 2:0:0:1 (500GB)
        2. /dev/sde at 2:0:0:2 (500GB)
        3. /dev/sdd at 2:0:0:3 (500GB)
        4. /dev/sde at 2:0:0:4 (1TB)

Select dedicated database volume for node 001: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]:
Select dedicated database volume for node 002: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]:
Select dedicated database volume for node 003: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]:
Select dedicated database volume for node 004: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default:
yes]:
Following volumes will be configured as dedicated database volumes:
    node 001: 4. /dev/sde at 2:0:0:4 (1TB)
```



```
node 002: 4. /dev/sde at 2:0:0:4 (1TB)
node 003: 4. /dev/sde at 2:0:0:4 (1TB)
node 004: 4. /dev/sde at 2:0:0:4 (1TB)
Do you want to continue? [Default: yes]?
```

10. Press **Enter** to continue the procedure.

If the prechecks are successful, HCP Setup recovers all of the logical volumes on the selected nodes or all nodes, as applicable.

If any of the prechecks fail, HCP Setup exits.

If that happens, fix the problem and then start the node recovery procedure again.

When the node recovery is complete, the **HCP Service Menu** is displayed.



**Note:** If HCP Setup exits before the node recovery processing is complete, contact your HCP support center. Do not try the node recovery again.

11. From the **HCP Service Menu**, enter `q` to return to the HCP Configuration Menu.
12. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.
13. From the **HCP Configuration Menu**, enter `q` to log out of the install shell.
14. In response to the confirming prompt, enter `y` or `yes` to confirm the move or `n` or `no` to try again.

---

## Chapter 6: Deploying the HCP VM system

This section covers how to use a local Linux machine and the Virtual Machine Manager application to create and configure an HCP VM system.

### Prerequisites

Before you deploy an HCP VM node, you need to:

- Install Virtual Machine Manager on your local machine
- Update your local machine Linux OS to either match or be a later version of the OS running on your KVM host

### Deploying the HCP VMsystem

To deploy the HCP VM system, you must deploy an HCP VM node on each KVM host. The following instructions explain how to deploy a single HCP VM node on a single KVM host and need to be repeated for each KVM host.

After the HCP VM connections have been added to Virtual Machine Manager, you can create the HCP VM nodes using the Virtual Machine Manager application or the command line. The instructions in this section explain how to deploy an HCP VM node using Virtual Machine Manager. For more information about using the command line to deploy an HCP VM, see [Creating an HCP VM node using the command line \(on page 95\)](#).

Before you deploy an HCP VM system:

- Install the KVM packages (see [Installing the KVM packages \(on page 17\)](#).)
- Create network bridges (see [Creating the bridge network files \(on page 21\)](#).)
- Copy and unzip a `.iso` file on your HCP node (see [Copying the .iso file and sending the Zip file to the KVM host \(on page 51\)](#).)

### Downloading the .iso installation files

Go to the HCP distributor download site and download the HCP software installation file (`HS222_release-number.iso.zip`).

## Copying the .iso file and sending the Zip file to the KVM host

### Procedure

1. Enter the following command to navigate to the folder on your local computer where you saved `HS222_release-number.iso.zip`:  

```
cd $(find / -name "HS222_release-number.iso.zip" | xargs  
dirname )
```
2. Enter the following command to copy and send the file to the `/var/iso` folder on your KVM host:  

```
scp HS222_release-number.iso.zip Username@Front-End-Node-  
IPAddress:/var/iso For example: scp HS222_8.0.0.824.iso.zip  
root@192.168.210.16:/var/iso
```
3. SSH into your KVM host.
4. From your KVM host, enter the following command to navigate to the folder with the saved .iso file:  

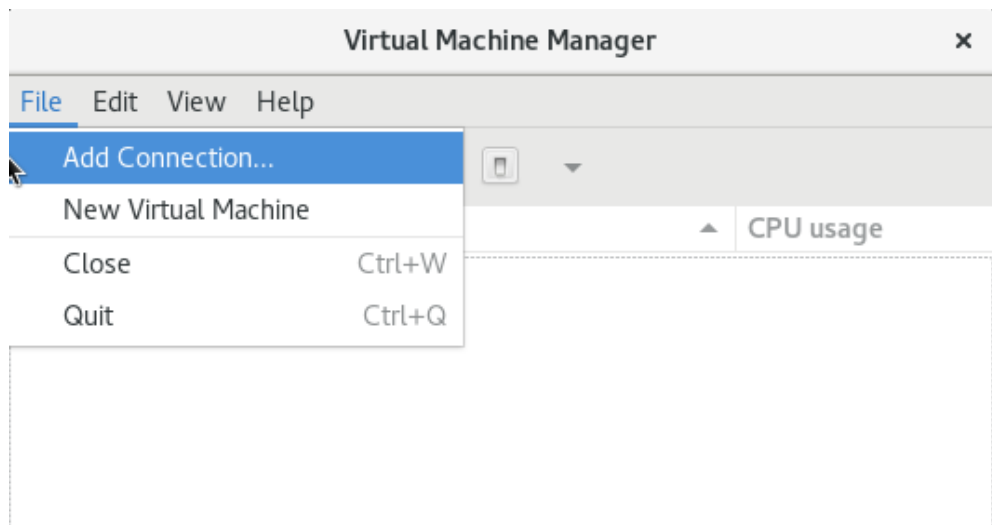
```
cd /var/iso
```
5. Enter the following command to unzip the file:  

```
unzip HS222_release-number.iso.zip
```

## Adding an HCP VM node connection

### Procedure

1. Open Virtual Machine Manager.
2. In Virtual Machine Manager, click **File > Add Connection**



3. To configure the **Add Connection** window:

**Add Connection** [X]

Hypervisor: QEMU/KVM ▼

☒ Connect to remote host

Method: SSH ▼

Username: root

Hostname: 172.21.150.153 ▼

Autoconnect: ☒

Generated URI: qemu+ssh://root@172.21.15...

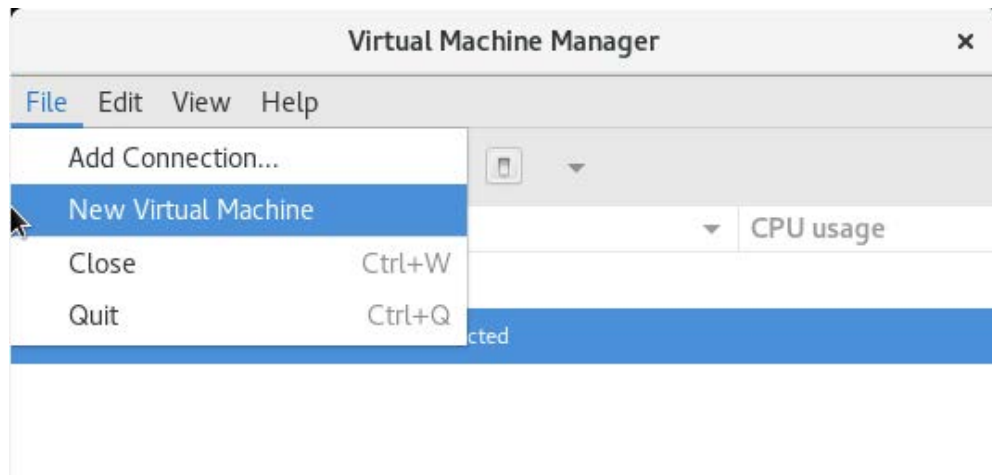
Cancel Connect

- a. In the **Hypervisor** menu, select **QEMU/KVM**.
  - b. Select **Connect** to a remote host.
  - c. In the **Method** menu, select **SSH**.
  - d. In the **Username** field, type the username for the KVM host.
  - e. In the **Hostname** field, type the hostname of the KVM host or its front-end IP address.
  - f. Select **Autoconnect**.
4. Click **Connect**.
  5. In the **OpenSSH window**, enter the password for your KVM host.
  6. Click **OK**.

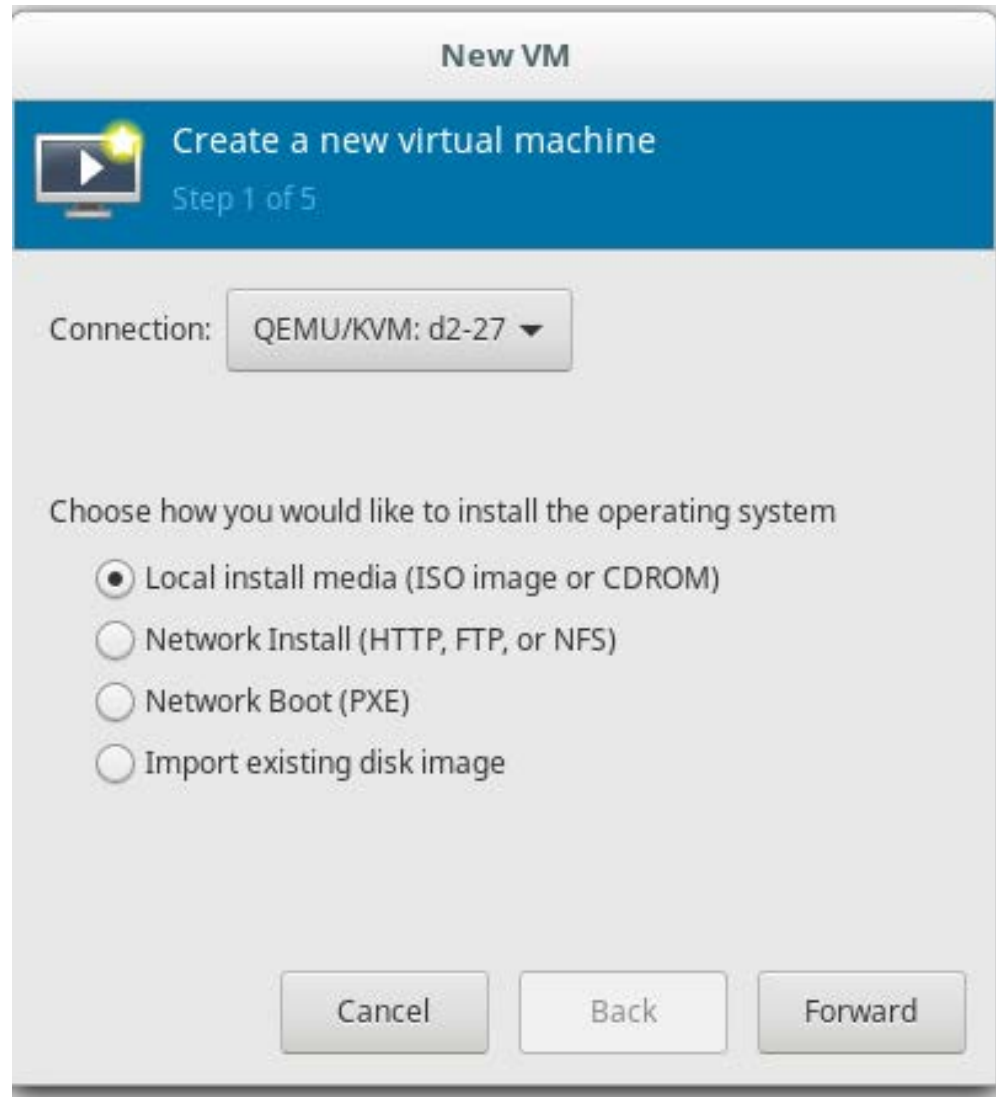
## Creating the HCP VM node

### Procedure

1. In Virtual Machine Manager, highlight the newly added connection, then click **File > New Virtual Machine**.
2. In the menu that opens, click **New Virtual Machine**.



3. To configure the **New VM Step 1 of 5** window:
- In the **Connection** field, select the host on which you are deploying the KVM.
  - Select **Local install media**.



4. Click **Forward**.
5. To configure the **New VM Step 2 of 5** window:
  - a. Select **Use ISO image**, then click **Browse**.
  - b. Navigate to the `HS222_release-number.iso` file on the KVM host, then click **OK**.
  - c. In the **OS type** field, select **Linux**.
  - d. In the **Version** field, select **Fedora 25**.

**New VM**

Create a new virtual machine  
Step 2 of 5

Locate your install media

☐ Use CDROM or DVD

No media detected (/dev/sr0) ▼

☒ Use ISO image:

/var/iso/HS222\_8.0.0.682.iso ▼ Browse...

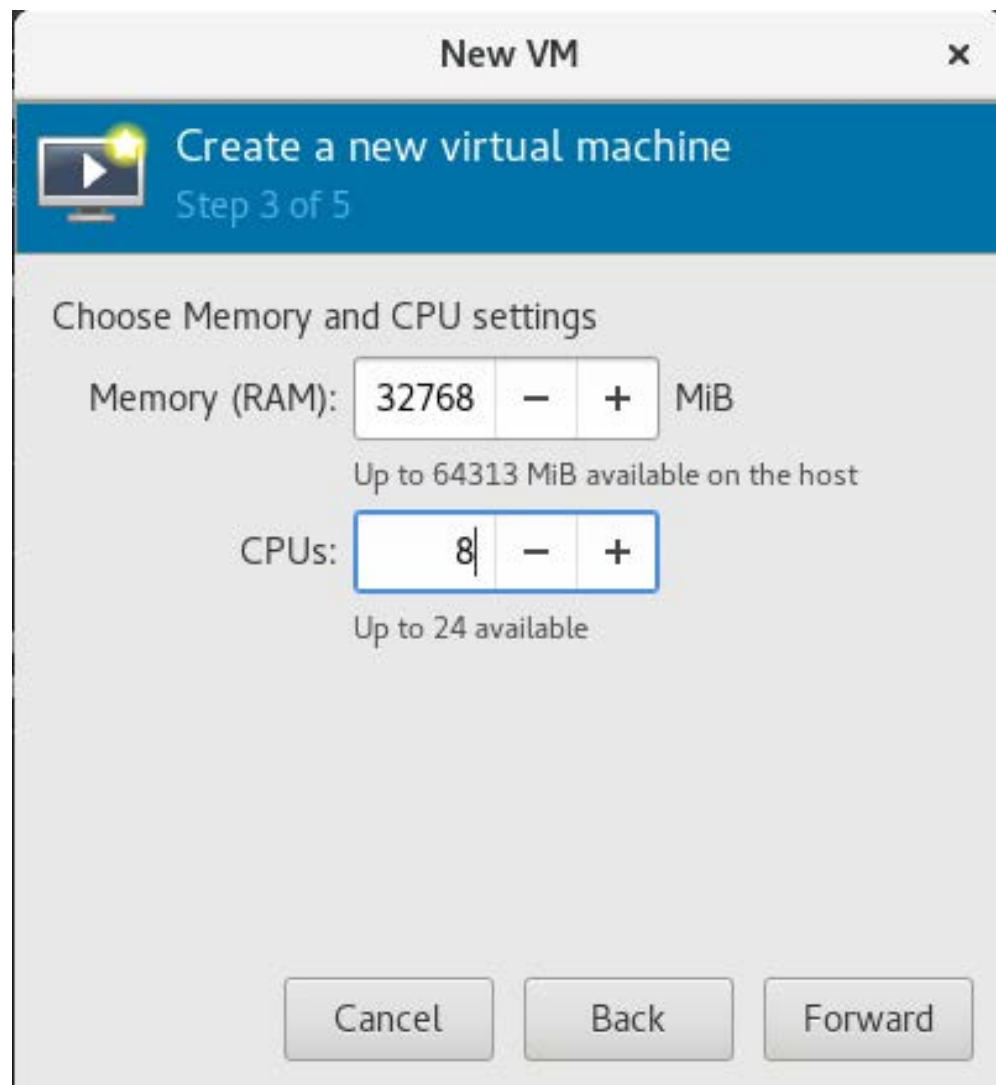
Choose an operating system type and version

OS type: Linux ▼

Version: Fedora 25 ▼

Cancel Back Forward

6. Click **Forward**.
7. To configure the **New VM Step 3 of 5** window:
  - a. In the **Memory (RAM)** field, type 32768 (32GB) for a standard HCP VM configuration, or type 16384 (16GB) for a small-instance HCP VM configuration.
  - b. In the **CPUs** field, type 8 for a standard HCP VM configuration, or type 4 for a small-instance HCP VM configuration.



8. Click **Forward**.
9. To configure the **New VM Step 4 of 5** window, perform either of these procedures:
  - If you are creating a disk image on the virtual machine:
    - a. Select **Enable storage for this virtual machine**.
    - b. Select **Create a disk image for the virtual machine**.
    - c. In the **Create a disk image for the virtual machine field**, enter 32.

**New VM**

Create a new virtual machine  
Step 4 of 5

☒ Enable storage for this virtual machine

☒ Create a disk image for the virtual machine

32|0 — + GiB

1134.0 GiB available in the default location

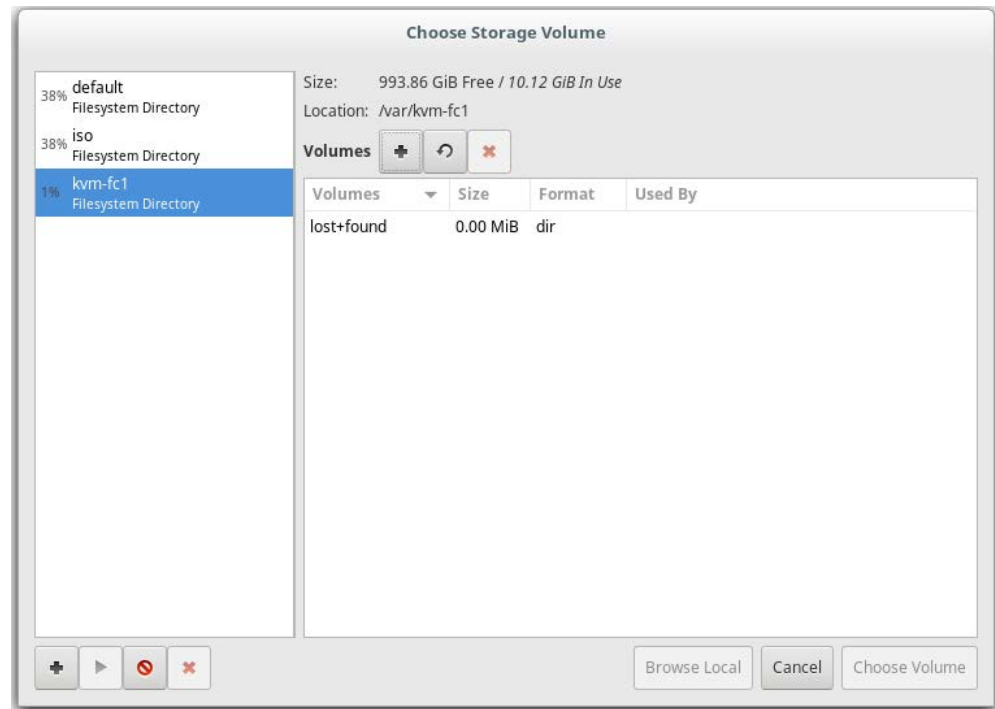
☐ Select or create custom storage

Manage...

Cancel Back Forward

- If you are using SAN storage and want to create a custom storage location:
  - a. Select **Select or create custom storage**.
  - b. Click **Manage**.
  - c. In the **Choose Storage Volume window** that opens, select the storage pool you want to host your storage volume.
  - d. Click the **Volumes:** plus sign.





- i. To configure the **Add a Storage Volume** window that opens:
  - 1) In the **Name** field, type a name for your storage volume.
  - 2) In the **Format** field, select **qcow2**.
  - 3) In the **Create a disk image for the virtual machine field**, enter a storage value of 32 or more.
- ii. Click **Finish**.

**Add a Storage Volume**

Create storage volume

Create a storage unit to be used directly by a virtual machine.

Name:

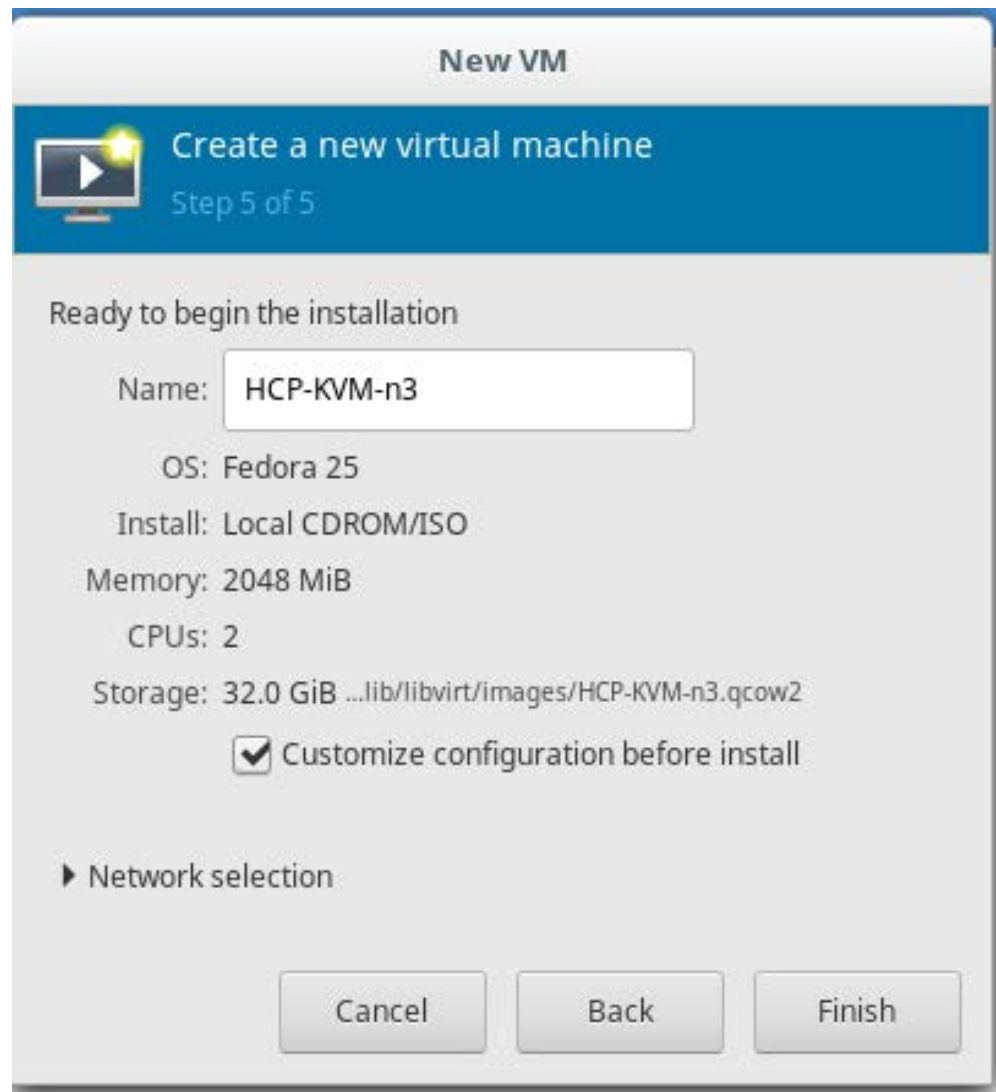
Format:

► Backing store

**Storage Volume Quota**  
kvm-fc1's available space: 993.86 GiB

Max Capacity:    GiB

10. Click **Forward**.
11. To configure the **New VM Step 5 of 5 window**:
  - a. In the **Name** field, type the name of your HCP VM.
  - b. Select **Customize configuration before install**.



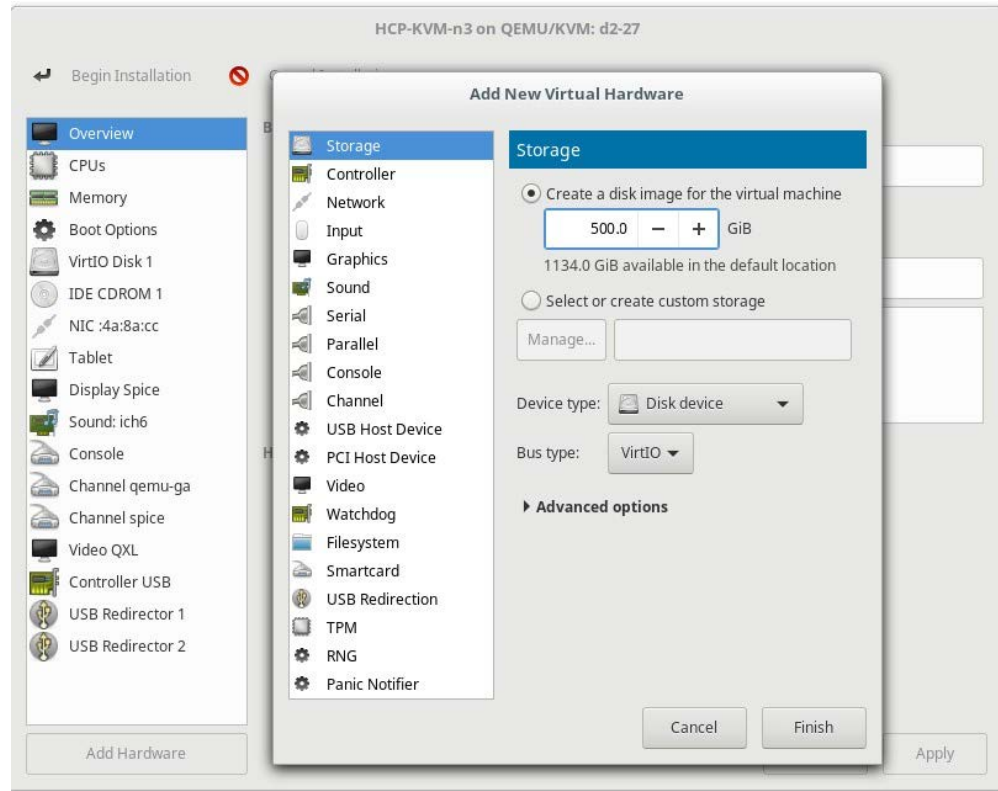
12. Click **Finish**. The customization window opens.

## Customizing the HCP VM node for HCP configuration

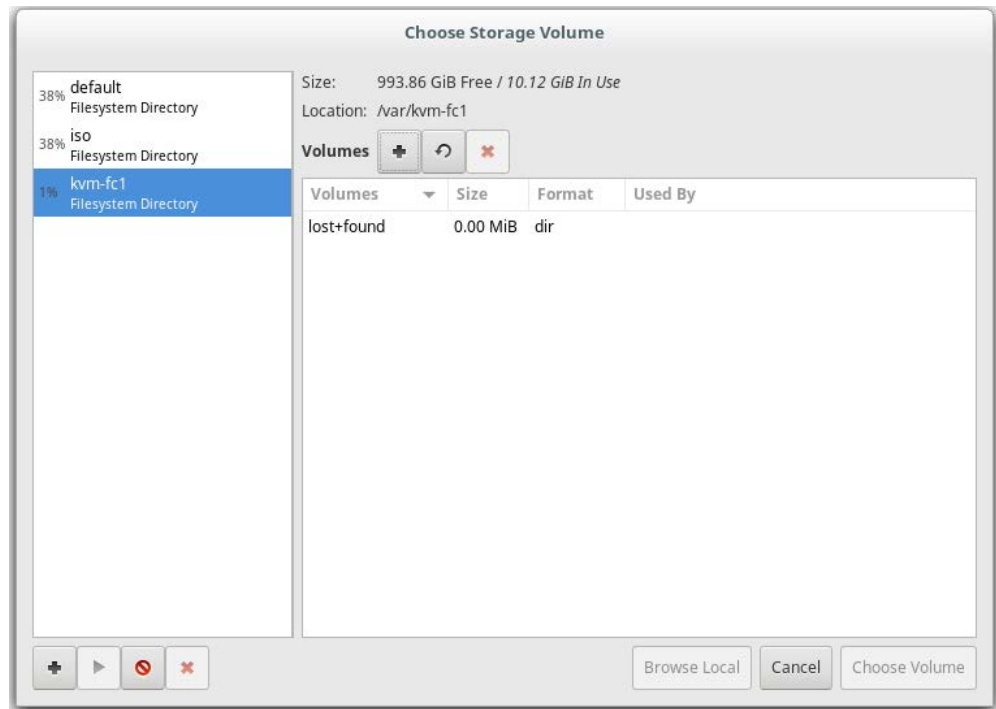
### Procedure

1. In the customization window, on the left navigation pane, right-click the **Overview** tab, and in the menu, click **Add Hardware**.
2. In the **Add New Virtual Hardware** window, on the left navigation pane, click the **Storage** tab.

3. To configure the **Storage Volume** window, perform one of these procedures:
  - If you are creating a disk image on the virtual machine:
    - a. Select **Create a disk image for the virtual machine**.
    - b. In the **Create a disk image for the virtual machine** field, enter a storage value of 500 or more.
    - c. In the **Device type** field, select **Disk device**.
    - d. In the **BUS type** field, select **VirtIO**.



- If you are using SAN storage and want to create a custom storage location:
  - a. Select **Select or create custom storage**.
  - b. In the **Device type** field, select **Disk device**.
  - c. In the **BUS type** field, select **VirtIO**.
  - d. Click **Manage**.
  - e. In the **Choose Storage Volume** window that opens, in the left panel, select the storage pool you want to host your storage volume.
  - f. Click the **Volumes** plus sign.



- g. To configure the **Add a Storage Volume** window that opens:
  - i. In the **Name** field, type a name for your storage volume.
  - ii. In the **Format** field, select **qcow2**.
  - iii. In the **Create a disk image for the virtual machine field**, enter a storage value of 500 or more.
- h. Click **Finish**.

**Add a Storage Volume**

Create storage volume

Create a storage unit to be used directly by a virtual machine.

Name:

Format:

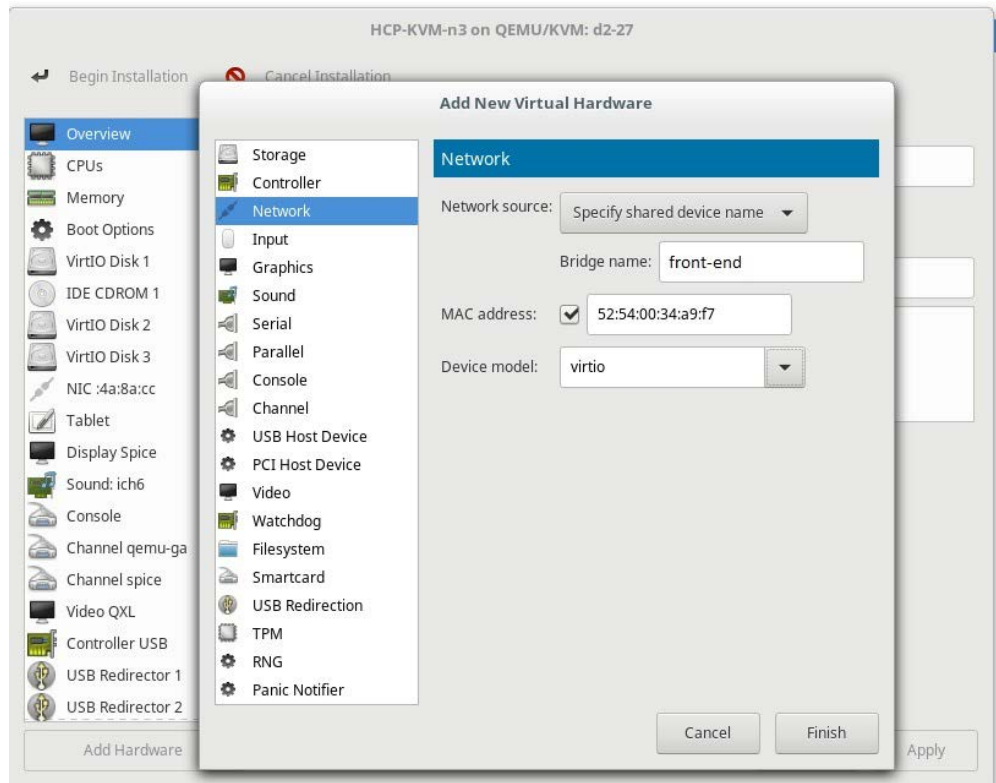
► Backing store

**Storage Volume Quota**  
kvm-fc1's available space: 993.86 GiB

Max Capacity:    GiB

4. Click **Forward**.
5. To configure the **Add New Virtual Hardware** window:
  - a. In the **Device type** field, select **Disk device**.
  - b. In the **BUS type** field, select **VirtIO**.
6. Click **Finish**.
7. After you return to the customization window, repeat Step 3, Substeps a through f for SAN storage, to create a second storage volume.
8. After you create two storage volumes, on the left navigation pane of the customization window, right-click the existing NIC and, from the menu, click **Delete**.
9. On the left navigation pane of the customization window, right-click the **Overview** tab and, in the menu, click **Add Hardware**.

10. To configure the **Add New Virtual Hardware** window that opens:
  - Click the **Network** tab, and in the **Network source** field, select **Specify shared device name**.
  - In the **Bridge name** field, type the name of your front-end network bridge.  
Following the example used in the networking chapter, the bridge file name for your front-end network will be `front-end`.
  - Select **MAC address** and leave the default MAC address.
  - In the **Device model** field, select **virtio**.



11. Click **Finish**.
12. After you return to the customization window, repeat Steps 9 through 11 to create a virtual back-end network.

In the **Bridge name** field, enter the name of your back-end network bridge instead of your front-end network bridge. The bridge file name for your back-end network will be `back-end`.

Following the example used in the networking chapter, the bridge file name for your back-end network will be **back-end**.

**Important:** Do not click **Apply**. If you click **Apply**, the installation fails, and you lose all of your configured settings.

13. Click **Begin Installation**.  
After you begin the installation, you are asked to configure the operating system.

## Performing the OS installation

To install the appliance on a node:

### Procedure

1. In Virtual Machine Manager, double-click the new HCP VM node.
2. Enter the KVM host password.
3. After the virtual machine starts, either press **Enter** or let the program default to the installation option after 75 seconds.

The installation program prompts whether to preserve or clear existing storage volumes.

```
P) Preserve storage volumes during installation
C) Clear storage volumes during installation
E) Exit the installation
```

```
Type your selection and press enter [pce]:
```

4. Enter `c` to clear existing storage volumes.
5. In response to the confirming prompt, enter `y`.

```
You have chosen to clear the storage volumes
THIS OPTION WILL DESTROY ANY DATA ON THE STORAGE VOLUMES.
Are you sure you want to clear the storage volumes (yN):
```

6. When prompted, enter the front-end network IP mode for the KVM host. The IP mode that you specify is used to set both the system-level IP mode and the [hcp\_system] network IP mode.  
Enter the front-end network IP mode ([IPv4], IPv6, Dual):
7. If the installer detects both BaseT and SFP+ network interface cards in this system, you are prompted to enter the front-end network interface types for the HCP system.  
Enter the front-end network interface type ([BaseT], SFP+)
8. When prompted, enter `y` to use a VLAN ID, or enter `n` if you don't want to provide a VLAN ID.  
Do you want to provide a VLAN ID for the front-end network? [n]:
9. If you entered `IPv4` or `Dual` in response to the prompt in Step 6, specify the IPv4 HCP VM node IP address, subnet mask, and gateway IP address for the front-end network; otherwise, go to Step 10.
  - a. When prompted, enter the IPv4 address assigned to the HCP VM node for the front-end network.



**Important:** Do not enter the front-end IP address for the KVM host. Enter the front-end network address for the HCP VM node.

```
Enter the front-end IPv4 IP address []:
```

```
--->
```



- b. When prompted, enter the IPv4 address subnet mask for the front-end network.

```
Enter the front-end IPv4 netmask [255.255.255.0]:
--->
```

- c. When prompted, enter the IPv4 gateway IP address for the front-end network.

```
Enter the front-end IPv4 gateway IP address [172.20.43.254]:
--->
```

If you entered IPv4 in response to the prompt in Step 6, you are finished entering front-end network configuration information for the node. Skip the next step in this procedure, and go to Substep 11.

If you entered IPv6 or Dual, proceed to Substep 10.

10. If you entered IPv6 or Dual in response to the prompt in Step 6, specify the primary IPv6 node IP address, prefix length, and gateway IP address for the front-end network:

- a. When prompted, enter the primary IPv6 address assigned to the HCP VM node for the front-end network.



**Important:** Do not enter the front-end IP address for the KVM host. Enter the front-end network address for the HCP VM node.

```
Enter the front-end IPv6 IP address []:
--->
```

- b. When prompted, enter the IPv6 address prefix length for the front-end network.

```
Enter the front-end IPv4 prefix length [64]:
--->
```

- c. When prompted, enter the IPv4 gateway IP address for the front-end network.

```
Enter the front-end IPv4 gateway IP address [172.20.43.254]:
--->
```

- d. When prompted, enter `y` to assign a secondary IPv6 address to the node for the front-end network, or `n` if you don't want to assign a secondary IPv6 address to the node for the front-end network.

```
Do you want to provide a second IP for the front-end IPv6
network? [n]:
```

- e. If you entered `y` in response to the prompt in Substep d, specify the secondary IPv6 node IP address, prefix length, and gateway IP address for the front-end network

11. When prompted, enter the back-end network IP address for the HCP VM node.



**Important:** Do not enter the back-end IP address for the KVM host.  
Enter the back-end IP address for the HCP VM node.

```
Enter the back-end IPv4 IP address [] :  
---->
```

The installation program displays your responses to all of the previous prompts and asks you to confirm them.

12. In response to the confirming prompt, enter `y` to confirm your responses or `n` to change any responses.

In this case, the installation program repeats the prompts, starting again with the front-end network bonding mode.

At this point, the installation program runs a precheck to see if LUNs 0 or 128 exist. If the precheck finds these LUNs, the install fails. You must remove the LUNs before proceeding.

The installation program reformats the system volume and installs the OS. This process takes several minutes. While installing the OS, the installer should report its progress to the console.

If the OS installation is not proceeding as expected, you can press **Alt+F2** to display a command prompt. This enables you to enter commands that can help you diagnose the problem. Pressing **Alt+F2** to display a command prompt works only during OS installation. To return to the OS installation display, press **Alt+F1**.

### Next steps

When the installation is complete, the HCP VM node shuts down. If it does not restart, power on the HCP VM node using Virtual Machine Manager.

After the HCP VM node restarts, you are asked to change the password for the install user.

## Changing the install user password

After installing the HCP OS, you must change the password on your HCP VM node. To change the password:

### Procedure

1. Log in to the HCP VM node console with the default login information:
  - Username: `install`
  - Password: `Chang3Me!`

```
Appliance Operating System release 6.8  
3.1.S-S.X86_64
```

```
Press ALT+F5 for Appliance Application Status  
Press ALT+F6 for Appliance Process Status  
press ALT. for Appliance Diagnostics
```

```
Press ALT+F1 to return to this login screen
```

```
aos login:
```

2. Change the password to `hcpinstall` (the last two characters are the numeral one).

```
Press ALT+F6 for Appliance Process Status
```

```
Press ALT+F8 for Appliance Diagnostics
```

```
Press ALT+F1 to return to this login screen
```

```
aos login: install
```

```
Password:
```

```
Identity added: /var/home/install/.ssh/id_dsa (/var/home/
install/.ssh/id_dsa)
```

```
Home: /var/home/install
```

```
Changing password for user install.
```

```
Changing password for install.
```

```
(current) UNIX password:
```

```
New password:
```

```
Retype new password:
```

```
Sorry, passwords do not match.
```

```
New password:
```

```
Retype new password:
```

```
Passwd: all authentication tokens updated successfully.
```

```
Password updated.
```

```
If you install your application from this node, the new password
will be propagated to all other nodes.
```

```
Press ENTER to continue:
```

3. Reenter the new password.
4. Press **Enter**.

## Installing the HCP software

The HCP installation is performed from the node with the highest last octet in its back-end IP address.

For example, if the four back-end IP addresses for a system are `172.21.150.150`, `172.21.150.151`, `172.21.150.152`, and `172.21.150.153`, perform the HCP software installation on node `172.21.150.153`.



**Note:** Although you can install the HCP system, you cannot enable data-at-rest encryption (DARE). DARE encrypts data on primary storage and data tiered to external storage pools. If you plan to use DARE features, contact your authorized HCP service provider before performing the software installation.

## Procedure

1. Access the Virtual Machine Manager.
2. Double-click the highest-numbered node to open the console.
3. Login to the HCP VM node console with the default login information:
  - Username: `install`
  - Password: `Chang3Me!`
4. Change the password to `hcpinstall` (thelast two characters are the numeral one).
5. Press **Enter** to display the **HCP Configuration** menu. This menu specifies the HCP version that you are running.

```
HCP 9.0 Configuration Menu
=====

[1] Get HCP Setup Files
[2] Install an HCP System
[3] Upgrade an HCP System
[4] Add a Node to an HCP System
[5] Perform Checks for Offline Upgrade
[6] Perform Checks for Online Upgrade
[v] Add Logical Volumes to an HCP System
[s] Perform a Service Procedure
[q] Log Out

Currently installed version:  9.0
Version on CD/DVD:           None
Extracted version:           9.0

Enter a selection: 2

You chose "2", is this correct? [Default: yes]:
```

6. Enter 2.
7. In response to the confirmation prompt, press **Enter**.  
The **New Install Menu** in the HCP Setup wizard is displayed.

```
HCP Setup: New Install Menu
=====

[1] HCP Nodes
[2] Distributor/OEM Key Access (Arizona)
[3] Networking Settings
```

```

[4] DNS Settings
[5] Time Settings
[6] Internal Configuration Settings
[7] Security and Encryption Settings

[c] Load HCP Configuration File
[r] Restore Default Configuration
[q] Review Current Configuration

[x] Install a New HCP System with This Configuration

[w] Exit/Write out Configuration File
[q] Return to Configuration Menu

Enter your choice.
[Default: 1]:

```

## Identifying the nodes in the HCP VM system

To identify the nodes in the HCP VM system:

### Procedure

1. From the **HCP 9.1.0 Configuration** menu, enter **3** to run the HCP Setup wizard.
2. In response to the confirming prompt, enter **y** or **yes** to confirm your entry or **n** or **no** to try again.

The HCP Setup wizard **New Install Menu** appears.

```

HCP Setup: New Install Menu
=====
[1] HCP Nodes

[r] Restore Default Configuration
[v] Review Current Configuration

[x] Install a New HCP System with This Configuration

[w] Exit/Write out Configuration File
[q] Return to Configuration Menu

Enter your choice.
[Default: 1]:

```

3. Enter **1** to identify the nodes in the HCP system. The **HCP Nodes** appears.

```

HCP Nodes Menu
=====

[1] Storage Node Back-end IP Addresses

```

```
[b] Go Back to the Previous Menu
[q] Return to the Configuration Menu

Enter your choice.
[Default: 1]:
```

4. From the **HCP Nodes Menu**, enter **1** to identify the storage nodes in the HCP system. Use the *back-end IP address* to identify each node.



**Tip:** If you chose to enter the node IP addresses as literal values, enter the IP address of the lowest-numbered node first. For subsequent IP addresses, HCP Setup gives a default value that is one greater than the previous IP address that you entered.

5. From the **HCP Nodes Menu**, enter **b** to return to the **New Install Menu**. The **New Install Menu** now includes additional options for configuring the HCP system.

```
HCP Setup: New Install Menu
=====
[1] HCP Nodes
[2] Distributor/OEM Key Access (Arizona)
[3] Networking Settings
[4] DNS Settings
[5] Time Settings
[6] Internal Configuration Settings
[7] Security and Encryption Settings

[c] Load HCP Configuration File
[r] Restore Default Configuration
[q] Review Current Configuration

[x] Install a New HCP System with This Configuration

[w] Exit/Write out Configuration File
[q] Return to Configuration Menu

Enter your choice.
[Default: 2]:
```

## Configuring the HCP system

### Procedure

1. From the **New Install Menu**, enter **2** to change the key access settings.

```
Distributor/OEM Key Access
=====

Please enter a valid distributor key for your company (supplied by
```

```
MDS).
```

Entering this key enables branding and other features specific to your company. If you do not need to enter a distributor key or are performing an NDS-Internal HCP deployment, accept the default. All keys are case sensitive.

Note: Control-C cancels input.

```
Enter distributor key.
[Default: Arizona]:
You chose: "Arizona", Is this correct?
(Default: yes):
```

2. Change the distributor key.



**Tip:** If this system is provided by Hitachi Vantara, keep the default Arizona key.

3. Enter `y` or `yes` to confirm the change.  
You are returned to the **New Install Menu**.
4. Enter `3` to configure the networking options.

```
HCP Networking Options
=====

[1] Gateway Router IP Address [172.28.27.254]
[2] Multicast Network [238.177.1.1]

[b] Go Back to the Previous Menu
[q] Return to Configuration Menu

Enter your choice.
[Default: 2]:
```

You are returned to the **New Install** menu.

5. Enter `1` and change the **Gateway Router IP address**.
6. Enter `2` and change the **Multicast Network**.
7. Enter `b`.

```
HCP DNS Options
=====

[1] Enable DNS (Yes)
[2] Domain Name for the System (None)
[3] DNS Server(s) (192.168.100.45)

[b] Go Back to the Previous Menu
[q] Return to Configuration Menu
```

```
Enter your choice.
[Default: 1]:
```

8. Enter 4 to configure the DNS options.
9. Enter 2 to input the domain name for the system.
10. Enter the system domain name.

```
Domain for the System
=====

Please enter the fully qualified name of the system from the corporate
DNS
configuration. If you are not using DNS, enter a dummy name to be used
for
system access.

Example: HCP1.example.com

Note: Control-C cancels input.

Enter system domain name:
[Default: None): cluster-vim-1.wilco.net

You chose: "cluster-vim-1.wilco.net", is this correct?
Default: [yes]:
```

11. Make sure **Option 1: Enable DNS** is set to **yes**.
12. Make sure **Option 3: DNS Servers** is set to the proper corporate DNS server.
13. Enter **b**.  
You are returned to the **New Install Menu**.
14. Enter 5 to configure the time settings.
15. Enter 1 and set the time configuration to a time server.  
Use the same time server that has been configured for all KVM hosts in the HCP VM system.

```
HCP Time Options
=====

[1] Time-Server Configuration (internal)
[2] Current Date and Time (not specified)
[3] Time Zone (America/New_York)
[4] Time Settings Compliance Mode (False)

[b] Go Back to the Previous Menu
[q] Return to Configuration Menu

Enter your choice.
[Default: 1]:
```



16. Specify an external time server or enter `internal`.

```
Time-Server Configuration
=====

What type of time server do you want the HCP system to use? You can
specify
"internal" or at most three time servers. You will be asked to specify
the
names or IP addresses one at a time. For you to specify an external
time
server, the HCP system must have connectivity to the time server
through the
front-end network.

Example (time.nist.gov): 192.13.244.18

Note: Control-C cancels input.

Internal or time server name or IP address.
[Default: internal]: 61.90.182.55

You chose: "64.96.182.55", is this correct?
[Default: yes):
```

17. Enter `y` or `yes` to confirm the change.  
You are returned to the **New Install Menu**.
18. Enter `6` to change the internal configuration settings.  
The **Internal Configuration Settings** menu is displayed.

```
Internal Configuration Settings
=====

[1] Storage Configuration (Not Set)
[2] HCP System Serial Number (00001)
[3] Enable Replication on This System (Yes)
[4] Reinstallation with DNS Failover in Effect (Not)
[5] Customer Support Contact Information

[b] Go Back to the Previous Menu
[q] Return to Configuration menu

Enter your choice.
[Default: 1]:
```

19. Enter `1` to set the storage configuration.

```
Storage Configuration
=====
```

```
What type of storage does this HCP system use? If the storage is
local/internal RAID, type "internal". If the storage is fibre channel
or other
SAN-attached storage, type "external".
```

Note: Control-C cancels input.

```
Enter internal or external.
(Default: internal): internal
```

```
You chose: "internal", is this correct?
(Default: yes):
Do you want to configure a dedicated database volume?
(Default: no): yes
```

```
You chose: "yes", is this correct?
(Default: yes):
```

20. Type `internal`.
21. In response to the confirming prompt, press **Enter**.  
Optionally, to configure a dedicated database volume, the system must have at least three drives per node. All dedicated database volume sizes must be at least 50 GB for a new installation. All dedicated database volumes must be the same size. HCP Setup asks you to configure a dedicated database volume only if your system meets the requirements mentioned previously.
22. If HCP Setup asks whether you want to configure a dedicated database volume, enter `yes` to configure a dedicated database volume; otherwise, enter `no`.
23. Press **Enter** to confirm your choices and return to the **Internal Configuration Settings** menu.
24. From the **Internal Configuration Settings** menu, enter **2** to set the serial number for the HCP system.

```
HCP System Serial Number
=====

Please enter a valid serial number. You will be prompted twice for
verification. The serial number can contain only letters, numbers,
spaces,
hyphens, underscores, and number signs and must not be blank.

Example: 00001
Note: Control-C cancels the input.
Enter a valid serial number.
[Default: 1001001]: 1001001

Please enter it again.
[Default: None: 1001001]
```

25. Enter the unique serial number for this HCP system.

26. Enter the serial number again for confirmation and return to the **Internal Configuration Settings** menu.



**Important:** The HCP system serial number is needed to license the system. If you omit the serial number, the system reports that you are in violation of your license agreement.

27. From the **Internal Configuration Settings** menu, enter 3 to confirm whether replication is enabled.

If you enter `yes` to enable replication, the wizard asks if this is a reinstallation of a primary system after a replication failover with DNS failover enabled. If you enter `yes` to this prompt, it requests that target replicated namespaces in this system continue to be redirected to the replica until data recovery is complete, provided that those namespaces are configured to accept such requests.



**Important:** Do not enable replication if you have not purchased this feature. Doing so makes the system violate your license agreement.

28. From the **Internal Configuration Settings** menu, enter 4 to configure whether reinstallation with DNS failover is enabled.
29. From the **Internal Configuration Settings** menu, enter 5 to set contact information. To specify no contact information, press the **Spacebar**.
30. Enter `b` to return to the **New Install Menu**.

## Running the HCP installation

### Before you begin

If you enabled encryption in the previous section, have your security administrator present for this procedure. The security administrator should be the only person to see the encryption key.

### Procedure

1. From the **New Install Menu**, enter `x`. If you run the installation as the install user, the wizard informs you that data-in-flight encryption is enabled.

The wizard asks for confirmation that it is legal to ship a system with data-in-flight encryption enabled to the country where the system is to be deployed.

```
Confirm Data in Flight Encryption / SSL
```

```
=====
```

```
Data-in-flight encryption has been enabled for this HCP system. Global
trade
compliance prohibits shipping HCP systems to restricted countries with
this
feature enabled. Are you sure it is legal to ship an HCP system with
data-in-
flight encryption enabled to the country where the system will be
deployed?
```

```
Note: Control-C cancels input.
```

```

Enter yes or no.
[Default: no]: yes
You chose: "yes", is this correct?
[Default: yes]:

```

2. Enter `yes` to continue.

3. Press **Enter** to confirm.

The wizard displays the configuration confirmation.

```

Configuration confirmation.
=====
DNS Server (s) = 172.18.4.46
Allow Data at Rest Encryption = No
Customer Support Contact Information = United States:
(800) 446-0744. Outside the United states: (858) 547-4526
Multicast Network = 238.177.1.1
Storage configuration = internal
Time Zone = America/New_York
Gateway Router IPv4 Address = 172.20.59.254
Current Date and Time = None
Domain Name for the System = hcp.example.com
Encrypt Data at Rest on Primary Storage = No
Reinstallation with DNS Failover in Effect = No
Allow Data in Flight Encryption / SSL = Yes
Time Settings Compliance Mode = No
HCP System Serial Number = 00001
Blade Servers = No
Distributor/OEM Key Access = Arizona
MOE Index-only volumes = No
Time Server(s) = internal
Gateway Router Secondary IPv6 Address = None
Gateway Router IPv6 Address = None
Enable DNS = Yes
Chassis = None
Enable Replication on This System = Yes
Configure Dedicated Database Volumes = Yes
Spindown Volumes = No
HCP Storage Nodes: 4
172.59.42.1
172.59.42.2
172.59.42.3
172.59.42.4

Use SHIFT+PGUP to review the Configuration.

IS this Configuration Correct?
(Default: no): yes

```

```
You chose: "yes", is this correct?
(Default: yes):
```

4. Review the configuration. Then complete one of the following actions:
  - If the configuration is not correct:
    - a. Enter `n` or `no`.
    - b. In response to the confirmation prompt, enter `y` or `yes`.
    - c. Correct the configuration information.
  - If the configuration is correct:
    - a. Enter `y` or `yes`.
    - b. In response to the confirmation prompt, enter `y` or `yes`.
5. After you confirm that the configuration information is correct, HCP Setup performs a set of installation prechecks.

```
You chose "yes", is this correct?
[Default: yes]:
Verifying system name
Verifying run location
Verifying running as install
Verifying node connections
Verifying SSH keys
Verifying SSH
Verifying systemwide SSH
Verifying total memory > 32GB
Verifying all network links
Verifying software versions
Verifying 64-bit hardware platform
Verifying drive size
Verifying disk space
Verifying nobody using /fcfs_*
Verifying nobody using /fs/*
Verifying multicast enabled
Syncing install password to all nodes.
Updating EULA
Syncing timezone to all nodes
Syncing date to all nodes.
Generating auth keys
Generating system UUID
Syncing HCP package to all nodes
Checking to see if we need to run update schemaupgrade scripts False
Updating schema scripts for upgrade.
```

6. (Optional) If you previously selected that you want to configure dedicated database volumes, complete the following substeps:
  - a. When prompted, select the dedicated database volume for the first node.
  - b. Press **Enter** to confirm your selection.
  - c. Repeat the Substeps a and b for each node in the system.

- d. After you select the dedicated database volumes for each node, HCP Setup confirms your selections and asks if you want to continue.

```
Select dedicated volume for each node.
Found these volumes:
  node 001:
    1. /dev/sdd at 2:0:0:1 (500GB)
    2. /dev/sde at 2:0:0:2 (500GB)
    3. /dev/sdd at 2:0:0:3 (500GB)
    4. /dev/sde at 2:0:0:4 (1TB)

  node 002:
    1. /dev/sdd at 2:0:0:1 (500GB)
    2. /dev/sde at 2:0:0:2 (500GB)
    3. /dev/sdd at 2:0:0:3 (500GB)
    4. /dev/sde at 2:0:0:4 (1TB)

  node 003:
    1. /dev/sdd at 2:0:0:1 (500GB)
    2. /dev/sde at 2:0:0:2 (500GB)
    3. /dev/sdd at 2:0:0:3 (500GB)
    4. /dev/sde at 2:0:0:4 (1TB)

  node 004:
    1. /dev/sdd at 2:0:0:1 (500GB)
    2. /dev/sde at 2:0:0:2 (500GB)
    3. /dev/sdd at 2:0:0:3 (500GB)
    4. /dev/sde at 2:0:0:4 (1TB)

Select dedicated database volume for node 001: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct?
[Default: yes]:
Select dedicated database volume for node 002: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct?
[Default: yes]:
Select dedicated database volume for node 003: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct?
[Default: yes]:
Select dedicated database volume for node 004: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct?
[Default: yes]:
Following volumes will be configured as dedicated database volumes:
  node 001: 4. /dev/sde at 2:0:0:4 (1TB)
  node 002: 4. /dev/sde at 2:0:0:4 (1TB)
  node 003: 4. /dev/sde at 2:0:0:4 (1TB)
  node 004: 4. /dev/sde at 2:0:0:4 (1TB)
Do you want to continue? [Default: yes]?
```

- e. Press **Enter** to continue.

If the prechecks are successful, the HCP software is installed on all nodes in the system.

Depending on the size of the logical volumes, this can take from several minutes to several hours.

If you enabled encryption in the system configuration, HCP Setup performs some initial setup tasks and then displays the encryption key.

Setup then prompts you to enter the key.



**Important:** Before entering the encryption key, record it. After you enter the key, HCP Setup completes the installation. You do not get a second chance to see the key, and it is not stored for later retrieval.

When the installation is complete, HCP Setup logs you out and restarts the nodes.

The console then displays the login prompt.

If HCP Setup exits before installation processing is complete, record all error messages. Then contact your authorized HCP service provider for assistance.

After the installation is complete, the HCP VM nodes restart, and, instead of the operating system login prompt, you should see an `hcp-node- nodeName` prompt.

You can also verify the run level of a node by pressing **Alt+F5** at the console prompt.

```
Every 30.0s: /sbin/system-info                               Fri May 15
12:29:58 2020
Host Name:                hcp-node-150.cluster-colo-089-
vol.lalo.arehivas.com
IS Mode:                   150

[hcp_system] IP:           172.21.159.158
[hcp_system] Mask:         255.255.255.0
[hcp_system] Gateway:     172.20.27.254
[hcp_backend] IP:         172.21.159.150
[hcpbackend] Mask:        255.255.255.0
Version:                   6.0.0.93

Operating System:          OS 6.0.0.514
Linux Kernel:              3.1.5-5.x06_64
Current Run Level:         4

12:29:58 up 22:47, 0 users, load average: 0.00, 0.01, 0.06
```

## Verifying the HCP software installation

### Procedure

1. Open the **System Management Console** by entering one of the following URLs in a client web browser:

- If the HCP system is configured for DNS:

`https://admin.hcp-domain-name:8000`

- If the HCP system is not configured for DNS:

`https://node-ip-address:8000`

where *node-ip-address* is the front-end IP address of a storage node in the HCP system.

If you enter `http` instead of `https`, the browser returns an error.

2. When prompted, accept the self-signed HCP SSL server certificate either permanently or temporarily.

Set a temporary certificate if you plan to install a trusted certificate later. The **System Management Console** login page is displayed.



**Tip:** If the browser cannot find the **System Management Console** login page, wait a few minutes and then try again. If the login page still does not open, contact your authorized HCP service provider.

3. Verify the serial number on the login page.

If it is incorrect, contact your authorized HCP service provider.

4. Log in to the **System Management Console** with the following username and password:

Username: `security`

Password: `Chang3Me!`

After you log in, the console displays either the **Change Password** page or the **Hardware** page. Perform one of the following actions:

- If the **Hardware** page is displayed, the nodes are still starting HCP. This process can take several minutes. When more than half of the nodes have completed the startup process, the console displays the **Change Password** page. Change your password.
- If the **Hardware** page remains displayed after several minutes, contact your authorized HCP service provider.

5. On the **Change Password** page, enter the following information:

- a. In the **Existing Password field**, enter `Chang3Me!`.
- b. In the **New Password field**, enter a new password, using the following criteria:
  - Must contain UTF-8 characters, including whitespaces
  - Minimum of six characters



- Maximum of 64 characters
  - Must include at least one character from two of the following groups: alphabetic, numeric, and special characters. Examples:
    - Valid password: P@sswOrd
    - Invalid password: password
- c. In the **Confirm New Password** field, retype your new password, then click **Update Password**.
6. In the top-level menu, click **Hardware**.
7. On the **Hardware** page, ensure the following statuses:
- Node status: **Available**.
  - Status of each logical volume: **Available**.
- To see the status of a logical volume, hover over the volume icon.
- If all the nodes and logical volumes are available, the installation was successful and you can begin creating tenants. However, you might not want to do this until all additional setup is complete.
- If any node has a status other than **Available**, or any logical volume associated with an available node has a status other than **Available** or **Spun down**, contact your authorized HCP service provider. Also contact your service provider if the number of logical volume icons for any node does not match the expected number of logical volumes for the node.
8. Log out of the System Management Console and close the browser window.

## Monitoring and alerting

HCP hardware appliance features such as redundant hardware, monitoring, alerting and failover behavior cannot be used by KVM. To maintain performance and data integrity, HCP VM system hardware must be monitored for failures outside of the virtual machine environment.

Hitachi servers and network components that are part of the HCP VM system can be connected to Hitachi Remote Ops for monitoring. For more information, see [Configuring HCP monitoring with Hitachi Remote Ops \(on page 84\)](#).

To monitor hardware supplied by vendors other than Hitachi, use a vendor-supplied or hardware-compatible software tool.

Any failures in the HCP VM infrastructure must be corrected as soon as possible. Drive failures in particular should be closely monitored, because of the possibility of long RAID rebuild times.

HCP Intelligent Platform Management Interface (IPMI) monitoring and Hitachi array monitoring is not available for HCP VMs.

## Software monitoring

HCP maintains a system log that logs all system events. You can view this log in the HCP System Management Console. You can send system log messages to syslog servers, System Network Management Protocol (SNMP) managers, and email addresses. Additionally, you can use SNMP to view and, when allowed, change HCP system settings.

You can generate chargeback reports to track system capacity and bandwidth usage at the tenant and namespace levels.

You can use Hitachi Remote Ops to monitor the health of the HCP software.

## HCP VM resource monitoring

HCP uses System Activity Reporter (SAR) data for resource usage reporting. SAR runs on each node in the HCP system. Every 10 minutes, SAR records statistics about the average use of resources in the node for the past time interval. The graphs on the resources page of the System Management Console show the statistics for a subset of those resources.

The resources that are monitored include the CPU, logical volumes, memory, and networks.



## HCP VM diagnostic menu

For any HCP VM node, you can run diagnostics that analyze and resolve issues with interactions between nodes and other components of the HCP environment.

HCP VM node diagnostics are available through the HCP System Management Console. The diagnostics let you:

- Ping: Test whether a selected device is accessible through the network.
- Traceroute: Display the network path used for communication between the node and a specified device.
- Dig: Query the DNS for the records that match a specified IP address or domain name.
- Route: Display the routing table for a node.
- Showmount: Display the NFS exports table for a specified device.

For more information about HCP system monitoring, see the HCP System Management Help.

---

## Chapter 7: Configuring HCP monitoring with Hitachi Remote Ops

Hitachi Remote Ops is a Hitachi Vantara product that enables remote monitoring of the nodes in an HCP VM system. This chapter assumes that Hitachi Remote Ops is installed and running according to the product documentation.

With Hitachi Remote Ops, you can view the status of nodes in an HCP VM system with a web browser. You can also configure Hitachi Remote Ops to send email notifications of error conditions as they occur.

Additionally, you can configure Hitachi Remote Ops to report error conditions to Hitachi Vantara support personnel.

Hitachi Remote Ops is used for monitoring and error notification only. It does not allow any changes to the system.

Hitachi Remote Ops is installed on a server that is separate from the HCP system. The program uses SNMP to retrieve information from HCP, so SNMP must be enabled in HCP.



**Note:** HCP supports IPv4 and IPv6 network connections to Hitachi Remote Ops servers. However, Hitachi Remote Ops support for IPv6 network connections varies based on the Hitachi Remote Ops server operating system. For requirements for Hitachi Remote Ops servers that support IPv6 networks, see the applicable Hitachi Remote Ops documentation.

### Enabling SNMP in HCP

To enable Hitachi Remote Ops to work with HCP, you must enable SNMP in the HCP System Management Console. When you enable SNMP, you can select version 1, 2c, or 3.

By default, Hitachi Remote Ops is configured to support SNMP version 1 or 2c with the community name `public`. If you change the community name in HCP, or if you select version 3, you must configure a new SNMP user in Hitachi Remote Ops to match the you specify in HCP. For more information, see the Hitachi Remote Ops documentation.

#### Procedure

1. Log in to the HCP System Management Console using the initial user account, which has the security role.
2. In the top-level menu of the console, go to **Monitoring** → **SNMP**.
3. In the **SNMP Settings** section on the SNMP page:
  - a. Select the **Enable SNMP at snmp.hcp-domain-name** option.

- b. Select either **Use version 1 or 2c** (preferred) or **Use version 3**.  
If you select **Use version 3**, specify a username and password in the **Username, Password, and Confirm Password** fields.
- c. (Optional) In the **Community** field, type a different community name.
4. Click **Update Settings**.
5. In the entry field in the **Allow** section, type the IP address that you want HCP to use to connect to the server where Hitachi Remote Ops is installed. Then click **Add**.
6. Log out of the System Management Console and close the browser window.

## Configuring Hitachi Remote Ops

To configure Hitachi Remote Ops to monitor the nodes in the HCP system, do the following:

### Procedure

1. Log into Hitachi Remote Ops.
2. Set the Hitachi Remote Ops base configuration, including the email addresses to which email about error conditions should be sent.
3. (Optional) Configure transport agents for reporting error conditions to Hitachi Vantara support personnel.
4. Identify the HCP system to be monitored.

## Log in to Hitachi Remote Ops

### Procedure

1. Open a web browser window.
2. In the address field, enter the URL for the Hitachi Remote Ops (using either the hostname or a valid IP address for the server) followed by the port number 6696; for example: <http://remoteops:6696>
3. In the **Select one of the following UserIds** field, select **Administrator**.
4. In the **Enter the corresponding password field**, type the case-sensitive password for the Administrator user. By default, this password is hds.  
If Hitachi Remote Ops is already in use at your site for monitoring other devices, this password may have been changed. In this case, see your Hitachi Remote Ops administrator for the current password.
5. Click the **Logon** button.

## Set the base configuration

The Hitachi Remote Ops base configuration specifies information such as the customer site ID, how frequently to scan devices, and whether to report communication errors that occur between Hitachi Remote Ops and monitored services. The Hitachi Remote Ops base configuration specifies information such as the customer site ID, how frequently to scan devices, and whether to report communication errors that occur between Hitachi

Remote Ops and monitored devices. The base configuration also specifies the addresses to which Hitachi Remote Ops should send email about error conditions.

If Hitachi Remote Ops is already in use at your site, the base configuration may already be set. In this case, you can leave it as is, or you can make changes to accommodate the addition of HCP to the devices being monitored.

To set the Hitachi Remote Ops base configuration:

### Procedure

1. In the row of tabs at the top of the Hitachi Remote Ops interface, click **Configuration**.
  - a. The Base page is displayed by default. To return to this page from another configuration page, click **Base** in the row of tabs below Configuration.
2. (Optional) In the **Device Monitoring** section:
  - a. Type your Hitachi Vantara customer ID, in the Site ID field.
  - b. Specify different values in the other fields to meet the needs of your site. For information about these fields, click the **Help on this table's entries** link above the fields.
3. In the **Notify Users by Email** section:
  - a. In the **eMail Server** field, type the fully qualified hostname or a valid IP address of the email server through which you want Hitachi Remote Ops to send email about error conditions.
  - b. In the **Local Interface** field, select the Ethernet interface that has connectivity to the specified email server. This is the interface on the Hitachi Remote Ops server.
  - c. In the **User List** field, type a comma-separated list of the email addresses to which Hitachi Remote Ops must send email about error conditions.
  - d. In the **Sender's Email Address** field, type a well-formed email address to be utilized in the From line of each email.

Some email servers require that the value in the **From** line be an email address that is already recognized by the server.
4. Click the **Submit** button.
5. Click the **Test Email** button to send a test email to the specified email addresses. This is an optional step.

## (Optional) Configure transport agents

If needed, your site can optionally specify one or more transport agents. A Hitachi Remote Ops transport agent transfers notifications of error conditions to a target location where Hitachi Vantara support personnel can access them. The transfer methods available are HTTPS, FTP, or dial up. For the destinations for each method, contact your authorized HCP service provider.

After they are configured, Hitachi Remote Ops tries each agent in the order in which they are listed until one is successful.

### Procedure

1. In the row of tabs below **Configuration**, click **Transport Agents**.
2. In the field below **Data Transfer Agents**, select the transfer method for the new transport agent.
3. Click **Create**.  
The new transport agent is displayed in the list of transport agents. A set of configuration fields is displayed below the list.
4. In the configuration fields, specify the applicable values for the new transport agent. For details, see the Hitachi Remote Ops documentation.
5. Click **Submit**.
6. You can change the order of multiple transport agents by moving them individually to the top of the list:
  - a. In the **Move to Top** column, select the transport agent you want to move.
  - b. Click **Submit**.

## Identify the HCP system

### Procedure

1. In the row of tabs at the top of the Hitachi Remote Ops interface, click **Summary**.  
The Summary page displays up to four tables that categorize the devices known to Hitachi Remote Ops - Device Errors, Communication Errors, Devices Okay, and Not Monitored. To show or hide these tables, click the checkboxes below the table names at the top of the page to select or deselect the tables, as applicable. Then click the **Refresh** button. While no tables are shown, the page contains an **Add a Device** link.
2. Follow one of these actions below:
  - a. If the **Summary** page doesn't display any tables, click the **Add a Device** link.
  - b. If the **Summary** page displays one or more tables, click the **Item** column heading in any of the tables.
3. In the **Select Device Type** field, select HCP. A set of configuration fields appear.
4. (Optional) In the **Name** field, type a name for the HCP system.  
The name can be from one through 40 characters long. special characters and spaces are allowed. Typically, this is the hostname of the system.
5. (Optional) In the **Location** field, type the location of the HCP system.  
The location can be from one through 40 characters long. Special characters and spaces are allowed.
6. (Optional) In the Group field, type the name of a group associated with the HCP system.  
The group name can be from one through 40 characters long. Special character spaces are allowed.  
For example: Finance Department.
7. In the **Site ID** field, type your Hitachi Vantara customer ID.  
If you don't know your customer ID, contact your authorized HCP service provider for help.

8. In the **IP Address or Name (1)** field, type a valid front-end IP address for the lowest-numbered storage node in the HCP system. In the **Local Interface** field, leave the value as -any-.
9. In the **IP Address or Name (2)** field, type a valid front-end IP address for the highest-numbered storage node in the HCP system. In the **Local Interface** field, leave the value as -any-.
10. In the **SNMP Access** field, select the SNMP user that corresponds to the SNMP configuration in HCP.  
Typically, this is public. For information about configuring SNMP in HCP, see [Enabling SNMP in HCP \(on page 84\)](#).
11. In the **Comms Error Reporting?** field, select one of these options to specify whether Hitachi Vantara must report communications errors that occur between Hitachi Vantara and HCP:
  - Yes- Report communication errors
  - No- Don't report communication errors
  - Local- Report communication errors only to the email addresses specified in the base configuration and not through the specified transport agents.
  - Default- Use the setting in the base configuration
12. Leave **Enabled?** selected.
13. Leave **Trace?** unselected.
14. Click the **Add** button.

If the operation is successful, the interface displays a message indicating that the HCP system has been added. Do not click the add button again. Clicking the add button twice adds the system a second time.



---

## Appendix A: Configuring SAN storage for the KVM host

This appendix covers how to create a SAN file system and connect it to your KVM hosts. You must mount a file system on your SAN storage for KVM to access before you can deploy an HCP VM. This section assumes you are using a Linux machine as your local computer.

### Logging in to the KVM host

#### Procedure

1. Open a new terminal.
2. Use SSH to log into the KVM host.

### Configuring multipathing

Before setting up DM-Multipath on your system, make sure your system includes the device-mapper-multipath package.

To configure multipathing:

#### Procedure

1. Enter the following command to list all block devices visible to the operating system:

```
lsblk
```

Here is a sample output:

```
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sdf 8:80 0 1T 0 disk
sdd 8:48 0 1T 0 disk
sdb 8:16 0 7.3T 0 disk
sdk 8:160 0 1T 0 disk
sdi 8:128 0 1T 0 disk
sdg 8:96 0 1T 0 disk
sde 8:64 0 1T 0 disk
sdc 8:32 0 7.2T 0 disk
sda 8:0 0 64G 0 disk
??sda2 8:2 0 63G 0 part
? ??fedora00-swap 253:1 0 6.4G 0 lvm [SWAP]
? ??fedora00-root 253:0 0 15G 0 lvm /
```

```

??sda1 8:1 0 1G 0 part /boot
sdj 8:144 0 1T 0 disk
sdh 8:112 0 1T 0 disk

```

2. Enter the following command to set up DM-Multipath for basic failover:

```
mpathconf --enable --with_multipathd y
```

3. Enter the following command to start the multipath service:

```
multipathd
```

4. Enter the following command to verify that SAN LUN devices are listed:

```
lsblk
```

Here is a sample output:

```

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sdf 8:80 0 1T 0 disk
??mpathb 253:3 0 1T 0 mpath
sdd 8:48 0 1T 0 disk
??mpatha 253:2 0 1T 0 mpath
sdb 8:16 0 7.3T 0 disk
sdk 8:160 0 1T 0 disk
??mpathc 253:4 0 1T 0 mpath
sdi 8:128 0 1T 0 disk
??mpathd 253:5 0 1T 0 mpath
sdg 8:96 0 1T 0 disk
??mpathc 253:4 0 1T 0 mpath
sde 8:64 0 1T 0 disk
??mpathd 253:5 0 1T 0 mpath
sdc 8:32 0 7.2T 0 disk
sda 8:0 0 64G 0 disk
??sda2 8:2 0 63G 0 part
? ??fedora-swap 253:1 0 6.4G 0 lvm [SWAP]
? ??fedora-root 253:0 0 15G 0 lvm /
??sda1 8:1 0 1G 0 part /boot
sdj 8:144 0 1T 0 disk
??mpathb 253:3 0 1T 0 mpath
sdh 8:112 0 1T 0 disk
??mpatha 253:2 0 1T 0 mpath

```

The devices should now have a multipath.

5. Enter the following command to list all multipath devices attached to the system:

```
multipath -l
```

## Creating the physical volume on the multipath disk

### Procedure

1. Enter the following command to create a physical volume on a multipath disk:  

```
pvccreate /dev/mapper/multipath-disk-name
```

For example:  

```
pvccreate /dev/mapper/mpatha
```
2. Enter the following command to create a volume group:  

```
vgcreate volume-group-name /dev/mapper/multipath-disk-name
```

For example:  

```
vgcreate mpathafcl /dev/mapper/mpatha
```
3. Enter the following command to create a logical volume on the volume group you created in the previous step:  

```
lvcreate -n lv-name -L 1T volume-group-name
```

For example:  

```
lvcreate -n kvmfcl -L 1023G mpathafcl
```
4. Enter the following command to create an ext4 file system on the multipath disk logical volume:  

```
mkfs.ext4 /dev/mapper/file-system-path
```

For example:  

```
mkfs.ext4 /dev/mapper/mpathafcl-kvmfcl
```
5. Enter the following command to verify that the physical volume is created:  

```
pvdiskdisplay
```
6. Enter the following command to verify that the volume group you created is listed:  

```
vgdisplay
```
7. Enter the following command to verify that the logical volume you created is listed:  

```
lvdisplay
```

## Mounting the file system

### Procedure

1. Enter the following command to create a folder with your file system:  

```
mkdir /var/mount-location-directory-name
```

For example:  

```
mkdir /var/kvm-fcl
```
2. Enter the following command to mount the new file system:  

```
mount /dev/mapper/file-system-path /var/volume-group-name-logical-volume-name
```

For example:  

```
mount /dev/mapper/kvmfcl /var/kvm-fcl/
```
3. Enter the following command to access the fstab:  

```
vi /etc/fstab
```
4. Press **I** to edit the file.

5. Add the following text to the existing file:  

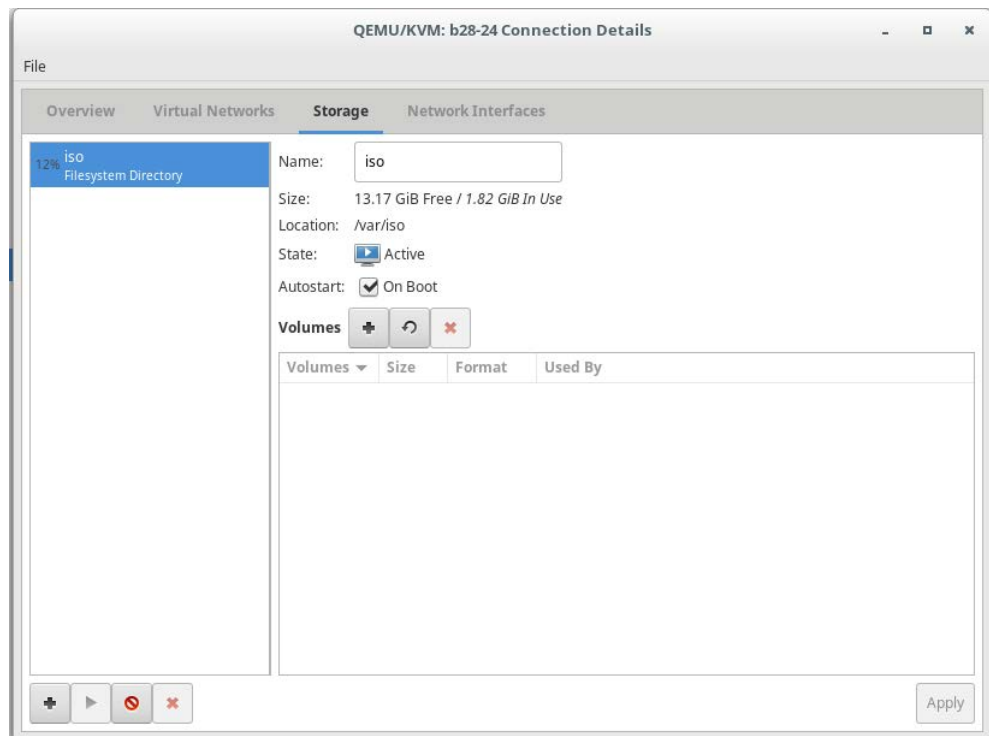
```
/dev/mapper/mpathafcl-kvmfc1 /var/kvm-fc1 ext4 defaults 1 2
```
6. Press **Esc**.
7. Enter the following command to save and exit the file:  

```
:wq
```

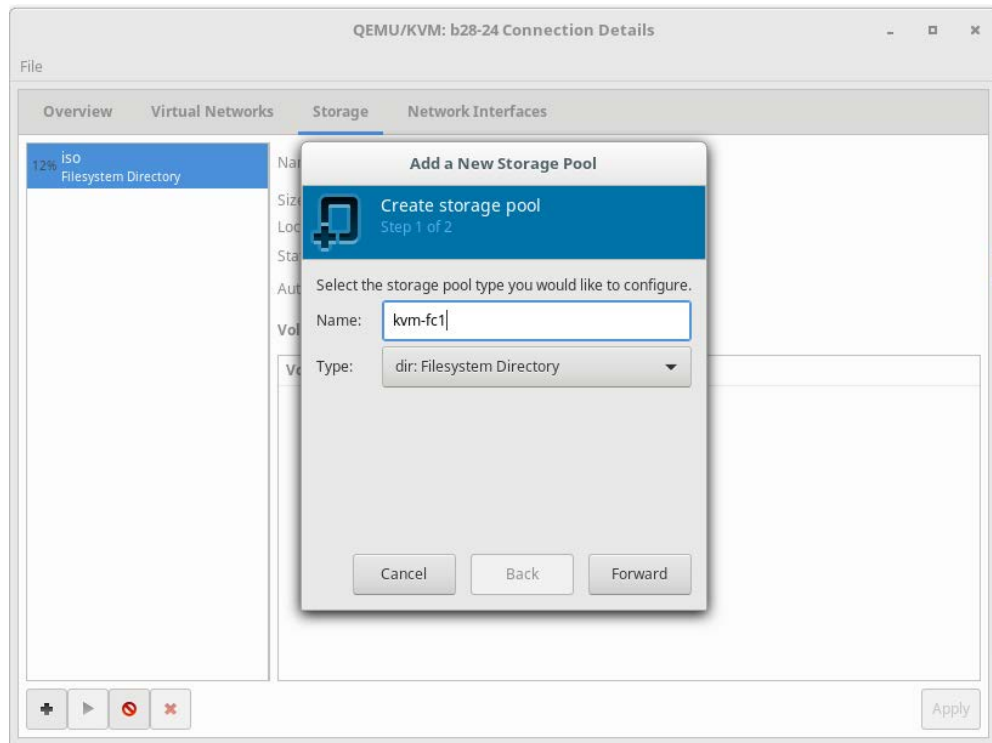
## Adding the new storage pool to the KVM host

### Procedure

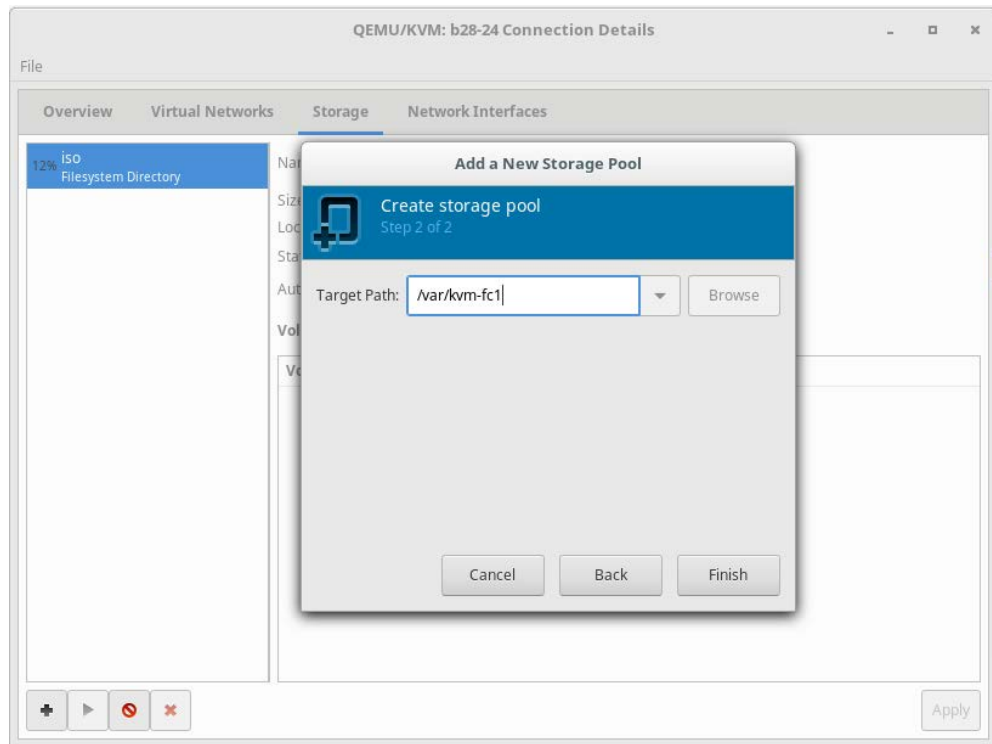
1. In Virtual Machine Manager, right-click the KVM host on which you mounted the volume.  
 The **Connection Details** window opens.
2. Click the **Storage** tab.



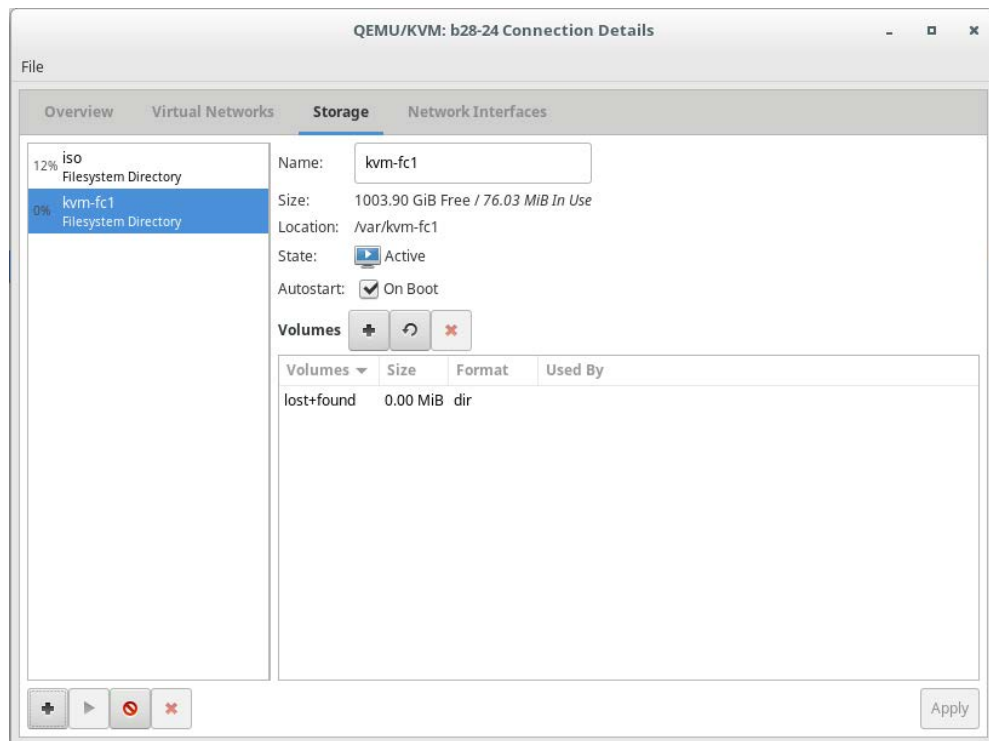
3. Click the **Plus** in the bottom left corner of the window.  
 The **Add a New Storage Pool** window opens.
4. In the **Name** field, type a name for your new storage pool.
5. In the **Type** field, select **dir: Filesystem directory**.
6. Click **Forward**.



7. Click **Browse**.
8. In the **Target Path** field, enter the file path for the mount point location.
9. Click **Finish**.



After the new storage pool is added, it can be carved into storage volumes for virtual machines.



---

## Appendix B: Creating an HCP VM node using the command line

This appendix covers how to deploy an HCP VM node using the command line. This appendix assumes you are using a Linux machine as your local computer. You must repeat these instructions to create one HCP VM node on each KVM host.

Before you deploy an HCP VM using the command line, do the following:

- Install the KVM packages. See [Installing KVM \(on page 16\)](#).
- Create network bridges. See [Configuring KVM networking \(on page 20\)](#).
- Copy and unzip a .iso file on your HCP node. See [Copying the .iso file and sending the Zip file to the KVM host \(on page 51\)](#).

### Logging in to the KVM host

#### Procedure

1. Open a new terminal.
2. Use SSH to log into the KVM host.

### Performing the command line initialization

To create, configure, and deploy the HCP VM, enter a virt-install command with the parameters shown:

```
virt-install -n hcp-vm-node-name -r memory-ram \  
--disk path=/var/lib/libvirt/images/os-disk- name.qcow2,bus=virtio,size=32 \  
\   
--disk path=/var/lib/libvirt/images/storage-disk-name- 1.qcow2,bus=virtio, \  
size=disk-size-1 \  
--disk path=/var/lib/libvirt/images/storage-disk-name- 2.qcow2,bus=virtio, \  
size=disk-size-2 \  
-c File-Path-To-HS222_Release-Number.iso\  
--network bridge=front-end-network-bridge-name,model=virtio \  
--network bridge=back-end-network-bridge-name,model=virtio \  
--noautoconsole -v --vcpus=number-of-virtual-cpus --os-variant=node- os
```

For example:

```
virt-install -n hcp_example -r 32768 \  
--disk path=/var/lib/libvirt/images/hcp_example- 01.qcow2,bus=virtio,  
size=32 \  
--disk path=/var/lib/libvirt/images/hcp_example- 02.qcow2,bus=virtio,  
size=500 \  
--disk path=/var/lib/libvirt/images/hcp_example- 03.qcow2,bus=virtio,  
size=500 \  
-c /var/iso/HS222_8.0.0.682.iso \  
--network bridge=front-end,model=virtio \  
--network bridge=back-end,model=virtio \  
--noautoconsole -v --vcpus=4 --os-variant=fedora25
```



## Hitachi Vantara



Corporate Headquarters

2535 Augustine Drive

Santa Clara, CA 95054 USA

[HitachiVantara.com](http://HitachiVantara.com) | [community.HitachiVantara.com](http://community.HitachiVantara.com)

Contact Information

USA: 1-800-446-0744

Global: 1-858-547-4526

[HitachiVantara.com/contact](http://HitachiVantara.com/contact)