

Hitachi Content Platform

9.1.0

Deploying an HCP VM System on KVM

This book is the setup guide for Hitachi Content Platform Virtual Machine systems. It describes how to deploy a virtualized HCP system in your Kernel-based Virtual Machine (KVM) environment.

© 2017, 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 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.

AlX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, 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 web site 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

Preface	vii
Intended audience	vii
Product version	vii
Release notes	vii
Related documents	vii
Accessing product documentation	ix
Getting help	ix
Comments	x
 Chapter 1: HCP system overview	 1
Introduction to Hitachi Content Platform	1
HCP VM system components and architecture	1
Host platform	2
Compute	2
Storage	3
Front-end network	3
Back-end network	3
Management port network	4
Dedicated database volume	4
Hardware monitoring and alerting	4
HCP software	5
Storage licensing	5
 Chapter 2: Configuration guidelines for the HCP VM environment	 7
Supported KVM versions	7
HCP VM hardware requirements	7
HCP VM system limits	9
HCP VM availability considerations	9

Chapter 3: Installing KVM	11
Configuring the BIOS for KVM	11
Step 1: Log into the KVM host	11
Step 2: Verify that virtualization technology is enabled	11
Step 3 (conditional): Connect a monitor and keyboard to your KVM host	12
Step 4 (conditional): Enter the BIOS	12
Step 5 (conditional): Enable virtualization technology	12
Installing the KVM packages	13
Step 1: Log into the KVM host	14
Step 2: Install the KVM packages	14
Step 3: Install Virtual Machine Manager on your local computer	14
Chapter 4: Configuring KVM networking	17
Resources you need before you start	17
Configuring the KVM host for networking	17
Step 1 (conditional): Connect a monitor and keyboard to your KVM host	17
Step 2: Back up the interface configuration files	18
Step 3: Enable the bonding kernel module	18
Step 4: Create the bridge network files	19
Step 5: Create the bond network files	21
Step 6: Enslave the front-end interface configuration files	23
Step 7: Enslave the back-end interface configuration files	25
Step 8: Restart networking on the KVM host	27
Step 9: Restart the individual KVM host	27
Step 10: Verify that the KVM host network has been updated	27
Chapter 5: Deploying the HCP VM system	31
Prerequisites	31
Deploying the HCP VM system	31
Step 1: Download the .iso installation file	32
Step 2: Copy the .iso file to the KVM host	32
Step 3: Add an HCP VM node connection	33
Step 4: Create the HCP VM node	34
Step 5: Customize the HCP VM node for HCP configuration	43
Step 6: Perform the OS installation	48
Step 7: Change the install user password	52
Installing the HCP software	53
Step 1: Identify the nodes in the HCP VM system	55

Step 2: Configure the HCP system	57
Step 3: Run the HCP installation	61
Step 4: Verify the HCP software installation	66
Monitoring and alerting	68
Software monitoring	69
HCP VM resource monitoring	69
HCP VM diagnostic menu	70
Chapter 6: Maintenance procedures	73
Adding logical volumes	73
Moving storage node databases to optimal volumes	78
Deleting databases from older database volumes	80
Adding HCP VM nodes	82
Recovering storage nodes	86
Recovering storage nodes and preserving volumes	86
Recovering storage nodes and clearing volumes	89
Chapter 7: Configuring HCP monitoring with Hitachi Remote Ops	95
Enabling SNMP in HCP	95
Configuring Hitachi Remote Ops	96
Step 1: Log in to Hitachi Remote Ops	97
Step 2: Set the base configuration	97
Step 3: (Conditional) Configure transport agents	99
Step 4: Identify the HCP system	99
Appendix A: Configuring SAN storage for the KVM host	103
Step 1: Log into the KVM host	103
Step 2: Configure multipathing	103
Step 3: Create the physical volume on the multipath disk	105
Step 4: Mount the file system	107
Step 5: Add the new storage pool to the KVM host	107
Appendix B: Creating an HCP VM node using the command line	113
Step 1: Log into the KVM host	113
Step 2: Perform the command line initialization	114
Glossary	115
Index	127



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 9.1.0 or later of HCP.

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

- *HCP System Management Help* — This help system is a comprehensive guide to administering and using an HCP system. The Help describes

how to configure, manage, and maintain 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 describes how to configure, manage, and maintain 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 software-supported protocols to store, retrieve, and deleting objects, as well as changing object metadata such as retention and shred settings.
- *Installing an HCP System* — This book describes how to install the software for a new HCP system. It explains what you need to know to configure the system and contains step-by-step instructions for the installation procedure.
- *Deploying an HCP VM System on ESXi* — This book describes how 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 RAIN System - Final On-site Setup* — This book contains instructions for deploying an assembled and configured 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. The book also provides instructions for assembling the components of an HCP RAIN system that was ordered without a rack and for configuring Hitachi Remote Ops to monitor the nodes in an HCP system.

- *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.
- *Third-Party Copyrights* — This book contains copyright and license information for third-party software distributed with or embedded in HCP.
- *Using HCP Data Migrator* — This book describes how 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 Windows interface and the set of command-line tools included in HCP-DM.
- *HCP DM Third-Party Copyrights and Licenses* — This book contains copyright and license information for third-party software distributed with or embedded in HCP Data Migrator.

Accessing product documentation

Product 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 Portal](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.

[Hitachi Vantara Community](https://community.hitachivantara.com/s) 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/s, register, and complete your profile.



Note: If you purchased your Hitachi Content Platform from a third party, please contact your authorized service provider.

Comments

Please send us your comments on this document:

HCPDocumentationFeedback@HitachiVantara.com

Include the document title and part number, including the revision (for example, -01), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

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 directory structure. Namespaces are owned and managed by tenants.

HCP provides access to objects through a variety of industry-standard protocols as well as through various HCP interfaces.

HCP VM system components and architecture

This section describes the components and architecture of an HCP VM system.

The next figure describes 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. Only one HCP VM node should run on a 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 like 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 can be backed by SAN storage with RAID 6 protection, which is recommended, or by storage that's internal to the KVM hosts. With SAN storage, each KVM host is connected to the storage array through two Fibre Channel switches, which helps ensure data availability in the event of a single host failure.

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 sees as internal drives, out of the storage pools. The HCP operating system and software are installed on the virtual machine disks.

With storage that's 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.

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 physical NICs (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. pNIC-1 on all KVM hosts 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 require 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**) needs to 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. For more information, see [Chapter 7: "Configuring HCP monitoring with Hitachi Remote Ops"](#) on page 95.

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 SAIN 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 RAIN and SAIN 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 provides two terabytes of active and HCP S Series storage. The basic storage license also provides two terabytes of extended storage. If you need additional storage, contact your Hitachi Vantara sales representative.

For more information about storage licensing, see HCP System Management Help.

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 Hitachi Vantara Support Connect:.

<https://knowledge.hitachivantara.com/Documents>

HCP VM hardware requirements

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

Standard HCP VM system

To deploy a standard HCP VM system, you need:

- A shared SAN storage. RAID 6 system is recommended.
- 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 virtual disks 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 system

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 directory 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:

- 40 HCP VM nodes
- 59 data LUNs on each HCP VM node
- 15.90 TB virtual machine disks

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

- 16 HCP VM nodes
- 59 data LUNs on each HCP VM node
- 15.90 TB virtual machine disks

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.

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 in order 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.

Step 1: Log into the KVM host

To log into the KVM host:

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

Step 2: Verify that virtualization technology is enabled

To verify that virtualization technology is enabled:

1. Enter the following command to check if 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 the text, `kvm_intel` or `kvm_amd`, enter the following command to check if KVM is disabled by the BIOS:

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

If the output contains the following line of text: `kvm: disabled by bios`, virtualization technology needs to be enabled on the BIOS.

Step 3 (conditional): Connect a monitor and keyboard to your KVM host

Connect your monitor and keyboard to the KVM host.

Step 4 (conditional): Enter the BIOS

To enter the BIOS, you need to restart the KVM host and access the BIOS as the KVM host powers on.

To enter the BIOS:

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 start up screen, press the button that accesses the BIOS.

Step 5 (conditional): Enable virtualization technology

To enable virtualization technology:

1. Once you are in the BIOS, navigate to **Processor Settings**.
2. Press Enter.
3. Navigate to **Virtualization Technology**.
4. Press Enter to change **Virtualization Technology** to **Enabled**.
5. Press Escape until the **Exit** window appears.
6. Navigate to **Save changes and exit**.
7. Press Enter.

Installing the KVM packages

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

You need to install the following packages:

- **libvirt** — is a virtualization API and toolkit that manages virtualization hosts. Libvirt brings together every server side RPM required by libvirt.
- **libvirt-daemon** — is a server side daemon which 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** — is a user interface for performing administrator tasks on virtual machines.
- **guestfs-browser** — is a graphic interface for browsing the virtual machine file system and disk images.
- **libguestfs-tools** — is a set of tools for accessing and modifying virtual machine disk images.
- **python-libguestfs** — is a libguestfs tools Python library.
- **virt-top** — monitors KVM guest virtual machine CPU, memory, and performance.
- **virt-install** — is CLI command support for creating guest virtual machines for KVM.
- **bridge-utils** — is 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.

Step 1: Log into the KVM host

To log into the KVM host:

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

Step 2: Install the KVM packages

To install and initiate the KVM packages:

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

```
sudo dnf install libvirt libvirt-daemon-kvm qemu-kvm virt-manager guestfs-browser  
libguestfs-tools python-libguestfs virt-top virt-install bridge-utils virt-viewer
```

2. Once 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
```

Step 3: Install Virtual Machine Manager on your local computer

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

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. Once 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
```


Configuring KVM networking

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

Resources you need before you start

To configure networking for a KVM host, you need:

- A USB keyboard
- A VGA monitor

Configuring the KVM host for networking

To configure the KVM host for networking, you need to create bond files that bond the two physical front-end ports together and bond the two physical back-end ports together. Then you need to 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.

In the examples used in this section, 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.

Step 1 (conditional): Connect a monitor and keyboard to your KVM host

Connect your monitor and keyboard to the KVM host.

Step 2: Back 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 the interface configuration files are named `ifcfg-eno1`, `ifcfg-eno2`, `ifcfg-eno3`, `ifcfg-eno4`.

To back up the interface configuration files:

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 directory:

```
ls
```

3. Verify that the directory 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 directory:

```
ls /root
```

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

Step 3: Enable the bonding kernel module

The bonding kernel module needs to be configured so that it's enabled every time the KVM host is powered on.

To configure the KVM host to enable the bonding kernel module when its powered on:

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` directory:

```
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 the *Escape* key.
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
```

Step 4: Create the bridge network files

You need to 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 back-end bridge network file is named `ifcfg-back-end`.

To create the bridge network files:

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 bridge 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 the *Escape* key.

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

```
:wq
```

7. Enter the following command to create the back-end bridge network file:

```
vi ifcfg-back-end
```

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:

```

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

```

10. Press the *Escape* key.**11.** Enter the following command to save and exit the file:

```
:wq
```

Step 5: Create the bond network files

You need to 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 back-end bond network file is named `ifcfg-bond1`.

To create the bond network files:

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=en0 mode=active-backup miimon=100
updelay=3000 downdelay=500"
BRIDGE="front-end"
```

5. Press the *Escape* key.
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-File2\
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=en02 mode=active-backup miimon=100
updelay=3000 downdelay=500"
BRIDGE="back-end"

```

10. Press the *Escape* key.**11.** Enter the following command to save and exit the file:

```
:wq
```

Step 6: Enslave the front-end interface configuration files

To bridge the front-end network, you need to enslave the front-end interface configuration files. In this example, the first front-end network configuration file is named `ifcfg-eno1` and second front-end network configuration file is named `ifcfg-eno3`.

To enslave the interface network configuration files:

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 the *Escape* key.
 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 the *Escape* key.

**11.** Enter the following command to save and exit the file:

```
:wq
```

## Step 7: Enslave the back-end interface configuration files

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

To enslave the interface network configuration files:

**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 the *Escape* key.

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 the *Escape* key.
11. Enter the following command to save and exit the file:

```
:wq
```

## Step 8: Restart networking on the KVM host

Once 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 performed the network configuration, you may experience connection issues after the network restarts.

## Step 9: Restart the individual KVM host

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

To restart the KVM host:

1. Press and hold the power button until the KVM host shuts down.
2. Press the power button again and wait for KVM host to restart.

The KVM host reboots. When the KVM host has finished rebooting, it is available for access.



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

```
sudo systemctl restart network
```

## Step 10: Verify 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.

To verify that the network has been updated:

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

```
ip link
```

2. 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>
```

3. 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
```

4. 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
```





# 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 VM system

To deploy the HCP VM system, you need to 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.

Once the HCP VM connections have been added to Virtual Machine Manager, you can create the HCP VM nodes using the 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 command line to deploy an HCP VM, see [Appendix B: "Creating an HCP VM node using the command line"](#) on page 113

Before you deploy an HCP VM system:

- Install the KVM packages (See [Chapter 3: "Installing KVM"](#) on page 11

- Create network bridges (See [Chapter 4: "Configuring KVM networking"](#) on page 17)
- Copy and unzip a .iso file on your HCP node (See [Step 2: "Copy the .iso file to the KVM host"](#) below)

## Step 1: Download the .iso installation file

To download the .iso files, go to the HCP distributor download site and download the HCP software installation file (HS222\_*release-number*.iso.zip)

## Step 2: Copy the .iso file to the KVM host

You need to use SCP to copy and send the HS222\_*release-number*.iso.zip file to the KVM host.

To move the HS222\_*release-number*.iso.zip file to the KVM host:

1. Enter the following command to navigate to the directory on your local computer where you saved the 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` directory 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 directory with the saved .iso file:

```
cd /var/iso
```

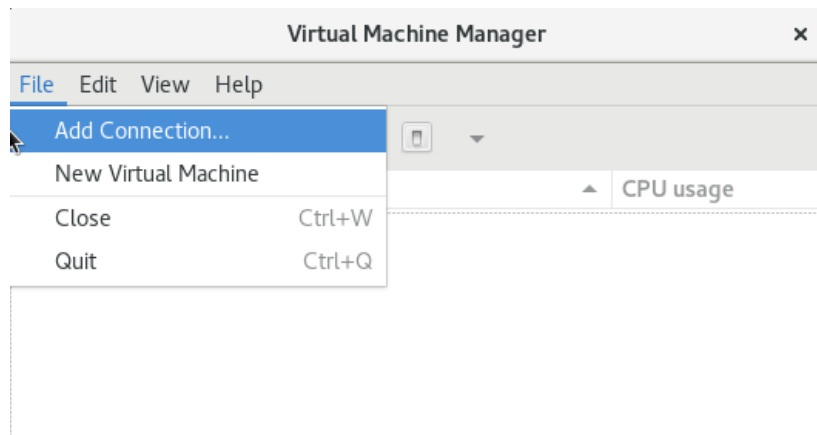
5. Enter the following command to unzip the file:

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

## Step 3: Add an HCP VM node connection

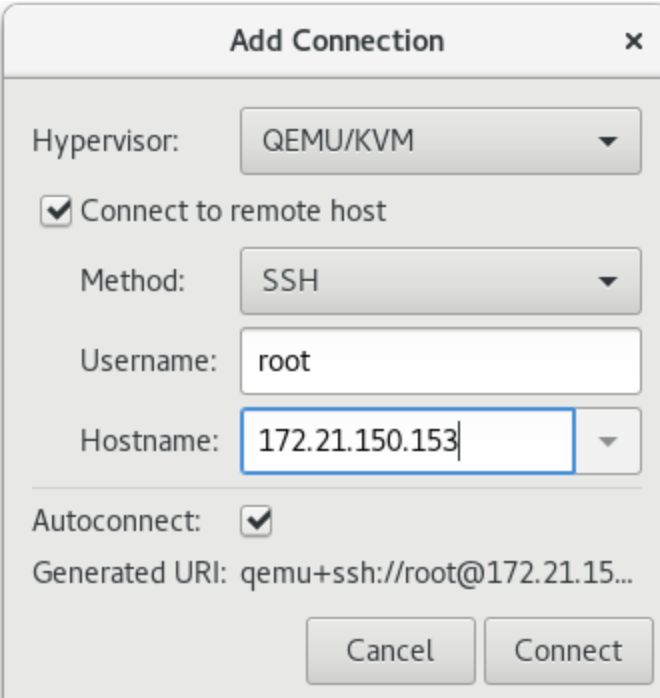
To add an HCP VM node connection:

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



3. To configure the **Add Connection** window:
  - In the **Hypervisor** drop down menu, select **QEMU/KVM**.
  - Select **Connect to remote host**.
  - In the **Method** drop down menu, select **SSH**.
  - In the **Username** field, type the username for the KVM host.
  - In the **Hostname** field, type the hostname of the KVM host or its front-end IP address.

- Select **Autoconnect**.



**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

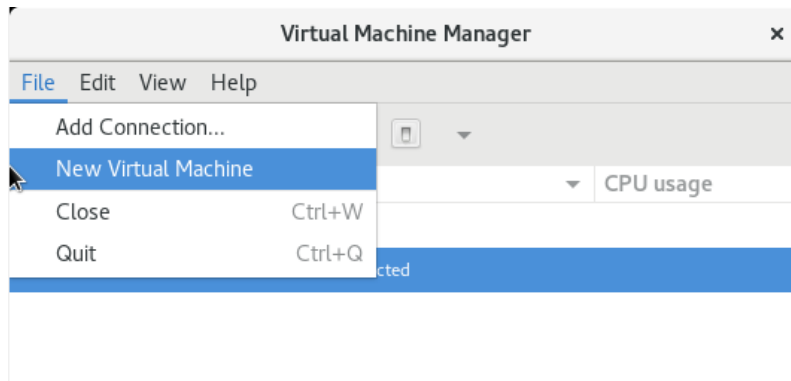
4. Click **Connect**.
5. In the **OpenSSH** window, enter the password for your KVM host.
6. Click **OK**.

## Step 4: Create the HCP VM node

To create an HCP VM node:

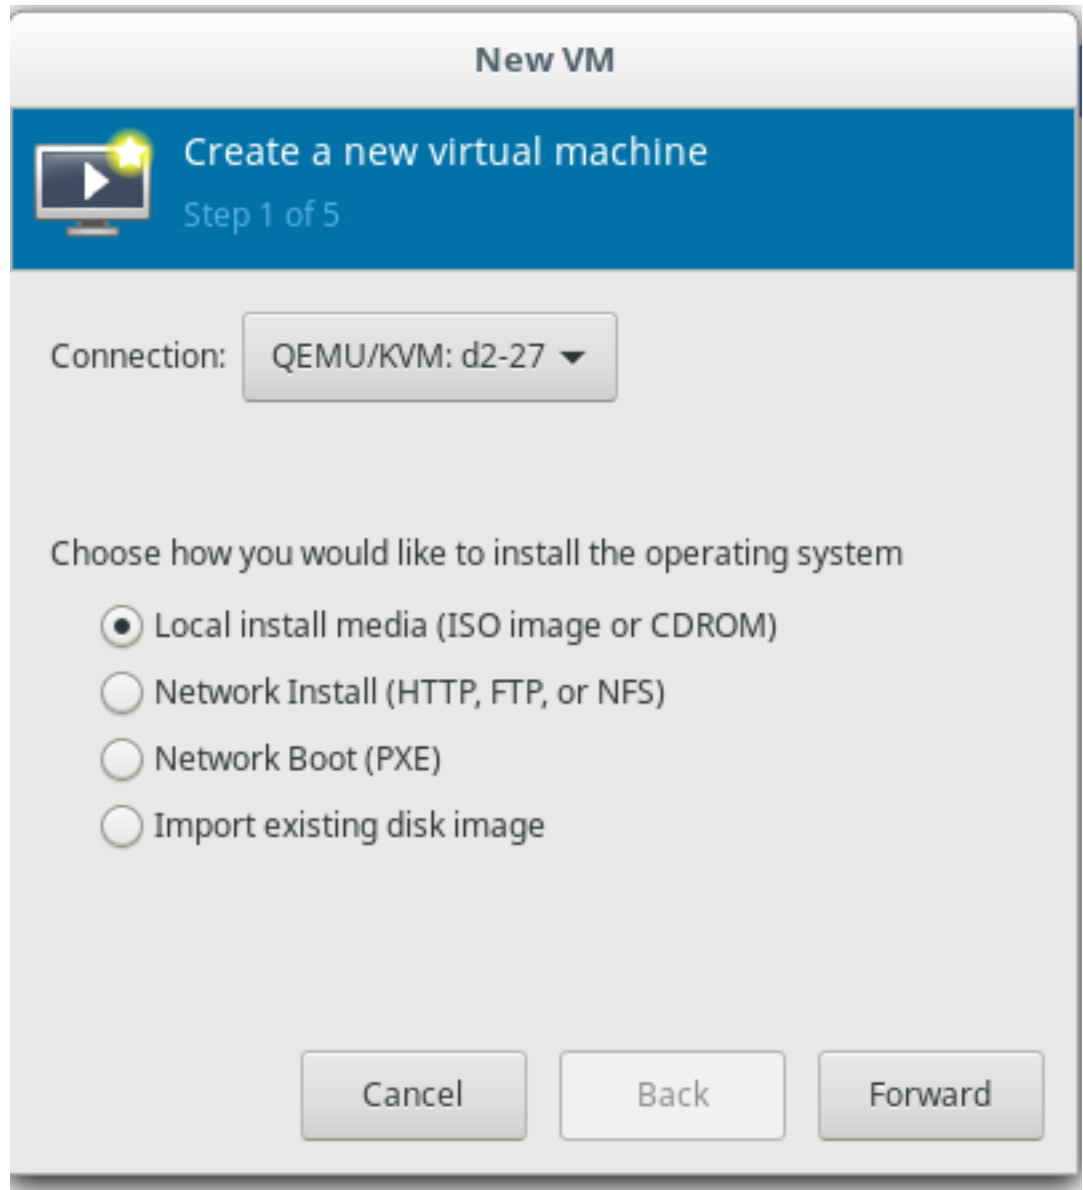
1. In **Virtual Machine Manager**, highlight the newly added connection and click **File ► New Virtual Machine**.

2. In the drop down 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:
  - Select **Use ISO image** and click **Browse**.
  - Navigate to the `HS222_release-number.iso` file on the KVM host and click **OK**.
  - In the **OS type** field, select **Linux**.

- 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 ▼

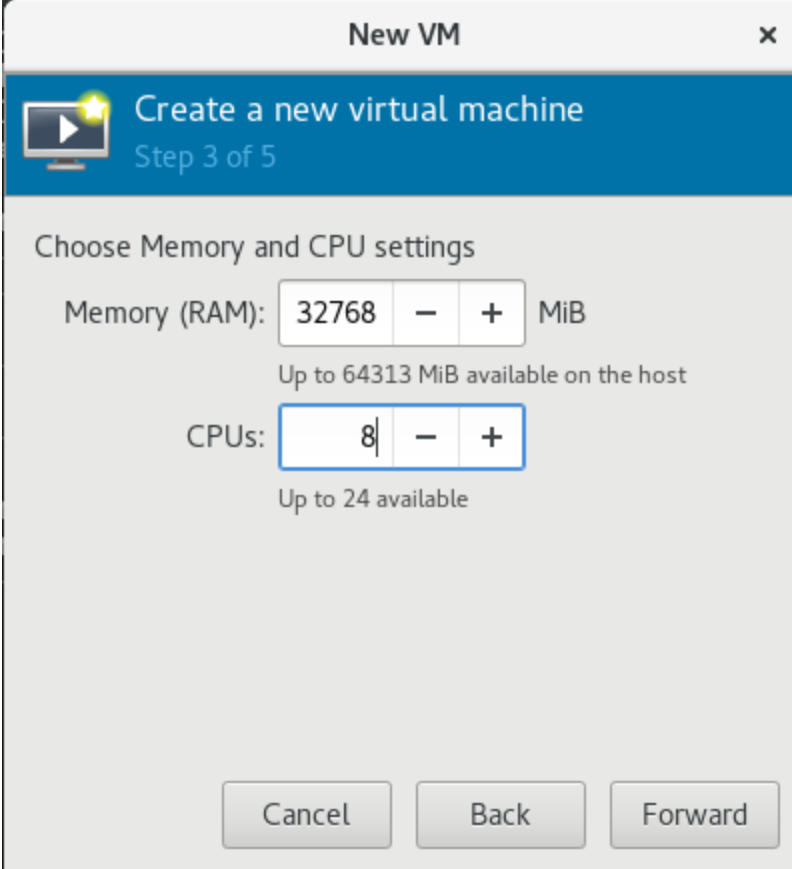
Cancel Back Forward

6. Click **Forward**.

7. To configure the **New VM Step 3 of 5** window:

- 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. For more information about HCP VM configurations, see ["HCP VM hardware requirements"](#) on page 7

- In the **CPUs** field, type 8 for a standard HCP VM configuration or type 4 for a small instance HCP VM configuration.

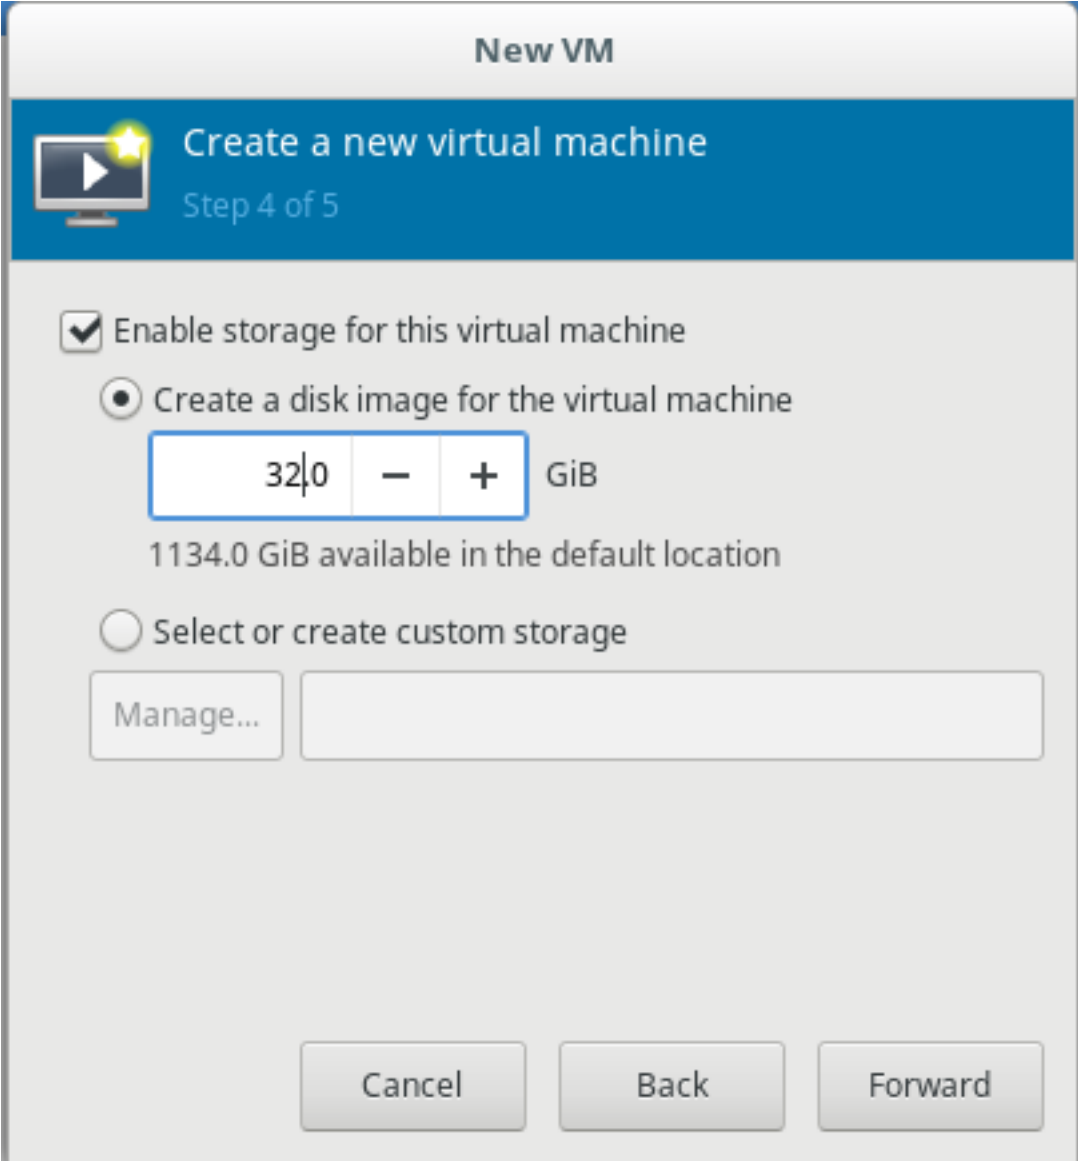


The screenshot shows a window titled "New VM" with a close button (X) in the top right corner. Below the title bar is a blue header with a play button icon and the text "Create a new virtual machine" and "Step 3 of 5". The main content area is titled "Choose Memory and CPU settings". It contains two settings: "Memory (RAM):" with a value of "32768" and "MiB" unit, and "CPUs:" with a value of "8". Both settings have minus and plus buttons for adjustment. Below the memory setting, it says "Up to 64313 MiB available on the host". Below the CPU setting, it says "Up to 24 available". At the bottom of the window are three buttons: "Cancel", "Back", and "Forward".


8. Click **Forward**.
9. To configure the **New VM Step 4 of 5** window, perform either of these steps:
  - If you are creating a disk image on the virtual machine:
    1. Select **Enable storage for this virtual machine**.
    2. Select **Create a disk image for the virtual machine**.



3. 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

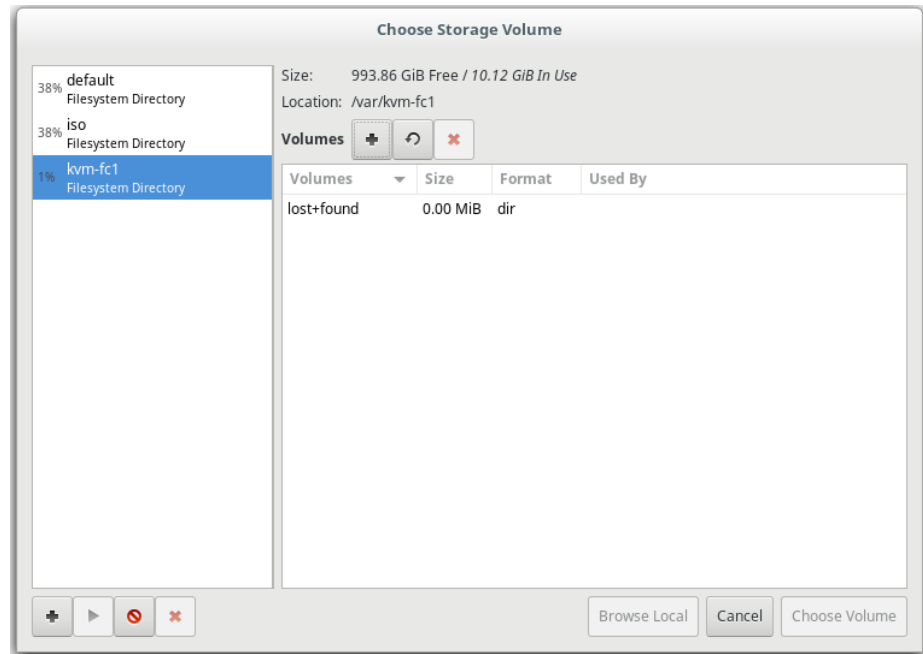
☐ Select or create custom storage

Manage...

Cancel Back Forward

- If you are using SAN storage and want to create a custom storage location:
  1. Select **Select or create custom storage**.
  2. Click **Manage**.
  3. In the **Choose Storage Volume** window that opens, in the left hand panel, select the storage pool you want to host your storage volume.

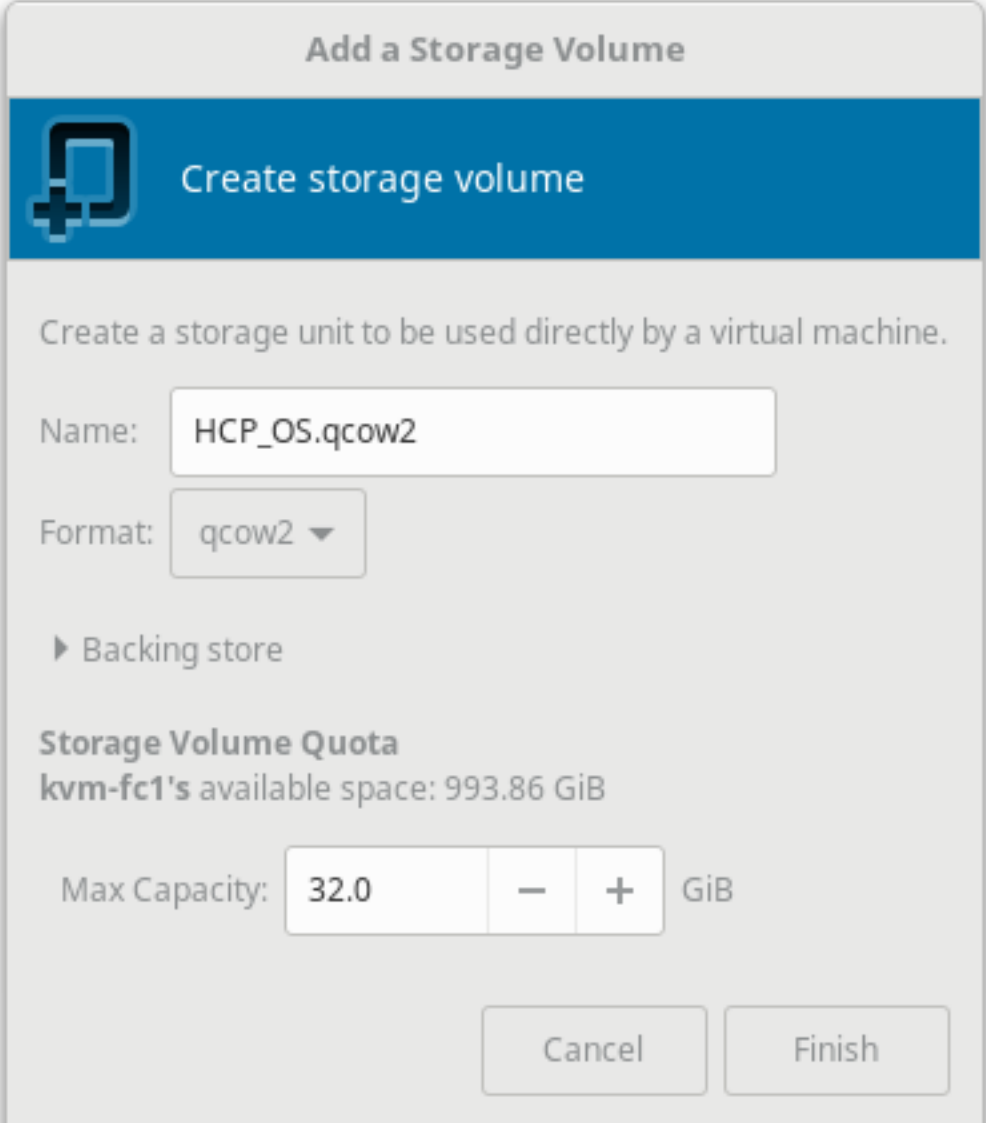
4. Click the **Volumes:** plus sign.



5. To configure the **Add a Storage Volume** window that opens:

- In the **Name** field, type a name for your storage volume.
- In the **Format** field, select **qcow2**.
- In the **Create a disk image for the virtual machine** field, enter a storage value of 32 or more.

- 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

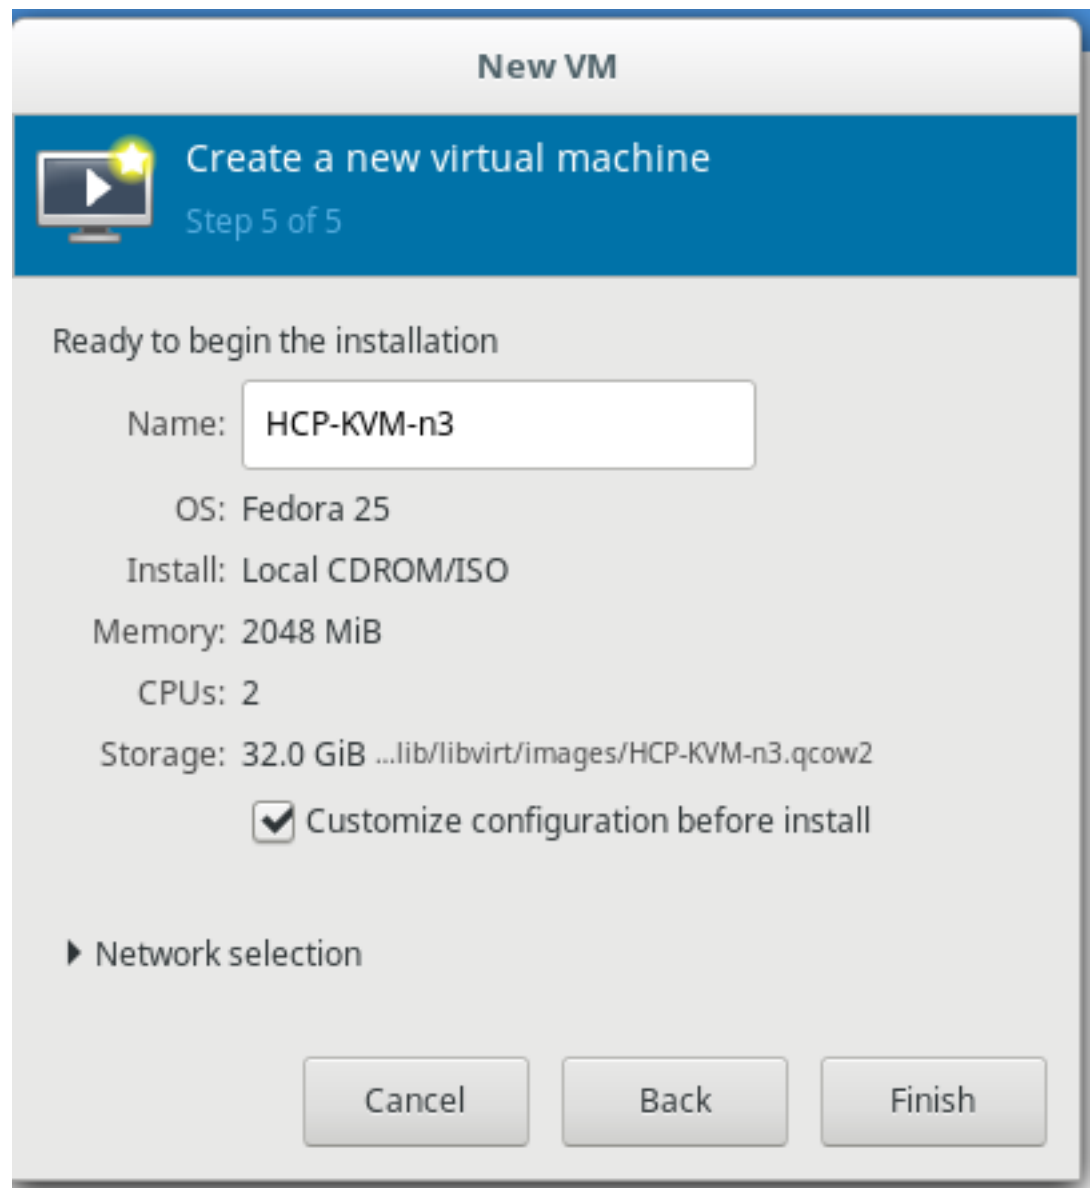
Max Capacity:    GiB

**10.** Click **Forward**.

**11.** To configure the **New VM Step 5 of 5** window:

- In the **Name** field, type the name of your HCP VM.

- Select **Customize configuration before install**.



12. Click **Finish**.

Once you click **Finish**, the customization window opens.

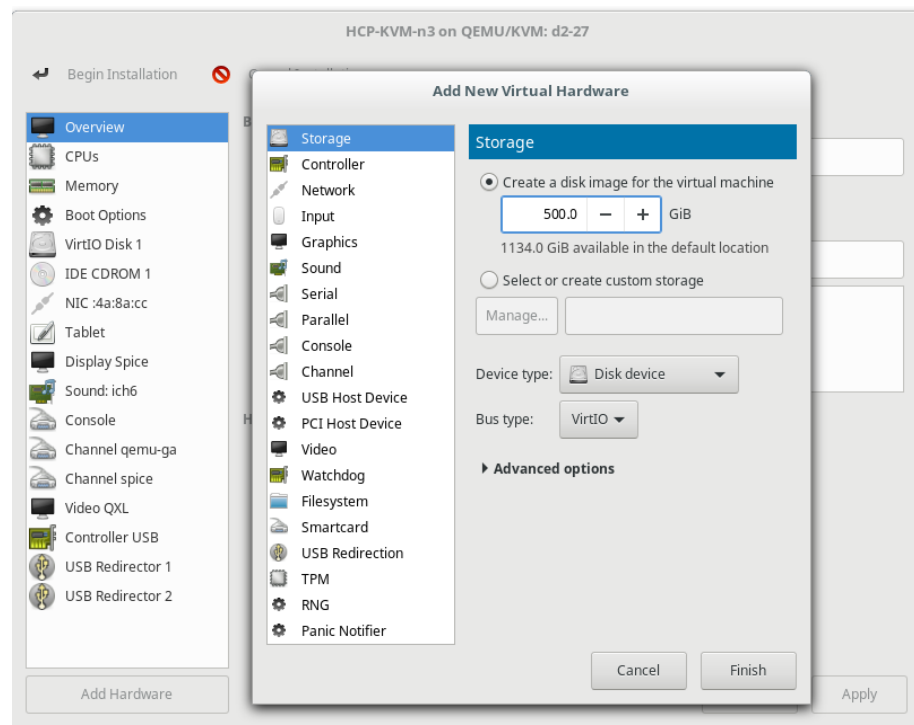
## Step 5: Customize the HCP VM node for HCP configuration

To customize the HCP VM node for HCP configuration:

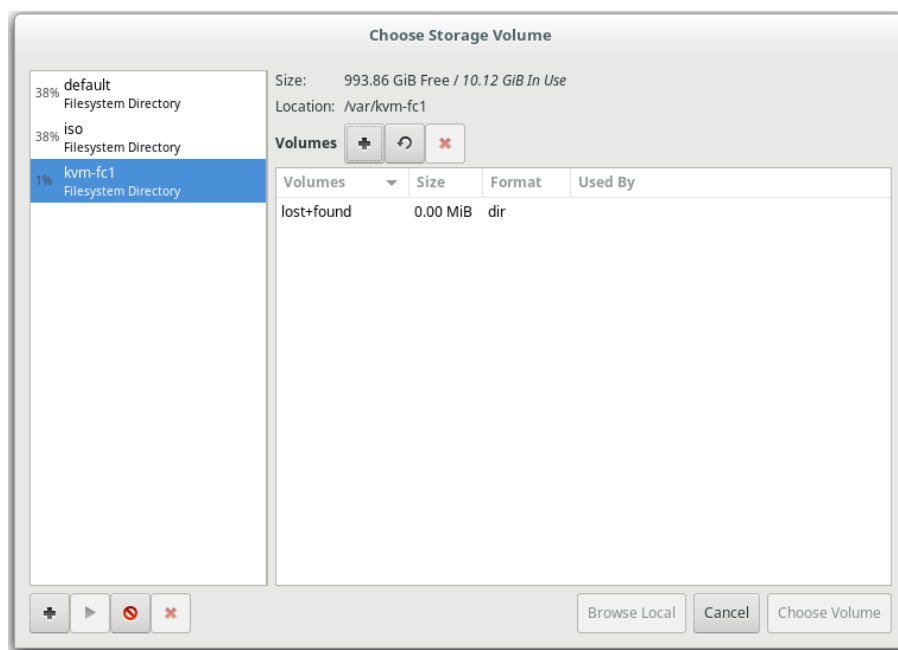
1. In the customization window, on the left hand navigation pane, right-click the **Overview** tab, and in the drop down menu, click **Add Hardware**.

The **Add New Virtual Hardware** window opens.

2. In the **Add New Virtual Hardware**, on the left hand navigation pane, click the **Storage** tab.
3. To configure the **Storage Volume** window, perform one of these steps:
  - If you are creating a disk image on the virtual machine:
    1. Select **Create a disk image for the virtual machine**.
    2. In the **Create a disk image for the virtual machine** field, enter a storage value of *500* or more.
    3. In the **Device type** field, select **Disk device**.
    4. In the **BUS type** field, select **VirtIO**.



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



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

- Click **Finish**.

**Add a Storage Volume**

Create storage volume

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

Name: HCP\_DATA\_01.qcow2

Format: qcow2 ▼

► Backing store

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

Max Capacity: 500.0 — + GiB

Cancel Finish

4. Click **Forward**.
5. To configure the **Add New Virtual Hardware** window:
  - In the **Device type** field, select **Disk device**.
  - In the **BUS type** field, select **VirtIO**.
6. Click **Finish**.

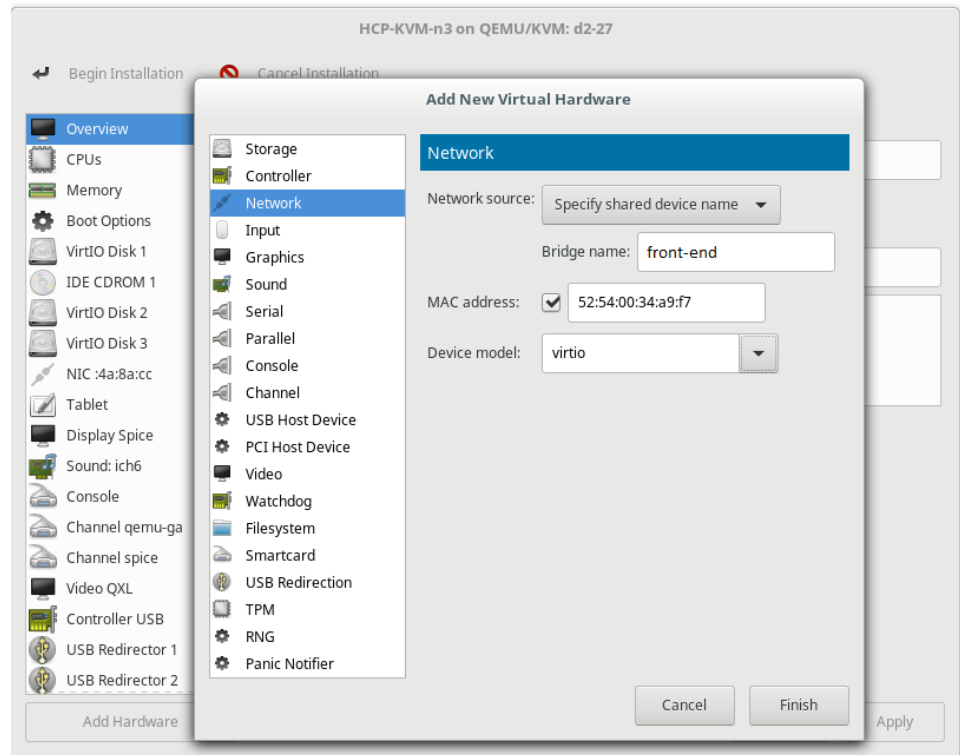
7. Once you return to the customization window, repeat steps 1 through 6 to create a second storage volume.
8. After creating two storage volumes, on the left hand navigation pane of the customization window, right-click the existing NIC and, from the drop down menu, click **Delete**.
9. On the left hand navigation pane of the customization window, right-click the **Overview** tab and, in the drop down menu, click **Add Hardware**.
10. To configure the **Add New Virtual Hardware** window that opens:
  - On the left hand navigation pane, 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`. For more information about the network bridge files, see [Step 4: "Create the bridge network files"](#) on page 19

- Select **MAC address** and leave the default MAC address.



- In the **Device model** field, select **virtio**.



**11.** Click **Finish**.

**12.** Once 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.

Following the example used in the networking chapter, the bridge file name for your back-end network will be `back-end`. For more information about the network bridge files, see [Step 4: "Create the bridge network files"](#) on page 19



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

**13.** Click **Begin Installation**.

Once you begin the installation, you are asked to configure the operating system.

## Step 6: Perform the OS installation

To install the Appliance OS on a node:

1. In **Virtual Machine Manager**, double click the new HCP VM node.
2. Enter the KVM host password.
3. Once the virtual machine starts up, 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 [pcel]: _
```

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 will be 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` if you want to use VLAN ID or enter `n` to indicate that 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 above, 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 IP address for the HCP VM.

```
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 above, you are finished entering front-end network configuration information for the node. Skip the next step in this procedure, and go to step 11. If you entered *IPv6* or *Dual* proceed to step 10.

10. If you entered *IPv6* or *Dual* in response to the prompt in step 6 above, 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 primary IPv6 address prefix length for the front-end network.

```
Enter the front-end IPv6 prefix length [64]:
--->
```

- c. When prompted, enter the primary IPv6 gateway IP address for the front-end network.

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

- d. When prompted, enter *y* to indicate that you want to assign a secondary IPv6 address to the node for the front-end network or *n* to indicate that 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 step d above, 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:
  - To confirm your responses, enter *y*.
  - To change any of your responses, enter *n*. 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 need to 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.

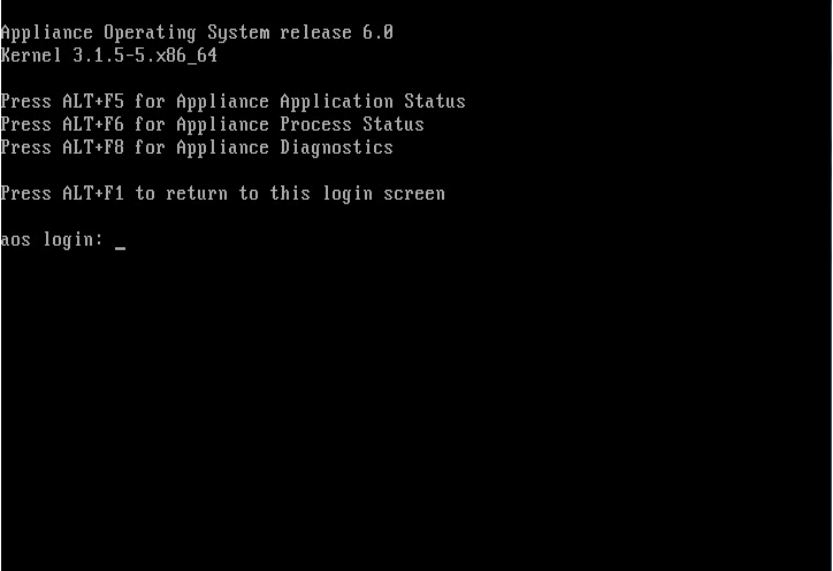
When the installation is complete, the HCP VM node shuts down. If it does not reboot, power on the HCP VM node using Virtual Machine Manager.

Once the HCP VM node reboots, you are asked to change the password for the install user.

## Step 7: Change the install user password

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

1. Login to the HCP VM node console with the default login information:
  - Username: *install*
  - Password: *Chang3Me!*

A screenshot of a terminal window showing the HCP VM node console. The text displayed is: "Appliance Operating System release 6.0", "Kernel 3.1.5-5.x86\_64", "Press ALT+F5 for Appliance Application Status", "Press ALT+F6 for Appliance Process Status", "Press ALT+F8 for Appliance Diagnostics", "Press ALT+F1 to return to this login screen", and "aos login: \_".

```
Appliance Operating System release 6.0
Kernel 3.1.5-5.x86_64

Press ALT+F5 for Appliance Application Status
Press ALT+F6 for Appliance Process Status
Press ALT+F8 for Appliance Diagnostics

Press ALT+F1 to return to this login screen

aos login: _
```

2. Change the password to *hcpinsta11* (The last two characters are the number 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*, you would 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.** Log in to the HCP VM node console with the default login information:

- Username: `install`
- Password: `Chang3Me!`

**4.** Change the password to `hcpinsta11` (the last two characters are the number one).

**5.** Press Enter to display the HCP Configuration menu.

The next screen is included for illustrative purposes. Your screen will specify the HCP version that you are running.

```
HCP 8.1 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: 8.1.0.1
Version on CD/DVD: None
Extracted version: 8.1.0.1

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
[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]:
```

## Step 1: Identify the nodes in the HCP VM system

To identify the nodes in the HCP-Vm system:

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.

When you enter *y* or *yes*, 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** menu appears.

```
HCP Nodes Menu
=====

[1] Storage Node Back-end IP Addresses
[b] Go Back to the Previous Menu
[q] Return to 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 presents a default value that's 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] Encryption Settings

[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]: _
```

## Step 2: Configure the HCP system

### Procedure

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

```
Distributor/DEM Key Access
=====

Please enter a valid distributor key for your company (supplied by HDS).
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
HDS-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]: _
```

5. Enter **1** and change the **Gateway router IP address**.

6. Enter **2** and change the **Multicast Network**.

7. Enter **b**.

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

8. Enter **4** to configure the DNS options.

```
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]: _
```

9. Enter **2** to input the domain name for the system.

10. Enter the system domain name.

```
Domain Name 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-vm-1.wilco.net

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

11. If **Option 1: Enable DNS** is not set to yes, change it to yes.
12. If **Option 3: DNS Servers** is not set to the proper corporate DNS server, change it accordingly.

13. Enter **bl**

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): 192.43.244.18

Note: Control-C cancels input.

Internal or time server name or IP address.
[Default: internal]: 64.90.182.55

You chose: "64.90.182.55", 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.

When you enter **6**, 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 (No)
[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, if you want to configure a dedicated database volume, the system needs to have at least three drives per node. Also, the dedicated database volume size needs to be at least 50 GB for a new installation. All dedicated database volumes need to be the same size. HCP Setup asks if you want to configure a dedicated database volume only if your system meets the above requirements.

22. If HCP Setup asks whether you want to configure a dedicated database volume, enter `yes` if you want to configure a dedicated database volume or `no` if you do not want to configure a dedicated database volume.
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 valid a 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 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 required to license the system. Omitting the serial number will cause the system to report that you are in violation of your license agreement.

---

27. From the **Internal Configuration Settings** menu, enter **3** to configure whether replication will be 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 will 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 will be enabled.
29. From the **Internal Configuration Settings** menu, enter **5** to set contact information. To specify no contact information, press the spacebar key.
30. Enter **b** to return to the **New Install** menu.

## Step 3: Run 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 running 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 will 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) = 172.18.4.45
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
MQE 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 next 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. Conditional:** If you previously selected that you want to configure dedicated database volumes, complete the next substeps.
  - a.** When prompted, select the dedicated database volume for the first node.
  - b.** Press Enter to confirm your selection.
  - c.** Repeat the above two substeps 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. It then prompts you to enter the key.



**Important:** Before entering the encryption key, record it. Once 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 reboots 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 reboot, and, instead of the operating system login prompt, you should see an **hcp-node-nodeNumber** prompt.

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

```
Every 30.0s: /sbin/system-info Fri Mar 1 12:29:58 2013
Hostname: hcp-node-150.cluster-colo-009-vm1.lab.archivas.com
NIS Node: 150
[hcp_system] IP: 172.20.27.150
[hcp_system] Mask: 255.255.255.0
[hcp_system] Gateway: 172.20.27.254
[hcp_backend] IP: 172.21.150.150
[hcp_backend] Mask: 255.255.255.0
Version: 6.0.0.93
Operating System: OS 6.0.0.514
Linux Kernel: 3.1.5-5.x86_64
Current Run Level: 4

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

## Step 4: Verify 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 on.

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 the nodes have completed the startup process, the Console displays the **Change Password** page.
  - 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.

Password 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 and 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 may 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 [Chapter 7: "Configuring HCP monitoring with Hitachi Remote Ops"](#) on page 95

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

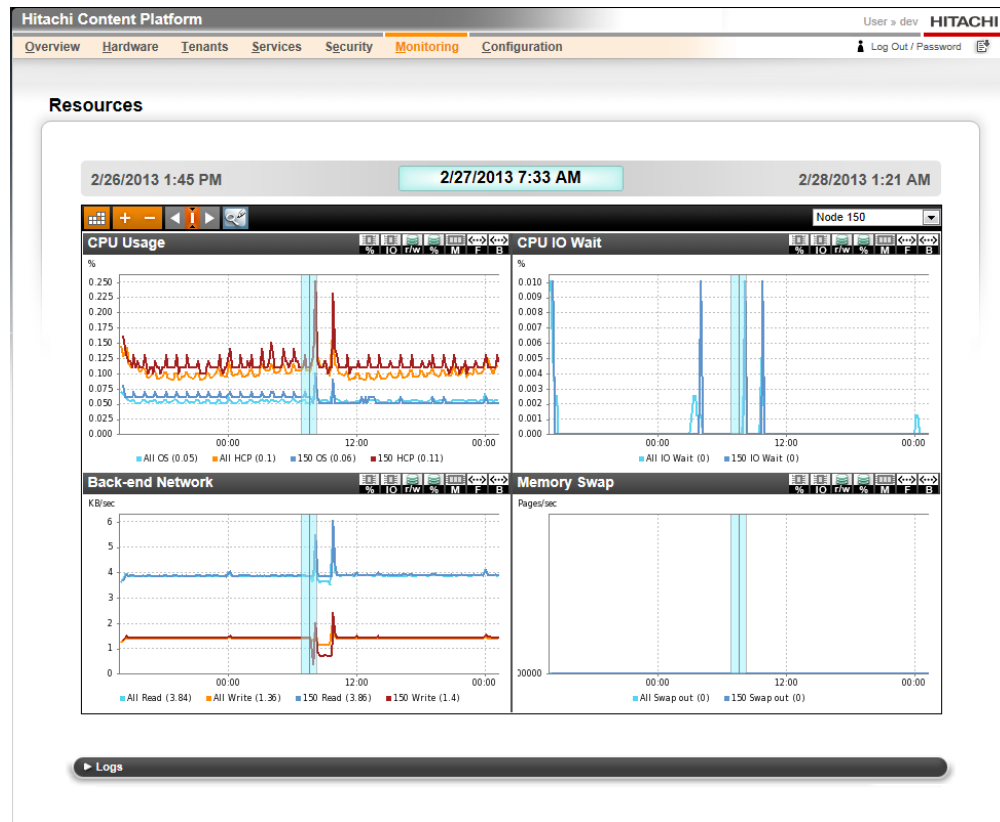
HCP maintains a system log that logs all system events. You can view this log in the HCP System Management Console. You can send system log messages to syslog servers, System Network Management Protocol (**SNMP**) managers, and email addresses. Additionally, you can use SNMP to view and, when allowed, 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. For more information, see [Chapter 7: "Configuring HCP monitoring with Hitachi Remote Ops"](#) on page 95

## 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 ten 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 if 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.



# Maintenance procedures

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 screen is included for illustrative purposes. Your screen will specify the HCP version that you are running.

```
HCP 8.1 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: 8.1.0.1
Version on CD/DVD: None
Extracted version: 8.1.0.1

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): 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.

On an HCP node, one new volume can be specified as the dedicated database volume.
All database files are moved to this dedicated volume. If the HCP system
already has dedicated database volumes, you can still add new volumes and
specify a new dedicated database volume for each HCP node.

Are you sure you want to continue?
(Default: no): yes

You chose: "yes", is this 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 003:
 1. /dev/sdd at 2:0:0:3 (500GB)

 node 004:
 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 would show storage added to all nodes.

Optional: If you want to configure a dedicated database volume, the volume size needs to be at least 50 GB and at least 1.5 times the size of the existing database for each node. All dedicated database volumes need to 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:
  - 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 1 11:51:59 2013 Current status:
node 150: 53% Complete (7/13): Running formatDrives
Fri Mar 1 11:52:15 2013 Current status:
node 150: 53% Complete (7/13): Formatting 27% complete (0 / 1 volumes)
Fri Mar 1 11:53:15 2013 Current status:
node 150: 53% Complete (7/13): Formatting 35% complete (0 / 1 volumes)
Fri Mar 1 11:54:15 2013 Current status:
node 150: 53% Complete (7/13): Formatting 42% complete (0 / 1 volumes)
Fri Mar 1 11:55:15 2013 Current status:
node 150: 53% Complete (7/13): Formatting 99% complete (0 / 1 volumes)
Fri Mar 1 11:56:15 2013 Current status:
node 150: 53% Complete (7/13): Formatting 100% complete (1 / 1 volumes)
Fri Mar 1 11:56:25 2013 Current status:
node 150: 76% Complete (10/13): Running create_volume_config
Fri Mar 1 11:56:30 2013 Current status:
node 150: 84% Complete (11/13): Running start_new_volumes
Fri Mar 1 11:56:46 2013 Current status:
node 150: 92% Complete (12/13): Running sync_new_local_volumes
Fri Mar 1 11:57:01 2013 Current status:
node 150: 100% Complete: Storage addition complete
Fri Mar 1 11:57:22 2013 Current status:
node 150: 100% Complete: Starting new volumes
Fri Mar 1 11:57:27 2013 Current status:
node 150: 100% Complete: Starting new volumes

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

```

- If you want to select a dedicated volume for the database, enter yes. Then complete the next substeps. , then:
  - a. If you want to select a dedicated database volume for the first node, enter yes.
  - b. If you entered yes, select the dedicated database volume for the first node.
  - c. Press Enter to confirm your selection.
  - d. Repeat the above three steps for each node in the system.
- 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 1 11:51:59 2013 Current status:
 node 150: 53% Complete (7/13): Running formatDrives
Fri Mar 1 11:52:15 2013 Current status:
 node 150: 53% Complete (7/13): Formatting 27% complete (0 / 1 volumes)
Fri Mar 1 11:53:15 2013 Current status:
 node 150: 53% Complete (7/13): Formatting 35% complete (0 / 1 volumes)
Fri Mar 1 11:54:15 2013 Current status:
 node 150: 53% Complete (7/13): Formatting 42% complete (0 / 1 volumes)
Fri Mar 1 11:55:15 2013 Current status:
 node 150: 53% Complete (7/13): Formatting 99% complete (0 / 1 volumes)
Fri Mar 1 11:56:15 2013 Current status:
 node 150: 53% Complete (7/13): Formatting 100% complete (1 / 1 volumes)
Fri Mar 1 11:56:25 2013 Current status:
 node 150: 76% Complete (10/13): Running create_volume_config
Fri Mar 1 11:56:30 2013 Current status:
 node 150: 84% Complete (11/13): Running start_new_volumes
Fri Mar 1 11:56:46 2013 Current status:
 node 150: 92% Complete (12/13): Running sync_new_local_volumes
Fri Mar 1 11:57:01 2013 Current status:
 node 150: 100% Complete: Storage addition complete
Fri Mar 1 11:57:22 2013 Current status:
 node 150: 100% Complete: Starting new volumes
Fri Mar 1 11:57:27 2013 Current status:
 node 150: 100% Complete: Starting new volumes

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

```

- 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`, HCP 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/100G) | (180G/200G)
172.20.59.125 | pgxlog | /RIS/archive34 | /RIS/archive95
 | | (450M/100G) | (180G/200G)
172.20.59.129 | pgxlog | /RIS/archive33 | /RIS/archive94
 | | (350M/100G) | (150G/200G)

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

Once the procedure is initiated, 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 12 09:51:15 2017 Current status:
 node 042: 40% Complete (2/5): Running arcShutdown
Thu Jan 12 09:52:13 2017 Current status:
 node 042: 60% Complete (3/5): Running mountDisks
Thu Jan 12 09:52:48 2017 Current status:
 node 042: 80% Complete (4/5): Running move_pgdata
Thu Jan 12 09:52:55 2017 Current status:
 node 042: 100% Complete: Deploy complete
Thu Jan 12 09:54:07 2017 Current status:
 node 042: 100% Complete: Rebooting node
Thu Jan 12 09:54:12 2017 Current status:
 node 042: 100% Complete: Waiting for node to become available.
Thu Jan 12 09:56:12 2017 Current status:
 node 042: 100% Complete: Waiting for node to become available.
Thu Jan 12 09:58:13 2017 Current status:
 node 042: 100% Complete: Waiting for node to become available.
Thu Jan 12 10:00:03 2017 Current status:
 node 042: 100% Complete: All nodes available
Thu Jan 12 10:00:38 2017 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 your entry 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 your entry or **n** or **no** to try again.

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

```
Manage Database Volumes
=====

[d] 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 your entry 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 your entry or **n** or **no** to try again.

Once 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. For more information, see [Chapter 3: "Installing KVM"](#) on page 11.
2. Unpack and upload the ISO files to the selected KVM hosts. For more information, see [Chapter 5: "Deploying the HCP VM system"](#) on page 31.
3. Create the new virtual machine. For more information, see [Chapter 5: "Deploying the HCP VM system"](#) on page 31.
4. Configure the HCP VM network on the newly deployed HCP VM nodes. For more information, see [Step 6: "Perform the OS installation"](#) on page 48.
5. From the highest active HCP VM node, run the **Add Node** service procedure. For more information, see *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 your entry, or `n` or `no` to try again.

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

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

[l1] Add Storage Nodes to the System (no updates)
[lv] Review Updated Configuration (disabled, no updates)
[xx] Add Nodes to an Existing HCP System (disabled, no updates)
[lq] Return to Configuration Menu

Enter your choice.
[Default: 11:
```

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 your entry, or `n` or `no` to back out.

The wizard prompts again for confirmation.

5. In response to the confirming prompt, enter `y` or `yes` to confirm your entry, 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) = 172.20.59.46
Allow Data at Rest Encryption = Yes
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
MQE 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.

- c. Select dedicated database volumes.

1. When prompted, select the dedicated database volume for the first node.
2. Press Enter to confirm your selection.
3. 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]?
...
```

- d. 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 rebooted. 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.

7. From the **HCP Configuration** menu, enter **q** to log out of the install shell.
8. In response to the confirming prompt, enter **y** or **yes** to confirm your entry, or **n** or **no** to try again.

## Recovering storage nodes

The following sections describes 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, refer to the *Installing and Maintaining an HCP System* manual.

### 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 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 **1** for recovery operations.
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*, 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 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. Be 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 your entry 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. In this case, fix the problem and then start the OS recovery procedure again.

When the node recovery is complete, HCP reboots 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 is automatically terminated when HCP reboots 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 your entry 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 your entry 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 your entry 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 operations.
4. In response to the confirming prompt, enter *y* or *yes* to confirm your entry 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 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. Be 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.

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 targetted 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. Conditional: 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 above two steps 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 the logical volumes on the selected nodes or all nodes, as applicable. If any of the prechecks fail, HCP Setup exits. In this case, fix the problem and then start the node recovery procedure again.

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



---

**Important:** 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 your entry 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 your entry or *n* or *no* to try again.





# 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 need to 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 need to configure a new SNMP user in Hitachi Remote Ops to match what 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, select **Monitoring ► SNMP**.
3. In the **SNMP Settings** section on the **SNMP** page:
  - Select the **Enable SNMP at `snmp.hcp-domain-name`** option.
  - Select either **Use version 1 or 2c** (recommended) or **Use version 3**.  
 If you select **Use version 3**, specify a username and password in the **Username**, **Password**, and **Confirm Password** fields.
  - 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, follow the steps outlined in the next table.

| Step | Activity                     | More information                                                          |
|------|------------------------------|---------------------------------------------------------------------------|
| 1    | Log into Hitachi Remote Ops. | <a href="#">Step 1: "Log in to Hitachi Remote Ops"</a> on the facing page |

| Step | Activity                                                                                                                           | More information                                                              |
|------|------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| 2    | Set the Hitachi Remote Ops base configuration, including the email addresses to which email about error conditions should be sent. | <a href="#">Step 2: "Set the base configuration"</a> below                    |
| 3    | Optional: Configure transport agents for reporting error conditions to Hitachi Vantara support personnel.                          | <a href="#">Step 3: "(Conditional) Configure transport agents"</a> on page 99 |
| 4    | Identify the HCP system to be monitored.                                                                                           | <a href="#">Step 4: "Identify the HCP system"</a> on page 99                  |

## Step 1: 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 server by using either the hostname or a valid IP address for the server, followed by port number 6696. For example:

`http://hitrack: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 **Logon**.

## Step 2: Set the base configuration

The Hitachi Remote Ops base configuration specifies information including the customer site ID, how frequently to scan devices, whether to report communication errors that occur between Hitachi Remote Ops and monitored devices, and the email addresses to send the error condition alerts to.

If Hitachi Remote Ops is already used at your site, the base configuration might already be set up. You can leave the configuration as is, or you can update it to accommodate additional HCP devices.

### Procedure

1. In the row of tabs at the top of the Hitachi Remote Ops interface, click **Configuration**.

The **Base** page is displayed.

2. In the **Device Monitoring** section:

- In the **Site ID** field, type your Hitachi Vantara customer ID. If you do not know your customer ID, contact your authorized HCP service provider.
- Optional: Update the field values to meet your site needs. For information about these fields, click **Help on this table's entries**.

3. In the **Notify Users by Email** section:

- 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.
- 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.
- In the **User List** field, type a comma-separated list of the email addresses to which Hitachi Remote Ops should send email about error conditions.
- In the **Sender's Email Address** field, type an email address to be used 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 known to the server.

4. Click **Submit**.
5. Optional: To send a test email to the specified email addresses, click **Test Email**.

## Step 3: (Conditional) Configure transport agents

An Hitachi Remote Ops transport agent transfers notifications of error conditions to a target location for access by Hitachi Vantara Support. The transfer methods available are HTTPS, FTP, or dial up. For the destinations for each method, contact your authorized HCP service provider.

You can specify multiple transport agents. Hitachi Remote Ops tries these agents 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 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 information about what to specify, see the Hitachi Remote Ops documentation.
5. Click **Submit**.

## Step 4: 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 (device errors, communication errors, devices ok, and not monitored) that categorize the devices known to Hitachi Remote Ops.

To show or hide these tables, toggle the check boxes below the table names at the top of the page and click **Refresh**.

2. Take one of these actions:
  - If the **Summary** page does not display any tables, an **Add a device** link is displayed. Click this link.

- 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 **Hitachi Content Platform (HCP)**.**

A set of configuration fields is displayed.

- 4. In the **Name** field, type a name for the HCP system. The name is typically the system host name. The name can be from one to 40 characters long. Special characters and spaces are allowed.**
- 5. In the **Location** field, type the location of the HCP system. The location can be from one to 40 characters long. Special characters and spaces are allowed.**
- 6. In the **Group** field, type the name of a group associated with the HCP system. The group name can be from one to 40 characters long. Special characters and spaces are allowed. For example, Finance Department.**
- 7. In the **Site ID** field, type your Hitachi Vantara customer ID. If you do not know your customer ID, contact your authorized HCP service provider.**
- 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.**
- 9. In the **Local Interface** field, leave the value as **-any-**.**
- 10. 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.**
- 11. In the **Local Interface** field, leave the value as **-any-**.**
- 12. In the **SNMP Access ID** field, select the SNMP user that corresponds to the SNMP configuration in HCP. Typically, this is **public**.**
- 13. In the **Comms Error Reporting** field, specify whether Hitachi Remote Ops should report communication errors between Hitachi Remote Ops and the HCP system. Values are:**
  - **Yes** — Report communication errors.
  - **No** — Do not report communication errors.

- **Local** — Report communication errors only to the email addresses specified in the base configuration, not through the specified transport agents.
- **Default** — Use the setting in the base configuration.

**12.** Ensure that **Enabled** is selected.

**13.** Ensure that **Trace** is cleared.

**14.** Click **Add**.

When the operation is successful, a message is displayed indicating that the HCP system has been added.



**Important:** Do not click **Add** again. Doing so will add the system a second time.

---





# 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 need to 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.

## Step 1: Log into the KVM host

To log into the KVM host:

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

## Step 2: Configure multipathing

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

To configure multipath:

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

```
lsblk
```

## Step 2: Configure multipathing

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 an mpath.

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

```
multipath -l
```

## Step 3: Create the physical volume on the multipath disk

To create the physical volume on the multipath disk:

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

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

For example:

```
pvcreate /dev/mapper/mpatha
```

### Step 3: Create the physical volume on the multipath disk

2. Enter the following command to create a volume group:

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

For example:

```
vgcreate mpathafc1 /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 kvmfc1 -L 1023G mpathafc1
```

4. Enter the following command to create an `ext4` filesystem on the multipath disk logical volume:

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

For example:

```
mkfs.ext4 /dev/mapper/mpatha1-kvmfc1
```

5. Enter the following command to verify that the physical volume is created:

```
pvdisk
```

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
```

## Step 4: Mount the file system

To mount the new file system:

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

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

For example:

```
mkdir /var/kvm-fc1
```

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/kvmfc1 /var/kvm-fc1/
```

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/mpatha1-kvmfc1 /var/kvm-fc1 ext4 defaults 1 2
```

6. Press the *Escape* key.

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

```
:wq
```

## Step 5: Add the new storage pool to the KVM host

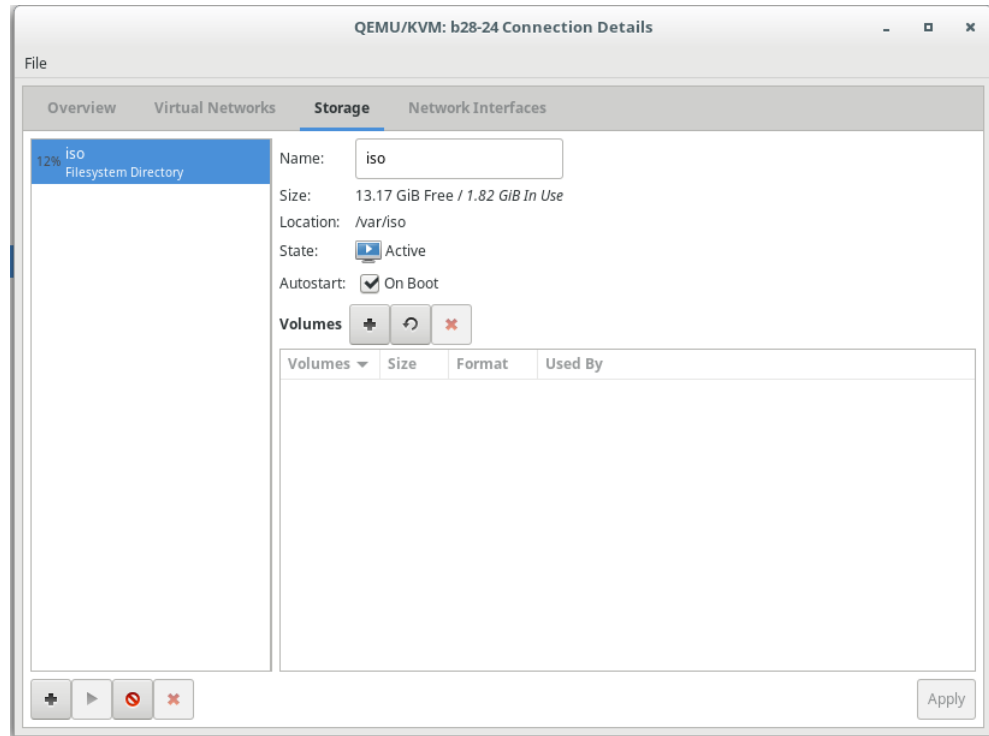
To add the new storage pool to the KVM host:

1. In **Virtual Machine Manager**, right-click the KVM host on which you mounted the volume.

The **Connection Details** window opens.

Step 5: Add the new storage pool to the KVM host

2. Click the **Storage** tab.




3. Click the **Plus** button 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**.

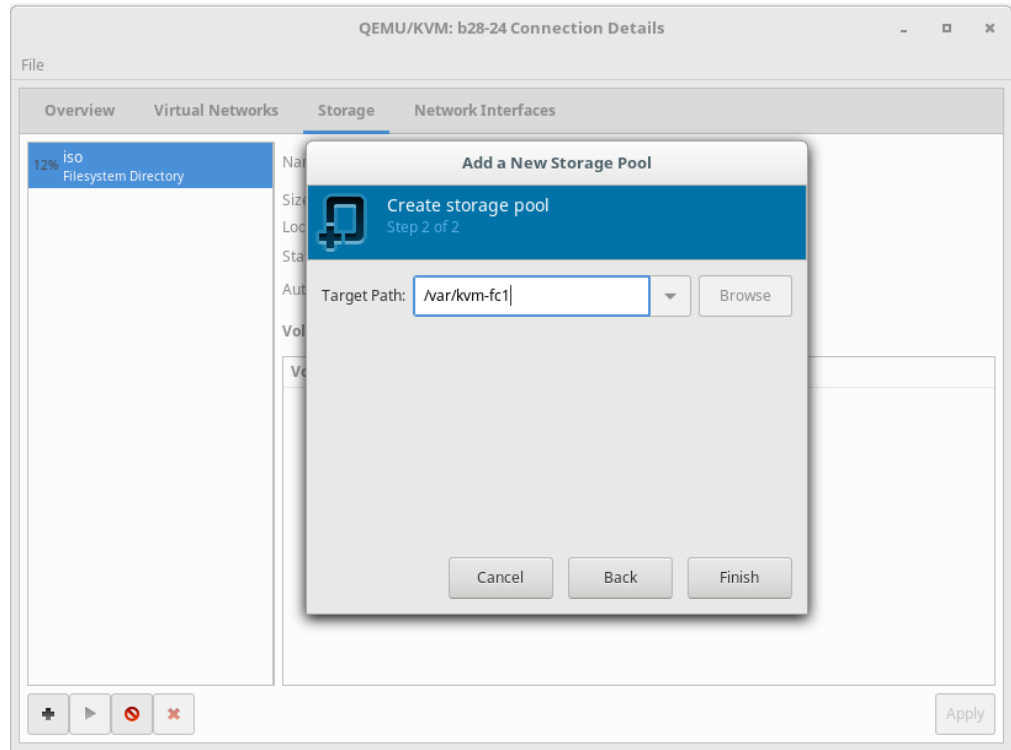


7. Click **Browse**.

8. In the **Target Path** field, enter the file path for the mount point location.

Step 5: Add the new storage pool to the KVM host

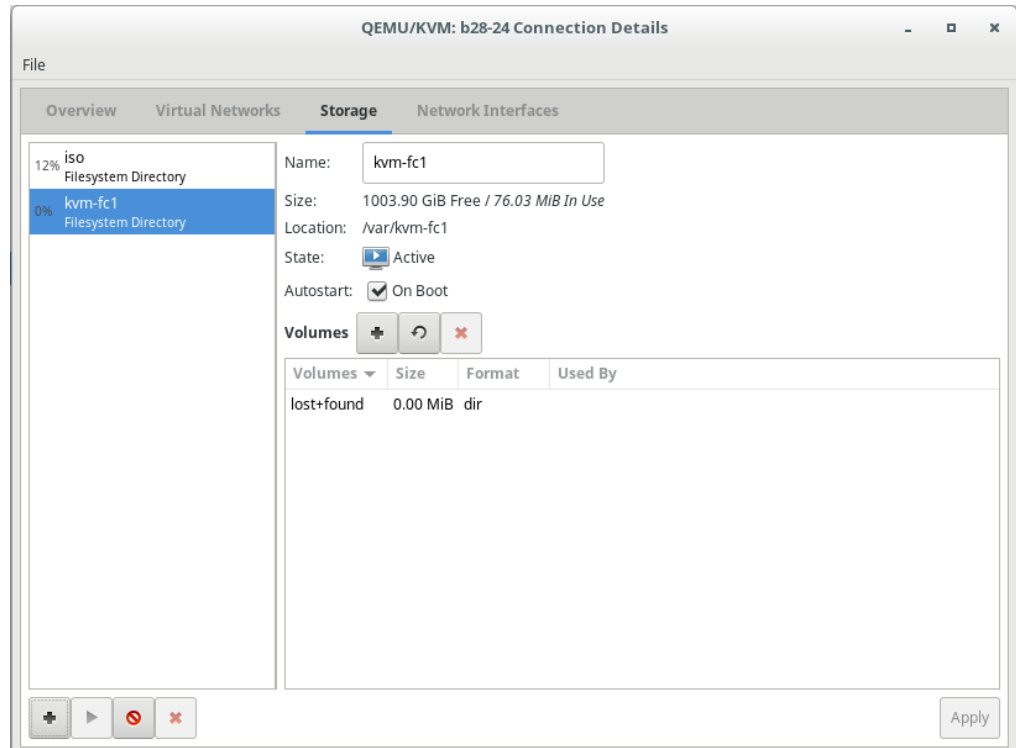
**9. Click **Finish**.**





## Step 5: Add the new storage pool to the KVM host

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





# 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 need to repeat these instructions to create one HCP VM node on each KVM host.

Before you deploy an HCP VM using the command line:

- Install the KVM packages (See [Chapter 3: "Installing KVM"](#) on page 11)
- Create network bridges (See [Chapter 4: "Configuring KVM networking"](#) on page 17)
- Copy and unzip a .iso file on your HCP node (See [Step 2: "Copy the .iso file to the KVM host"](#) on page 32)

## Step 1: Log into the KVM host

To log into the KVM host:

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

## Step 2: Perform 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
```



# Glossary

## A

### **access control list (ACL)**

Optional metadata consisting of a set of grants of permissions to perform various operations on an object. Permissions can be granted to individual users or to groups of users.

ACLs are provided by users or applications and are specified as either XML or JSON in an XML request body or as request headers.

### **ACL**

See [access control list \(ACL\)](#).

### **Active Directory (AD)**

A Microsoft product that, among other features, provides user authentication services.

### **AD**

See [Active Directory \(AD\)](#).

### **alert**

A graphic that indicates the status of some particular element of an HCP system in the System or Tenant Management Console.

## C

### **capacity**

The total amount of primary storage space in HCP, excluding the space required for system overhead for all data to be stored in primary running storage, including the fixed-content data, metadata, any redundant

data required to satisfy services plans, and the metadata query engine index.

## **CIFS**

Common Internet File System. One of the namespace access protocols supported by HCP. CIFS lets Windows clients access files on a remote computer as if the files were part of the local file system.

## **custom metadata**

User-supplied information about an HCP object. Custom metadata is specified as one or more annotations, where each annotation is a discrete unit of information about the object. Users and applications can use custom metadata to understand repurpose object content.

## **D**

### **database**

An internal component of an HCP VM system that contains essential data about the system, users, and user's files. The database is maintained by one node and copied to the other.

### **data center**

In VMware vSphere, a logical unit for grouping and managing hosts.

### **data protection level (DPL)**

The number of copies of the data for an object HCP must maintain in the repository. The DPL for an object is determined by the service plan that applies to the namespace containing the object.

### **datastore**

A representation of a location in which a virtual machine stores files. A datastore can represent a location on a host or an external storage location such as a SAN LUN.

### **domain**

A group of computers and devices on a network that are administered as a unit.

### **domain name system**

A network service that resolves domain names into IP addresses for client access.

## **DNS**

See [domain name system](#).

## **DPL**

See [data protection level \(DPL\)](#).

## **E**

### **ESXi**

See ["VMware ESXi"](#).

## **H**

### **Hitachi Content Platform (HCP)**

A distributed object-based storage system designed to support large, growing repositories of fixed-content data. HCP provides a single scalable environment that can be used for archiving, business continuity, content depots, disaster recovery, e-discovery, and other services. With its support for multitenancy, HCP securely segregates data among various constituents in a shared infrastructure. Clients can use a variety of industry-standard protocols and various HCP-specific interfaces to access and manipulate objects in an HCP repository.

### **HCP VM system**

An HCP VM in which the nodes are virtual machines running in a KVM or VMware vSphere environment.

## **HDDS**

See ["Hitachi Data Discovery Suite \(HDDS\)"](#)

### **Hitachi Data Discovery Suite (HDDS)**

A Hitachi product that enables federated searches across multiple HCP systems and other supported systems.

## **host**

A physical computer on which virtual machines are installed and run.

## K

### KVM

KVM hypervisor is the virtualization layer in Kernel-based Virtual Machine (KVM) which runs on Linux nodes.

## L

### logical unit number (LUN)

A number used to identify a logical unit, which is a device addressed by the Fibre Channel.

### logical volume

A logical unit of storage that maps to the physical storage managed by a node. The physical storage can be storage that's managed by HCP or storage on an external NFS device.

### LUN

See ["logical unit number \(LUN\)"](#).

## M

### metadata

System-generated and user-supplied information about an object. Metadata is stored as an integral part of the object it describes, thereby making the object self-describing.

### multipathing

In SAIN systems, multiple means of access to a logical volume from a single node.

## N

### namespace

A logical partition of the objects stored in an HCP system. A namespace consists of a grouping of objects such that the objects in one namespace are not visible in any other namespace. Namespaces are configured independently of each other and, therefore, can have different properties.



**network**

In an HCP system that supports virtual networking, a named network configuration that identifies a unique subnet and specifies IP addresses for none, some, or all of the nodes in the system.

**network file system**

One of the namespace access protocols supported by HCP. NFS lets clients access files on a remote computer as if the files were part of the local file system.

**network interface controller (NIC)**

A hardware interface that connects the computer to its appropriate network. NICs can be physical (pNIC) or virtual (vNIC).

**NFS**

See [network file system](#).

**NIC**

See "[network interface controller \(NIC\)](#)".

**node**

A server or virtual machine running HCP VM software. Two nodes are networked together to form an HCP VM system.

**O****object**

An exact digital representation of data as it existed before it was ingested into HCP, together with the system and custom metadata that describes that data. Objects can also include ACLs that give users and groups permission to perform certain operations on the object.

An object is handled as a single unit by all transactions, services, and internal processes, including shredding, indexing, versioning, and replication.

**open virtualization format (OVF)**

Standard file style for packaging and distributing virtual software.

**OVF**

See "[open virtualization format \(OVF\)](#)".

**ping**

A utility that tests whether an IP address is accessible on the network by requesting a response from it. Also, to use the ping utility.

**pNIC**

See ["network interface controller \(NIC\)"](#).

**Q****query**

A request submitted to HCP to return metadata for objects or operation records that satisfy a specified set of criteria. Also, to submit such a request.

**R****RAIN**

See [redundant array of independant nodes \(RAIN\)](#).

**redundant array of independant nodes (RAIN)**

An HCP system configuration in which the nodes use internal or direct-attached storage.

**replication**

A process by which selected tenants and namespaces are maintained on two or more HCP systems and the objects in those namespaces are managed across those systems. Typically, the systems involved are in separate geographic locations and are connected by a high-speed wide area network. This arrangement provides geographically distributed data protection (called **geo-protection**).

**repository**

The aggregate of the namespaces defined for an HCP system.

**running storage**

Storage on continuously spinning disks.

# S

## **SAIN**

See ["SAN-attached array of independent nodes \(SAIN\)"](#).

## **SAN-attached array of independent nodes (SAIN)**

An HCP system configuration in which the nodes use SAN-attached storage.

## **search console**

The web application that provides interactive access to HCP search functionality. When the Search console uses the hcp metadata query engine for search functionality, it is called the Metadata Query Engine Console.

## **search facility**

An interface between the HCP Search console and the search functionality provided by the metadata query engine or HDDS. Only one search facility can be selected for use with the Search Console at any given time.

## **secure shell**

A network protocol that lets you log into and execute commands in a remote computer. SSH uses encrypted keys for computer and user authentication.

## **secure sockets layer**

Secure Sockets Layer. A key-based Internet protocol for transmitting documents through an encrypted link.

## **service**

A background process that performs a specific function that contributes to the continuous tuning of the HCP system. In particular, services are responsible for optimizing the use of system resources and maintaining the integrity and availability of the data stored in the HCP repository.

## **service plan**

A named specification of an HCP service behavior that determines how HCP manages objects in a namespace. Service plans enable you to tailor service activity to specific namespace usage patterns or properties.

**simple network management protocol (SNMP)**

A protocol HCP uses to facilitate monitoring and management of the system through an external interface.

**SNMP**

See ["simple network management protocol \(SNMP\)"](#).

**SNMP trap**

A type of event for which each occurrence causes SNMP to send notification to specified IP addresses. SNMP traps are set in management information base (MIB) files.

**spin-down storage**

Storage on disks that can be spun down and spun up as needed.

**SSH**

See ["secure shell"](#).

**SSL**

See [secure sockets layer](#).

**SSL server certificate**

A file containing cryptographic keys and signatures. When used with the HTTP protocol, an SSL server certificate helps verify that the web site holding the certificate is authentic. An SSL server certificate also helps protect data sent to or from that site.

**storage node**

An HCP node that manages the objects that are added to HCP and can be used for object storage. Each storage node runs the complete HCP software (except the HCP search facility software).

**subdomain**

A subset of the computers and devices in a domain.

**switch**

A device used on a computer network to connect devices together.

**syslog**

A protocol used for forwarding log messages in an IP network. HCP uses syslog to facilitate system monitoring through an external interface.

**system management console**

The system-specific web application that lets you monitor and manage HCP.

**T****tag**

An arbitrary text string associated with an HCP tenant or namespace. Tags can be used to group tenants or namespaces and to filter tenants or namespace lists.

**tagged network**

A network that has a VLAN ID.

**tenant**

An administrative entity created for the purpose of owning and managing namespaces. Tenants typically correspond to customers or business units.

**tenant management console**

The tenant-specific web application that lets you monitor and manage tenants and namespaces.

**transaction log**

A record of all create, delete, purge, and disposition operations performed on objects in any namespace over a configurable length of time ending with the current time. Each operation is represented by an operation record.

**U****unix**

Any UNIX-like operating system (such as UNIX itself or Linux).

**upstream DNS server**

A DNS server to which HCP routes the outbound communications it initiates (for example, for sending log messages to syslog servers or for communicating with Active Directory).

**user account**

A set of credentials that gives a user access to one or more of the following interfaces: HCP System Management Console, HCP Tenant Management Console, HCP management API, HCP Search Console, namespace content through the namespace access protocols, HCP metadata query API, and HCP Data Migrator.

**user authentication**

The process of checking that the combination of a specified username and password is valid when a user tries to log into the System Management Console, Tenant Management Console, HCP Search Console, tries to access the HCP system through the management API, or tries to access a namespace.

**V****vCenter**

See ["VMware vCenter Server"](#).

**versioning**

An optional namespace feature that enables the creation and management of multiple versions of an object.

**virtual local area network (VLAN)**

A distinct broadcast domain that includes devices within different segments of a physical network.

**virtual machine**

A piece of software that emulates the functionality of a physical computer.

**VLAN**

See Virtual Local Area Network (VLAN).

**VLAN ID**

An identifier that's attached to each packet routed to HCP over a particular network. This function is performed by the switches in the physical network.

**vmNIC**

A representation in VMware vSphere of one of the physical NICs on a host.

**VMware ESXi**

The underlying operating system for the VMware vSphere product.

**VMware vCenter Server**

A VMware product that allows you to manage multiple ESXi hosts and the virtual machines that they run.

**vNIC**

See "[network interface controller \(NIC\)](#)".

**Z****zero-copy failover**

The process of one node automatically taking over management of storage previously managed by another node that has become unavailable.





# Index

## B

- back-end network
  - node IP addresses, setting during installation 50
- Back-end network 3
- base configuration, Hitachi Remote Ops 97

## C

- Compute 2
- configuring Hitachi Remote Ops 96
- Console pages
  - SNMP 96

## D

- database
  - recovering 88
- diagnostic 70

## E

- email, Hitachi Remote Ops 98
- enabling SNMP 95

## F

- Front-end network 3

## H

- HCP 1
- HCP-VM 1, 9
- HCP-VM nodes 82
- HCP systems
  - enabling SNMP 95
- Hitachi Remote Ops
  - about 95
  - base configuration 97
  - configuring 96
  - email 98

- logging in 97
- transport agents 99

## I

- IP addresses
  - back-end network, setting during installation 50
- IPMI 69

## L

- logging in to Hitachi Remote Ops 97
- logical volumes 73

## N

- Namespaces 1

## P

- pNIC 3

## R

- Repository 1

## S

- SAR 69
- SNMP 69
- SNMP page 96
- SNMP, enabling 95
- System management console 69

## T

- transport agents, Hitachi Remote Ops 99

## V

- Vmware 7
- vNIC 4



**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)