

Hitachi Content Platform

v10.0

Deploying an HCP VM System on ESXi

This book is the setup guide for Hitachi Content Platform Virtual Machine systems. It describes how to deploy a virtualized HCP system in your VMware vSphere® environment.

© 2007, 2025 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., Hitachi Vantara, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

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

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

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

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

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

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

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

AIX, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, IntelliMagic, IntelliMagic Vision, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z17, 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, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

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

The open source content used in Hitachi Vantara products may be found within the Product documentation or you may request a copy of such information (including source code and/or modifications to the extent the license for any open source requires Hitachi make it available) by sending an email to OSS_licensing@hitachivantara.com.

Contents

Preface.....	6
Intended audience.....	6
Product version.....	6
Release notes.....	6
Document conventions.....	6
Related documents.....	9
Accessing product documentation.....	9
Getting help.....	10
Comments.....	10
Chapter 1: HCP system overview.....	11
Introduction to Hitachi Content Platform.....	11
HCP VM system architecture and components.....	11
HCP VM architecture.....	11
HCP VM system components.....	12
Host platform.....	12
Compute.....	12
Storage.....	12
HCP network connectivity.....	14
Front-end network.....	14
Back-end network.....	14
Management port network.....	15
Storage network.....	15
Dedicated database volume.....	15
Hardware monitoring and alerting.....	15
HCP software.....	15
HCP upgrades.....	16
HCP VM node failover (vCenter and vSphere HA).....	16
Storage licensing.....	17
Chapter 2: Configuration guidelines for the HCP VM environment.....	18
Supported VMware versions.....	18
Supported VMware functionality.....	18
HCP VM hardware requirements.....	19
HCP VM system limits.....	21

HCP VM availability considerations.....	22
Chapter 3: Configuring the HCP VM environment.....	23
ESXi considerations.....	23
Enabling NTP for the ESXi hosts.....	23
Configuring a vSphere HA cluster for HCP VM.....	24
Provisioning HCP VM storage.....	25
Adding VMFS datastores to a vSphere HA cluster.....	28
About NFS datastores.....	29
Creating an NFS datastore.....	29
Configuring networking.....	30
Configuring networking for front-end switching.....	30
Configuring networking for back-end switching.....	30
(Optional) Configuring networking for management port switching.....	31
Chapter 4: Creating the HCP VM system.....	32
Unpacking and uploading the ISO zip file.....	32
Creating and configuring the new virtual machine.....	32
Creating the virtual machine.....	32
Configure the network adapters.....	34
Delete the ISO image connection.....	35
Installing the Appliance Operating System.....	35
Installing the HCP software.....	37
Step 1: Identify the nodes in the HCP system.....	40
Step 2: Configure the HCP system.....	41
Step 3: Run the HCP installation.....	46
Step 4: Verify the HCP software installation.....	50
Setting additional configuration options.....	51
Monitoring and alerting.....	52
Software monitoring.....	52
HCP VM resource monitoring.....	53
HCP VM diagnostic menu.....	53
Chapter 5: Maintaining the HCP VM environment.....	55
Adding logical volumes.....	55
Moving storage node databases to optimal volumes.....	60
Deleting databases from older database volumes.....	63
Adding HCP VM nodes.....	64
Recovering storage nodes and preserving volumes.....	68
Chapter 6: Configuring HCP monitoring with Hitachi Remote Ops.....	72
Enabling SNMP in HCP.....	72
Configuring Hitachi Remote Ops.....	73

Step 1: Log in to Hitachi Remote Ops.....	73
Step 2: Set the base configuration.....	74
Step 3: (Conditional) Configure transport agents.....	74
Step 4: Identify the HCP system.....	75
Appendix A: Changing the HCP VM network adapters.....	77
About network adapters.....	77
Disabling LRO on the ESXi host for VMXNET3.....	77
Changing the HCP VM network adapter.....	78
Step 1: Power off the HCP VM.....	78
Step 2: Remove the previous network adapters.....	78
Step 3: Configure the front-end network adapters.....	78
Step 4: Configure the back-end network adapters.....	79
Step 5: Power on the HCP VM.....	79
Appendix B: Modifying the DRS settings.....	80
Appendix C: Configuring the HCP VM small instance.....	82
Appendix D: Managing failover.....	83

Preface

This book is the setup guide for Hitachi Content Platform (HCP) Virtual Machine (VM) systems. It provides the information you need to deploy a virtualized HCP system in your VMware vSphere® environment. To complete the installation, there are instances where you may want to reference other materials.

Intended audience

This guide is intended for personnel who are responsible for deploying an HCP VM system on a VMware vSphere® environment at a customer site. It assumes that the audience has experience with computer networking and creating virtual machines, familiarity with VMware products and concepts, and a basic understanding of HCP systems.

Product version

This guide applies to Hitachi Content Platform release 10.0 or later.

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 the Hitachi Vantara documentation website: <https://docs.hitachivantara.com>.




Document conventions




This document uses the following typographic conventions:

Convention	Description
Bold	<ul style="list-style-type: none">Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK.Indicates emphasized words in list items.

Convention	Description
<i>Italic</i>	<ul style="list-style-type: none"> Indicates a document title or emphasized words in text. Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: <pre>pairdisplay -g group</pre> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>
< > angle brackets	Indicates variables in the following scenarios: <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> Variables in headings.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Important	Highlights information that is essential to the completion of a task.

Icon	Label	Description
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	CAUTION	Warns the user of a hazardous situation that, if not avoided, could result in major or minor injury.
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Related documents

The following documents contain additional information about Hitachi Content Platform:

- *HCP System Management Help* — This help system is a comprehensive guide to administering and using an HCP system. The Help 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, and managing search and indexing. 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 delete 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.
- *Installing HCP Evaluation Edition* — This book contains instructions for setting up a virtual HCP system for product evaluation purposes.
- *Installing an HCP with Internal Storage System - Onsite Setup Guide* — This book contains instructions for deploying an assembled and configured HCP with internal storage (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 with internal storage system that was ordered without a rack and for configuring Hitachi Remote Ops to monitor the nodes in an HCP system.
- *Installing an HCP SAN-attached Storage System - Onsite Setup Guide* — This book contains instructions for deploying an assembled and configured single-rack HCP with SAN-attached storage (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.

Accessing product documentation

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

Getting help

The [Hitachi Vantara Support Website](#) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to the Hitachi Vantara Support Website. For additional contact methods, go to <https://support.hitachivantara.com/en/contact-support.html>.

[Hitachi Vantara Community](#) 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, register, and complete your profile.

Comments

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

Thank you!

Chapter 1: HCP system overview

This chapter introduces Hitachi Content Platform. It describes the architecture of an HCP VM system installed in a VMware vSphere 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 gives access to objects through a variety of industry-standard protocols and through various HCP interfaces.

HCP VM system architecture and components

This section describes the architecture and components of an HCP Virtual Machine (HCP VM) system.

HCP VM architecture

The following figure describes the architecture of an HCP system running on VMware infrastructure.



HCP VM system components

Host platform

In an HCP VM system, each HCP VM node runs in a virtual machine on an VMware ESXi host.

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

HCP VM storage infrastructure is highly available and fault tolerant. It is best practice for the physical servers that the ESXi hosts run on to be connected to shared SAN storage with RAID 6 protection or Hitachi NAS (HNAS).

The HCP SAN storage must have at least two paths to each logical unit number (LUN), and each LUN must have the same LUN number (HLUN) on each ESXi host.

A datastore will be created from each LUN or export, creating one Virtual Machine File System (VMFS) volume per LUN or export. A single datastore is not shared by HCP VM nodes. However, HCP VM nodes can have multiple datastores. Each datastore is carved into one or multiple Virtual Machine Disks (VMDK) which are presented to the HCP OS as local disks. The HCP OS recognizes its storage as internal drives similar to an HCP appliance with a local storage configuration. The disks are controlled by the VMware Paravirtual SCSI controller (PVSCSI). VMware recommends PVSCSI for better overall performance.



Tip: The PVSCSI adapter reduces CPU utilization and potentially increases throughput compared to default virtual storage adapters.

Each VMDK can be a maximum size of 15.90TB.

In addition to the RAID 6, shared SAN storage configuration, and HNAS datastores, HCP VM also supports the following for storage configurations:

- Shared SAN arrays with virtual volumes created from Dynamic Provisioning (DP) pools. This configuration does not support thin provisioning.
- Shared SAN arrays with LUNs configured using Raw Device Mapping (RDM) in vSphere® Client or vCenter™. The RDM is to be configured in Physical Mode.
- Best Practices: HCP VM Placement Strategy. To ensure optimal availability and resilience for your HCP virtual machine (VM) nodes, we recommend implementing the following host-affinity rules. Implementing these rules will significantly enhance the fault tolerance and stability of your HCP deployment.:
 - Prevent co-location: Configure host affinity rules to ensure that individual HCP VM nodes are distributed across different physical ESXi hosts. This prevents a single host failure from impacting multiple HCP instances.
 - Balanced Placement for Limited Hosts: In scenarios where the number of physical ESXi hosts is less than the total number of HCP VM nodes, these affinity rules should be designed to provide a balanced distribution. The goal is to maximize resource sharing while still maintaining the critical “N/2 + 1” quorum target for HCP instances, ensuring continued operation even with host failures.
- VMware vSAN datastores.
 - VMware vSAN aggregates direct attached storage (DAS) devices to create a single storage pool that is shared across all VMware vSphere ESXi hosts in the vSAN cluster.
 - Choose vSAN datastore for storing VMDK files for HCP node VMs.
 - Use thick eager-zeroed VMDK to avoid overprovisioning storage that is used by HCP VMs.
 - Follow VMware best practices to configure vSAN storage-backed ESXi environments.
 - For optimal performance, use VM storage policy to ensure that all VM files for an HCP VM node reside on storage of the same ESXi server that is running the HCP VM node.

- Other devices like HNAS that export NFS v3 shares, which are mounted and used as NFS datastore.
 - Thick, eager zero is needed when formatting NFS datastores, so additional ESXi plug-ins might be required from your vendor. Hitachi provides a VAAI plug-in that enables this functionality on the HNAS platform.
 - Best practice is to not have multiple datastores on the same file system or the same underlying disk due to performance and availability considerations.

Follow the vendor's best practice for configuring NFS datastores.
- RAID-protected storage that's internal to the ESXi hosts. Each LUN created from this storage corresponds to a Virtual Machine File System datastore. This configuration does not support vSphere High Availability (HA). The underlying storage in this configuration must be RAID protected.
- HCP VM also supports an ESXi datastore that resides on SSD storage.

If you deviate from the best-practice configuration, you need to consider the possible ramifications to performance, availability, backup, security, ease of management, and data integrity. Ensure that you completely understand failure scenarios, HDD failure rates, RAID protection levels, RAID rebuild times, and support windows before changing configurations. System health must be closely monitored to prevent service failures and ensure that the underlying storage does not fail.

For information about supported storage vendors and devices, see the applicable VMware documentation.

HCP network connectivity

HCP VM network connectivity is provided to the HCP guest operating system by VMware VMXNET3 vNICs, vSwitches, and distributed switches.

For the VMXNET3 vNIC, the back-end vSwitch must be configured to provide access to one vmNIC and the front-end vSwitch must be configured to provide access to a different vmNIC.

Front-end network

The HCP front-end network is used for client and management access. For HCP front-end networks, it is best practice for the ESXi host to create one vNIC mapping to two physical NICs (pNICs) on the ESXi host. Having two pNICs ensures redundancy and consistent performance.

Back-end network

The HCP private back-end network is used for internode communication and data transfer. The ESXi host has one vmNIC that maps to two physical NICs on the ESXi host 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 ESXi hosts must connect to the same physical switch (switch1), and pNIC-2 on all ESXi 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.

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



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

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

Management port network

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

Storage network

Best practice is for the VM kernel network be set up in a private network or with a unique VLAN ID that gives network isolation.

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 [Configuring HCP monitoring with Hitachi Remote Ops \(on page 72\)](#).

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 with internal storage and HCP with SAN-attached storage 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 with internal storage and HCP with SAN-attached storage 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.

HCP upgrades

HCP v5.0 introduced HCP Evaluation Edition for proof of concept (POC) and test activities at Hitachi Vantara partner and customer sites. Upgrades from the Evaluation Edition single node and Evaluation Edition multi-node to HCP VM are not supported. HCP VM supports upgrades from the initial 6.0 release to future releases of HCP VM.

HCP VM node failover (vCenter and vSphere HA)

To configure automatic failover in the event of an ESXi host failure, an instance of the VMware vCenter Server must be available. Failover functionality is provided by a vSphere High Availability (HA) cluster.

A vSphere HA cluster allows a collection of ESXi hosts to work together to optimize their levels of availability. The system administrator is responsible for configuring the cluster to respond to host and virtual machine failures.

The system administrator will configure each ESXi host participating in an HCP VM system to be part of a single vSphere HA cluster in vCenter. This enables high availability if one or more servers or ESXi hosts fail. When the primary host detects a server or ESXi host failure, it can restart the HCP VM node that was running on the failed server or ESXi host to other healthy ESXi hosts in the cluster.

The primary host monitors the status of secondary hosts in the cluster through network heartbeat exchanges every second. If the primary host stops receiving heartbeats from a secondary, it checks for liveness before declaring a failure. The liveness check determines if the secondary is exchanging heartbeats with a datastore.

The HCP VM system administrator will not configure the HCP VM vSphere HA cluster to automatically move the failed-over HCP VM node back to its original ESXi host after the host is available. The system administrator will manually shut down the HCP VM node, and the vCenter administrator will manually move the HCP VM node onto the preferred ESXi host and start the HCP VM node. After the node starts, it will rejoin the HCP VM system.

In the case of network isolation, the HCP VM vSphere HA cluster will be configured to leave the HCP VM node powered on. This means that the HCP VM node will still be able to communicate over its private back-end network with the other HCP VM nodes in the system. As in the case of a physical HCP node, the HCP VM node and the data it is managing will remain available to the system through the front-end of the other nodes in the HCP VM system.

The vCenter Server used to configure the vSphere HA cluster of ESXi hosts for the HCP VM system either can be a pre-existing server in the user environment, or can be allocated as part of the HCP VM HA cluster of ESXi hosts. It is a best practice for the vCenter Server to be separate from the HCP VM HA cluster. The vCenter Server can consume a fair amount of resources on the ESXi host that can be used by the HCP VM nodes.

The rules for creating a vSphere HA cluster for use with HCP VM are very specific. If the HCP VM system will be added to an existing HA cluster, ensure that the cluster is configured exactly to the specifications in this guide.

Storage licensing

HCP VM systems come with a basic storage license that provides 2 terabytes of active and HCP S Series storage. The basic storage license also provides 2 terabytes of extended storage. If you need additional storage, contact your Hitachi Vantara sales representative.

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

Chapter 2: Configuration guidelines for the HCP VM environment

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

Supported VMware versions

HCP VM supports multiple versions of VMware.

For a list of supported VMware versions, see the release notes on Hitachi Vantara Support Connect:

https://docs.hitachivantara.com/search/all?filters=Product_custom~%2522Content+Platform%2522*PublicationType~%2522Release+Notes%2522*Version_custom~%25229.6.x%2522&content-lang=en-US

Supported VMware functionality

HCP VM supports the following VMware functionality:

- vSphere HA cluster.
- HCP VM node shutdown from the vCenter management console. This functionality is part of the VMware tools package, which is included in the HCP OS with HCP VM.



Note: Pausing live migration and other functionality enabled by the VMware tools package are not currently supported.

- DRS when used in a manual capacity to assist with virtual machine-to-host affinity. See [Modifying the DRS settings \(on page 80\)](#).
- The current version of HCP VM does not support the following VMware functionality:
 - Other failover capabilities provided by VMware, such as vMotion, storage vMotion, and DRS with FT
 - Software used for VM replication

The following HCP features are specific to the physical HCP appliances (HCP with internal storage system, HCP with SAN-attached storage system) and are not applicable to HCP VM:

- **Autonomic tech refresh:** Gives you the capability of migrating a VM to a different host, allowing for server refresh. The raw storage layer is obscured from HCP in the VMware environment; any storage update must be handled at the VMware layer.
- **Zero-copy failover:** VMware HA replaces this capability by restarting an HCP VM on an ESXi host after it is lost due to a host failure. This storage availability is provided by shared SAN storage.
- **Spindown, Indexing-only LUNs:** Spindown is not compatible with the VMware environment. Indexing-only LUNs are not available in HCP VM with this release. Shared-index LUNs are standard as with all other HCP systems.
- **HCP integrated HDvM monitoring:** The raw storage layer is obscured from HCP in the VMware environment. Storage connected to HCP VM needs to be monitored at the customer site using their preferred mechanism.
- **VLAN tagging:** Active-active NIC Teaming in VMware is designed for load balancing and redundancy. Both physical NICs must be configured with the same VLAN tagging. Also, VMware vSwitch is a layer 3 switch and will not route traffic out of physical NICs per VLAN tagging. You cannot configure physical vmNIC2 to be tagged on VLAN 20 and physical vmNIC3 to be tagged on VLAN 30 so that VMware will route HCP traffic out of the appropriate physical NIC.

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
- A minimum of four 1.2 TB LUNs for the default VMDK size deployment



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



Note: An HCP VM solution requires a minimum of four nodes. Running two or more HCP VM nodes of a single HCP system on a single ESXi host is supported as long as the four HCP VM nodes meet the minimum requirements described in this section.



WARNING: If two or more HCP VM nodes are hosted on a single ESXi host, depending on the ESXi and HCP environment, losing a single ESXi host may result in:

- Temporary or permanent loss of object data, loss of system metadata, loss of custom metadata (annotations), or all of the above, if the VMware ESXi environment architecture stores HCP node virtual machine files on ESXi hosts.
- Loss of access to all data until the failed VMware ESXi host is recovered.
- The HCP system possibly becoming read-only, preventing any updates or object uploads if the HCP environment falls below quorum.
- Temporary or permanent loss of the entire HCP system, including all object data, system metadata, custom metadata (annotations), and object retention information stored in the HCP system.

For these reasons, thorough system architecture design and risk assessment are required before multiple HCP VM nodes are placed on a single VMware ESXi host, with consideration for data availability, system resiliency, and single points of failure in the HCP system. At a minimum, increased Data Protection Level (DPL) within the HCP system and HCP replication to another HCP system are recommended measures to mitigate some of the risks inherent in the architecture if multiple HCP VM nodes are run on a single VMware ESXi host.

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



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

- Suggested volume size for NFS datastores: Contact your storage array vendor or NFS server vendor for information about the maximum NFS volume size.
- Two-port Fibre Channel HBA cards (or VMware compatible IO device) for shared SAN storage connectivity (when applicable).

- ESXi needs a minimum of 2 GB of physical RAM. VMware suggests providing at least 8 GB of RAM to take full advantage of ESXi features and run virtual machines in typical production environments.
- To ensure continuous availability and fault tolerance for your HCP cluster, an ESXi server can host one HCP VM. In addition, the same ESXi server can host one or more workloads unrelated to HCP. However, you must ensure that the HCP VM has access to the necessary CPU, memory, network, and disk resources to operate. For more information, see [HCP VM availability considerations \(on page 22\)](#).

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

An HCP VM has system limits.

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

- 40 HCP VM nodes
- 59 data LUNs on each HCP VM node: (ESXi guest OS limitation)
- Maximum open VMDK storage per host (ESXi limitation). See the VMware documentation for the ESXi limitations associated with the version that you are running.

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

HCP VM availability considerations

An HCP cluster is considered in a state of continuous availability if there is one HCP VM node per ESXi host, and if more than half of the HCP VM nodes are healthy and running. When the HCP cluster is in this state, the system can survive a single ESXi 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 ESXi host and one of your ESXi 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 ESXi hosts can cause HCP system instability.

A vSphere HA cluster is suggested, but not required.

Chapter 3: Configuring the HCP VM environment

This chapter describes how to provision the VMware environment for an HCP VM deployment.

ESXi considerations

You might want to deploy the HCP VM system on existing VMware ESXi hosts in your environment. Before doing this, make sure that the hosts meet the minimum requirements for compute and memory catalogued in [HCP VM hardware requirements \(on page 19\)](#).

Depending on the versions of VMware ESXi, VMware vCenter Server, and VMware vSphere that you are running, steps and terminology might be different in the GUI. This document applies to VMware vSphere versions 6.5 and later; however, you can also use it as a reference for earlier versions. For information about ESXi versions 6.0 and earlier, see the revision of this document with HCP version 8.0.

For a list of supported VMware versions for a specific HCP version, see the HCP Release Notes associated with that version.

For information about VMware functionality that HCP VM supports, see [Supported VMware functionality \(on page 18\)](#).

When provisioning storage to use with HCP VM, make sure to review and follow the VMware ESXi storage guides and the relevant storage vendor's VMware best practices guide. For more information about provisioning storage to use with HCP VM, see [Provisioning HCP VM storage \(on page 25\)](#).



Important: All ESXi hosts that will contain an HCP VM node must have Network Time Protocol (NTP) enabled. To enable NTP, use the Time Configuration option in the vSphere client on each ESXi host.

Enabling NTP for the ESXi hosts

Procedure

1. Access your vSphere client.
2. Select the ESXi host for which you want to enable NTP.
3. Click **Configure**.
4. Under **System**, click **Time Configuration**.
5. Click **Edit**.

6. Specify whether you want to manually configure the date and time for the host or whether you want to enable the NTP client. Enabling the NTP client enables date and time synchronization with an NTP server.
7. For the **NTP service startup policy**, select **Start and stop with host**.
8. In the **NTP Servers** section, enter the time server or servers you want to use.
9. Click **OK** and then start or restart the NTP service.
10. Repeat the procedure with the same time server for all ESXi hosts that reside on an HCP VM node.



Tip: Write down the NTP server used in your ESXi hosts. You will use this information during the HCP VM installation.

Configuring a vSphere HA cluster for HCP VM

A VMware vSphere HA cluster lets a collection of ESXi hosts work together to optimize their levels of availability. You are responsible for configuring the cluster to respond to host and virtual machine failures.

Procedure

1. Create a datacenter:
 - a. Access the vSphere client.
 - b. In the left navigation bar, right-click your server and select **New Datacenter**.
 - c. Enter a name for your HCP VM datacenter: for example, `hcp-vm-center-1`.
 - d. Click **OK**.
2. Add a cluster:
 - a. In the left navigation bar, right-click the datacenter you just created and select **New Cluster**.
 - b. Enter a name for the cluster: for example, `hcp-vm-cluster-1`.
 - c. Select the check box to turn on vSphere HA.



Important: Do NOT select the option to turn on DRS. You can enable DRS later to define VM affinity to a particular host or group of hosts. This function does not provide further failover automation. The setting assists with keeping virtual machines on a given host, and alert you if the rule cannot be followed. For information about the required DRS settings, see [Modifying the DRS settings \(on page 80\)](#).

- d. Click **OK**.
3. Add ESXi hosts:

For information about the maximum number of ESXi hosts allowed in a vSphere HA cluster, see the VMware documentation.

 - a. In the left navigation bar, right-click the cluster you created and select **Add Host**.
 - b. In the **Add Host** wizard, enter the ESXi host connection information.
 - c. Enter username and password for the ESXi host.

- d. Review the host summary.
- e. Enter the license information for the ESXi host if it is not already assigned.
- f. (Optional) To increase security on your ESXi host and prevent remote users from logging in directly, select **Lockdown mode**.



Note: You are responsible for the decision whether to implement lockdown mode.

- g. Review your choices, then click **Finish**.
The ESXi host is added to the vSphere HA cluster.
- h. Repeat this procedure for each ESXi host in the system.

Provisioning HCP VM storage

When provisioning storage to use with HCP VM, make sure to review and follow the VMware ESXi storage guides and the relevant VMware best practices guide from the storage vendor.


You can provision HCP VMs in a local storage configuration or a shared SAN storage configuration. Local storage is not the best option due to its increased data availability risk. For this reason, if you provision HCP VM in a local storage configuration, the best practice is to set your data protection level (DPL) to 2. For more information about data protection levels, see *HCP System Management Help*.

The following guidelines are for provisioning shared SAN storage for use with HCP VM:

- Datastores used for HCP VM nodes must be backed by shared RAID 6 storage.
- Datastores can reside on HDD or SSD storage.
- Each datastore should only consist of one LUN.
- HCP VM nodes cannot share datastores.
- All LUNs will be mapped to ESXi hosts in the vSphere HA cluster.
- All LUN IDs must be consistent across hosts. For example, LUN 1 should be mapped to host 1, host 2, host 3 and host 4 as LUN 1.
 - This is also true for VMDK and RDM.
 - For Network File System (NFS), all ESXi hosts must mount the export with the same datastore name.
- All SAN LUNs will have at least two paths (*multipathing*) presented to the ESXi host.
- If fabric is connected, redundant FC switches will be deployed as part of the HCP VM storage environment to ensure maximum availability.
- To ensure maximum data security, use WWN zoning (not port) for HCP VM Zones.
- If loop is connected, redundant controllers must be provisioned for the HCP VM storage environment to ensure maximum availability. Do not use different ports on the same array controller.

The following figure shows a sample SAN layout for VMDK and RDM. The number of storage controller ports dedicated to an HCP VM system depends on the capabilities of the storage array. For Hitachi Vantara midrange storage, the best practice is to spread host access across all cores.

Consult the storage vendor documentation for sizing and configuration options.



FC Switch 1, HCP VM path 1

Zone name	Zone member wwpn	Zone member wwpn
HCP_VM_cluster_1_path_1	Storage controller 0	ESXi_host1_port0
	Storage controller 0	ESXi_host2_port0
	Storage controller 0	ESXi_host3_port0
	Storage controller 0	ESXi_host4_port0

FC Switch 2, HCP VM path 2

Zone name	Zone member wwpn	Zone member wwpn
HCP_VM_cluster_1_path_2	Storage controller 1	ESXi_host1_port1
	Storage controller 1	ESXi_host2_port1
	Storage controller 1	ESXi_host3_port1
	Storage controller 1	ESXi_host4_port1

The following tables are sample Host Group/LUN layouts that display the same LUNs mapped with the same HLUN to each ESXi host.

These examples assume that the ESXi OS LUN has already been provisioned, but it can be provisioned from the SAN as well. In the case of the OS LUN being provisioned on the SAN, only the ESXi host that is starting from the LUN should be granted access.

Table 1 Array path 1

Host Group Name	Hosts	HLUN	ArrayLUN	VMware dataStore
HCP_VM_cluster_1_path_1	ESXi-1	1	10	hcp-vm_cluster-1_node_1_datastore_1
	ESXi-2			
	ESXi-3	2	11	hcp-vm_cluster-1_node_2_datastore_1
	ESXi-4			
	4	12	hcp-vm_cluster-1_node_3_datastore_1	
	5	13	hcp-vm_cluster-1_node_4_datastore_1	

Array path 2

Host Group Name	Hosts	HLUN	ArrayLUN	VMware dataStore
HCP_VM_cluster_1_path_2	ESXi-1	1	10	hcp-vm_cluster-1_node_1_datastore_1
	ESXi-2			
	ESXi-3	2	11	hcp-vm_cluster-1_node_2_datastore_1
	ESXi-4			
	4	12	hcp-vm_cluster-1_node_3_datastore_1	
	5	13	hcp-vm_cluster-1_node_4_datastore_1	

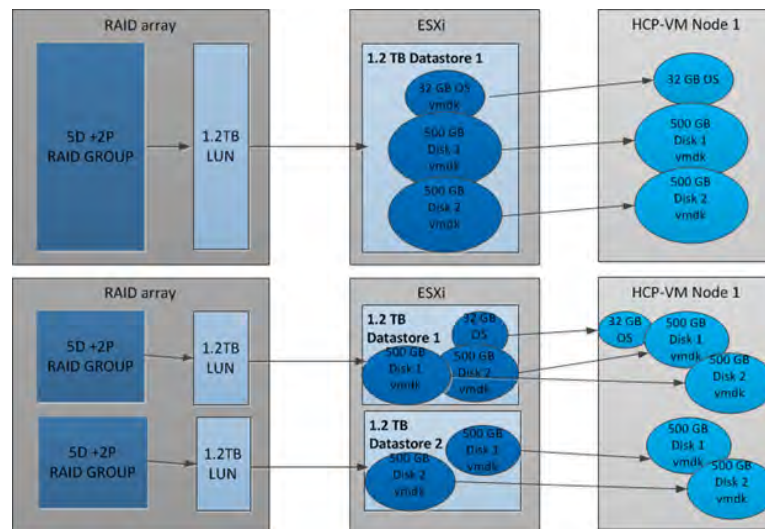
Adding VMFS datastores to a vSphere HA cluster

As a best practice, use only one LUN from a RAID Group in the HCP VM system. Adding multiple LUNs from the same RAID Group increases the risk of data loss in the event of a failure.

A datastore can only be set for one HCP VM node, but each HCP VM node can have multiple datastores.

In an HCP VM system using VMware version 5.5 or later, the largest a disk can be is 16 TB.

The following is a visual depiction of the cluster layout:



Procedure

1. Access your vSphere Client.
2. In the **Datastores** section, click the option to create a new datastore.
3. To specify the datastore type, select **VMFS**.
4. Enter a meaningful name for the datastore: for example, `hcp-vm_cluster_1_node_1_datastore_1`.
5. Select the VMFS version for the datastore (VMFS 6 is advised).
6. Specify partition configuration details for the datastore.
7. Review your choices, then click **Finish** to create the VMFS datastore.
The datastore should now be initialized and mounted. If it is, then in the **Recent Tasks** section of the vSphere Client, a **Rescan VMFS** alarm should be issued for all other ESXi hosts in the cluster.

The new datastore should be automatically added to the inventory of all the other ESXi hosts.
8. Repeat this procedure for any other datastore LUNs with all the same values and verification except for the datastore name.

Here are examples of other identifiable datastore names you can use:

- LUN2 = hcp-vm_cluster_1_node_2_datastore_1
- LUN3 = hcp-vm-cluster_1_node_3_datastore_1
- LUN4 = hcp-vm-cluster_1_node_4_datastore_1

About NFS datastores

You can configure HNAS file systems and their underlying storage in several different ways. To achieve the best performance, follow these guidelines for configuring HNAS in a VMware vSphere environment:

- In general, a 4 KB file system block size is best practice; 32 KB can be used in instances where all VMs on a specific HNAS file system perform large block requests.
- Set `cache-bias` to `large` (`cache-bias --large-files`).
- Disable shortname generation and access time maintenance (`shortname -g off`, `fs-accessed-time --file-system <file_system> off`).
- Disable the quick start option for HNAS read ahead when VM IO profiles are primarily random (`read-ahead --quick-start disable`).
- NFS exports: Do not export the root of the file system.
- File system utilization: Maintain at least 10% free space in each file system used by ESXi hosts.
- Storage pools: Do not mix disk types in the same storage pool.
- Limit ownership of all file systems that are created on a storage pool to one EVS.
- Configure a minimum of four (4) System Drives (SD) in a storage pool.
- Configure one (1) LULDEV per RAID group consuming all space (if possible).

Creating an NFS datastore

Procedure

1. Access your vSphere Client.
2. In the **Datastores** section, click the option to create a new datastore.
3. To specify the datastore type, select **NFS**.
4. Enter a meaningful name for the datastore: for example, `hcp-vm_cluster_1_node_1_datastore_1`.
5. Select the NFS version for the datastore.
6. Enter the server name and the mount point folder name.
If you are using NFS version 4.1, you can optionally enable and configure Kerberos authentication to secure NFS messaging.
7. Select the hosts that need access to the datastore.

- Review your choices, then click **Finish** to create the NFS datastore.



Important: Make sure that you mount datastores with the same volume label on all vSphere ESXi hosts within VMware HA environments.

Configuring networking

Networks should be configured for specific switches in the system before the ISO file is deployed. HCP supports both VMware vSphere standard virtual switches (VSS) and vSphere distributed virtual switches (VDS).

Make sure to review *HCP System Management Help* for the latest information about HCP network administration.

To configure standard VMware vSwitch networking for front-end switching, back-end switching, and management port switching, complete the following applicable procedures.

To configure distributed vSwitch networking for VMware vCenter server in your HCP environment, follow VMware best practices. Dedicated distributed vSwitches are required for each front-end switching, back-end switching, and management switching HCP network.

Configuring networking for front-end switching

Procedure

- From the vSphere Client, select an ESXi host.
- Click **Configure**.
- In the **Networking** section, click **Virtual Switches**.
- Click **Add Host Networking**.
- Select **Virtual Machine Port Group for a Standard Switch**, then click **Next**.
- Click **Select an existing standard switch**, then click **Browse**.
- Select the default **vSwitch0**, then click **OK**. Then click **Next**.
- For the **Network label**, enter **Front-end Network**, then click **Next**.
- Review your settings, then click **Finish**.

Configuring networking for back-end switching

Procedure

- Click **Add Host Networking**.
- Select **Virtual Machine Port Group for a Standard Switch**, then click **Next**.
- Click **New standard switch**, then click **Next**.
- Click the **Add adapters** icon.
- Select your configured back-end network adapter, then click **OK**. Then click **Next**.
- For the **Network label**, enter **Back-end Network**, then click **Next**.
- Review your settings, then click **Finish**.

(Optional) Configuring networking for management port switching

Procedure

1. Click **Add Host Networking**.
2. Select **Virtual Machine Port Group for a Standard Switch**, then click **Next**.
3. Click **New standard switch**, then click **Next**.
4. Click the **Add adapters** icon.
5. Select your configured management port network adapter, then click **OK**. Then click **Next**.
6. For the **Network label**, enter **Management Port Network**, then click **Next**.
7. Review your settings, then click **Finish**.

Chapter 4: Creating the HCP VM system

For general installation guidelines, before performing the HCP software installation on an HCP VM system, review the documentation for *Installing an HCP System*.

Unpacking and uploading the ISO zip file

Before you can create a new virtual machine, you need to unpack the `ISO.zip` archive and upload the ISO file into to a datastore.

Procedure

1. Download the `HS222_x.x.x.x.iso.zip` archive file onto your computer.
2. Unpack the `ISO.zip` archive to extract the ISO file.
3. Start your VMware vSphere client.
4. Click **Datastores**, then select the datastore to which you want to upload the ISO file.
5. Navigate to the datastore.
6. Click the **Upload a file to the Datastore** icon.
7. Locate the ISO file that you unpacked, then upload the file to the datastore.

Creating and configuring the new virtual machine

To create and configure the new virtual machine, you must complete the following procedures for each node in your HCP system. The procedures use the New virtual machine wizard in your VMware vSphere client. Make sure you have reviewed the configuration guidelines in [Configuration guidelines for the HCP VM environment \(on page 18\)](#).

Creating the virtual machine

Procedure

1. Log in to your vSphere client.
You should see the datacenters, clusters, and VMware ESXi hosts that were previously added to VMware vCenter.
2. Connect to the vCenter server where you configured the vSphere HA cluster for HCP VM.
3. In the navigation bar, right-click the ESXi host that you plan to use for the deployment and select the option to create a new virtual machine.
The **New virtual machine wizard** appears.

4. Select **Create a new virtual machine**, then click **Next**.
5. Enter a name for the new virtual machine, select a datacenter, then click **Next**.
6. Select a host, then click **Next**.
7. Select the datastore, then click **Next**.
8. Select the compatibility for the virtual machine, then click **Next**.



Note: Make sure that you select the version that corresponds with the version of ESXi host software that you are using.

9. For **Guest OS family**, select **Linux**.
10. For **Guest OS version**, select **Red Hat Enterprise Linux 6 (64-bit)**, then click **Next**.
11. Specify the following information:
 - For **CPU**, select **8**.
 - For **Cores per socket**, select **4**.
 - For **Memory**, specify **32 GB** or **64 GB**.
 - For **Hard disk 1** (the OS disk), specify at least **32 GB**.
 - For **Disk provisioning**, select **Thick provisioned, eager zeroed**.
 - Select **SCSI controller 0** and **SCSI (0:0)**.
12. Create additional VMDK or RDM data disks as needed.



Note: You must create a minimum of two additional hard disks, or three if you need a dedicated database volume.

- To create additional VMware VMDK hard disks:
 - a. From the **New device** menu, select **New Hard Disk**, then click **Add**.
 - b. For the first additional hard disk, select **SCSI (0:1)**. For the second additional hard disk, select **SCSI (0:2)**, and so on.
 - c. Specify a size of at least **500 GB** for a user database volume or at least **50 GB** for a dedicated database volume.
 - d. For **Disk provisioning**, select **Thick provisioned, eager zeroed**.
- To create additional RDM hard disks:
 - a. From the **New device** list, select **SCSI Controller**, then click **Next**.
For better performance, it is a best practice to use a separate SCSI controller for the RDM hard disk.
 - b. From the **New device** list, select **RDM Disk**, then click **Add**.
 - c. For the first additional hard disk, select **SCSI (1:0)**. For the second additional hard disk, select **SCSI (1:1)**, and so on.
 - d. Specify a size of at least **500 GB**.
 - e. For **Disk provisioning**, select **Thick provisioned, eager zeroed**.

Configure the network adapters

You must configure at least two network adapters. The first adapter is for the front-end network, and the second adapter is for the back-end network.

Procedure

1. Create the network adapter for the front-end network.
 - a. In the **New virtual machine** wizard, from the **New Network** list, select the front-end network you defined in [Configuring networking \(on page 30\)](#).
 - b. Specify the following information:
 - Set the first network adapter to the **Front-end Network**.
 - Select **Connect At Power On**.
 - Set **Adapter Type** to **VMXNET 3**.
 - Set **MAC Address** to **Automatic**.
2. Create a network adapter for the back-end network:
 - a. From the **New device** menu, select **Network**, then click **Add**.
 - b. Specify the following information:
 - Set the second network adapter to the **Back-end Network**.
 - Select **Connect At Power On**.
 - Set **Adapter Type** to **VMXNET 3**.
 - Set **MAC Address** to **Automatic**.
3. (Optional) Create a network adapter for the HCP management port network.
 - a. From the **New device** list, select **Network**, then click **Add**.
 - b. Specify the following information:
 - Set the third network adapter to the **Management Port Network**.
 - Select **Connect At Power On**.
 - Set **Adapter Type** to **VMXNET 3**.
 - Set **MAC Address** to **Automatic**.
4. From the **New CD/DVD Drive** list, select **Datastore ISO file**.
5. Next to **CD/DVD Media**, click **Browse**, select the ISO file that you have unpacked and uploaded in [Unpacking and uploading the ISO zip file \(on page 32\)](#), and then click **OK**.
6. Select **Connect At Power On**.
7. Click **Next**.
8. Review your settings, then click **Finish**

Result

The new virtual machine is created.

Delete the ISO image connection

After the new virtual machine is created, delete the connection to the ISO image.



Important: Failure to delete the connection to the ISO image can lead to accidental reinstallation of the appliance operating system, or nodes failing to start.

Procedure

1. In the **New Virtual Machine** wizard, right-click the virtual machine that is connected to the ISO image, then click **Edit settings > VM Hardware**.
The **Edit Settings** page opens to the **Virtual Hardware** tab.
2. Expand **New CD/DVD Drive <drive number>**.
3. Clear the **Connect At Power On** option.
4. Clear the **Connected** check box in the same row as **New CD/DVD Drive <drive number>**.
5. Click the **x** icon next to the **Connected** check box.
The connection to the ISO image is deleted.

Installing the Appliance Operating System

After deploying the ISO file, you must complete the following procedure for each HCP VM node in the vSphere cluster.

Before you begin

You will need the front-end IP addresses, network mask, default gateway and back-end IP addresses from the network administrator at your site. All back-end IP addresses must be on the same subnet. For easier installation and support, request that the last octet of the front end and back end be sequential.

Procedure

1. Power on the first node.
2. Start the vSphere Client.
3. In the left navigation bar, right-click the lowest-numbered node, then click **Open Console**.
The installation program prompts for the installation mode.

```

Please select boot device:
-----
UEFI: Built-in EFI Shell
(Bus 03 Dev 00)PCI RAID Adapter
FLASH Drive SM_USB20 1100
Enter Setup
-----
↑ and ↓ to move selection
ENTER to select boot device
ESC to boot using defaults

```

4. Either press **Enter** or let the program default to the installation option after 75 seconds. The installation program prompts for the procedure you want to perform.

```

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

```

Type your selection and press enter [pce]:

5. Enter **c** to clear existing storage volumes.
6. In response to the confirmation 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):

7. When prompted, enter the front-end network IP mode for the node. The IP mode that you specify is used to set both the system-level IP mode and the [hcp_system] network IP mode. Valid responses are IPv4, IPv6, and Dual .

Enter the front-end network IP mode (I[IPv4], IPv6,Dual):

8. When prompted, enter **y** to indicate that you want to provide a VLAN ID for the [hcp_system] network, or **n** to indicate that you do not want to provide a VLAN ID.

Do you want to provide a VLAN ID for the front-end network [n]:

9. If you entered **y** when prompted, enter the VLAN ID for the [hcp_system] network. Valid values are integers in the range 0001 through 4094. The VLAN ID you specify can include leading zeroes but cannot be more than four digits long.

Enter the front-end network VLAN ID [0000]:

If you entered **n**, the installation program does not prompt you to enter a VLAN ID.

10. If you entered **IPv4** or **Dual**, specify the IPv4 node IP address, subnet mask, and gateway IP address for the front-end network.

- a. When prompted, enter the IPv4 address assigned to the node for the front-end network.

```
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]:
    --->
```

11. If you entered IPv6 or Dual, respond to the additional configuration prompts.
12. When prompted, enter the back-end network IP address for the node.

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

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

13. Respond to the prompt as appropriate:
- To confirm your responses and start the installation, enter `y`.
 - To change any of your responses, enter `n`. In this case, the installation program repeats the prompts, starting with the front-end network bonding mode.

When the installation is complete, the node restarts and displays the Appliance OS login window on the console. The node is ready for the HCP software installation.



Note: If the node does not restart automatically, the Appliance OS installation failed. Contact your authorized HCP service provider for help.

14. Complete the previous steps on each node of the HCP VM system.

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 the example 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 vSphere client.
2. In the left navigation bar, select the node with the highest last octet in its back-end IP address.
3. Right-click the VM and then click **Open Console**.
4. Log in to the HCP VM node console with the default login information:

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

5. Change the password to `Hcpinstall_<HCP system's serial number>` (the last two characters in `Hcpinstall` are the numeral one). For example, if the system serial number is 12345, type the following: `Hcpinstall_12345`

If you choose a different password, be sure it is at least 14 characters long and contains at least one uppercase character, one lowercase character, one number, and one special character. Changing the `install` user password is a one-time event. The password has no expiration configured. If you decide to change the password again, you must wait at least 7 days after changing the password before you can change it again.



Note: Upgraded clusters may have the password set to `hcpinstall` (the last two characters are the numeral one).

6. Press **Enter** to display the **HCP Configuration Menu**.
Your screen will specify the HCP version that you are running.

```
HCP Configuration Menu
=====

[1] Get HCP Setup Files
[q] Log Out

Currently installed version:  None
Version on CD/DVD:           9.1.0.122
Extracted version:           None

Enter a selection:
```

7. Enter 1.
8. Press **Enter** to extract the files.

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

```

[1] Get HCP Setup Files
[q] Log Out

Currently installed version:  None
Version on CD/DVD:           9.1.0.122
Extracted version:          None
Enter a selection: 1
You chose: "1" Is this correct? [Default: yes]:
Preparing to extract HCP Core package.
Getting files:      [#####] 100%

All files retrieved. Press ENTER to continue: _

```

9. Press **Enter** to return to configuration menu.

```

HCP Configuration Menu
=====

[1] Get HCP Setup Files
[2] Configure HCP Network
[3] Install an HCP System
[q] Log Out

Currently installed version:  None
Version on CD/DVD:           9.1.0.122
Extracted version:          9.1.0.122

Enter a selection: _

```

10. Press 3.

The **New Install Menu** in the HCP Setup wizard 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]:

```

Step 1: Identify the nodes in the HCP system

Procedure

1. In the HCP Setup wizard **New Install Menu**, enter 1.

```
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]:
```

When you enter 1, 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]:
```

2. 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 shows a default value that is one greater than the previous IP address that you entered.

3. From the **HCP Nodes Menu**, enter b to return to the **New Install Menu**.

```
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
```

```

[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: 2]:

```

Step 2: Configure the HCP system

Procedure

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

```

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

[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: 2]:

```

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
[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 automatic DNS zone management (Yes)
[2] Domain Name for the System (None)
[3] Upstream DNS Server(s)

[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: HCP1example.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. Enable **[set to Yes] Option 1** if automatic DNS zone management is required.

12. To configure the corporate upstream DNS server(s), choose option 3, and specify the list of IPv4 and IPv6 addresses the HCP needs to use to forward DNS requests to each of your corporate servers.



Note: Each DNS server IP address that you specify must be routable from the [hcp_system] network.



Tip: If the front-end network IP mode is set to `Dual`, specify both IPv4 and IPv6 DNS server IP addresses. Doing so ensures that all client requests can use the HCP system domain name for access to the system, regardless of which types of IP addresses each client is configured to use.

13. If Option 3, **Upstream DNS Server(s)**, is not set to the proper corporate DNS server, change it accordingly.

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

14. Enter `b`.
15. Enter `5` to configure the time settings.
16. Enter `1` and set the time configuration to a time server. Use the same time server that has been configured for all ESXi hosts in the HCP VM system.

You previously configured the time server for all hosts in [Enabling NTP for the ESXi hosts \(on page 23\)](#).

```
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]:
```

17. 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):

Note: Control-C cancels input.
```

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

18. Enter `y` or `yes` to confirm the change.

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

19. Enter `6` to change the internal configuration settings.

When you enter, `6`, the **Internal Configuration Settings** menu appears.

```
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 (United States:
      (800) 446-0744. Outside the United States: (858) 547-4526

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

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

20. 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]:
```

21. Type `internal`.

22. Press **Enter** in response to the confirming prompt.

Optionally, to configure a dedicated database volume, the system must have at least three drives per node. Also, the dedicated database volume size must be at least 50 GB for a new installation. All dedicated database volumes must be the same size. HCP Setup asks if you want to configure a dedicated database volume only if your system meets these requirements.

23. If HCP Setup asks whether you want to configure a dedicated database volume, enter `yes` to configure a dedicated database volume or `no` if you do not want to configure a dedicated database volume.
24. Press **Enter** to confirm your choices and return to the **Internal Configuration Settings** menu.
25. From the **Internal Configuration Settings** menu, enter `2` to set the serial number for the HCP system.

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

Please enter 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
```

26. Enter the unique serial number for this HCP system.
27. 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.

28. 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.

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

Step 3: Run the HCP installation

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 are running the installation as the `install` user, the wizard informs you that data-in-flight encryption is enabled. The wizard then 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) =
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 =
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

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.
5. Complete one of the following actions:
 - If the configuration is not correct:
 - a. Enter `n` or `no`.
 - b. In response to the confirmation prompt, enter `y` or `yes`.
 - c. Correct the configuration information.
 - If the configuration is correct:
 - a. Enter `y` or `yes`.
 - b. In response to the confirmation prompt, enter `y` or `yes`.

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 following substeps:
 - a. When prompted, select the dedicated database volume for the first node.
 - b. Press **Enter** to confirm your selection.
 - c. Repeat the previous two substeps for each node in the system.
 - d. After you select the dedicated database volumes for each node, HCP Setup confirms your selections, then 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. This can take from several minutes to several hours, depending on the size of the logical volumes.

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. After you enter the key, HCP Setup completes the installation. You do not get a second chance to see the key, and it is not stored for later retrieval.

When the installation is complete, HCP Setup logs you out and restarts the nodes. The console then displays the login prompt.

If HCP Setup exits before installation processing is complete, make a note of all error messages. Then contact your authorized HCP service provider for help.

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

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

```

Every 30.0s: /sbim/system-info                               Fri Mar  1 12:29:58

Hostname:                               hcp-node-150.cluster-colo-009-
vml.lab.archives.com
RIS Mode:                                150
[hcp_system] IP:
[hcp_system] Mask:
[hcp_system] Gateway:
[hcp_backend] IP:
[hcp_backend] Mask:
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` in the URL, 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 appears.



Tip: If the browser cannot find the System Management Console login page, wait a few minutes, 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 the serial number is incorrect, contact your authorized HCP service provider.
4. Log in to the System Management Console with the default username and password.

```
Username: security
Password: Chang3Me!
```

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

- If the **Hardware** page is displayed, the nodes are still starting HCP. This process can take several minutes. When more than half of the nodes have completed the startup process, the Console displays the **Change Password** page.
 - 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 the default password.
 - b. In the **New Password** field, enter a new password.

Password criteria:

- Must contain UTF-8 characters, including white spaces
- 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@ssw0rd*
- Invalid password: *password*

- c. In the **Confirm New Password** field, type your new password again.
- d. Click **Update Password**.

6. In the top-level menu, click **Hardware**.
7. On the **Hardware** page, confirm the following status:
 - Node status: **Available**
 - Status of each logical volume: **Available**

To see the status of a logical volume, hover your cursor over the volume icon.

If all of the nodes and logical volumes are available, the installation was successful and you can begin creating tenants. However, you might not want to do this until all additional setup is complete.

If any node has a status other than Available, or if 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. Complete either of the following steps:
 - Set additional configuration options, as described in [Setting additional configuration options \(on page 51\)](#). You can set additional configuration options only if the installation was successful.
 - Log out of the System Management Console and close the browser window. This action ensures that no one else can use the Console without logging in.

Setting additional configuration options

After verifying that the HCP system was correctly installed, you can set additional configuration options, such as enabling the management port network, enabling syslog logging, or disabling ping.

Procedure

1. Log into the HCP System Management Console as the security user.

2. Create a new user account with the administrator role.
Alternatively, you can add the administrator role to the security user account, then go to step 4.
3. Log out of the Administration Console. Then log in again using the new account with the administrator role.
4. Perform the configuration activities.
5. 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 VMware environment. To maintain performance and data integrity, HCP VM system hardware must be monitored for failures outside of the virtual machine environment.

Hitachi servers and network components that are part of the HCP VM system can be connected to Hitachi Remote Ops for monitoring. For more information, see [Configuring HCP monitoring with Hitachi Remote Ops \(on page 72\)](#).

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 are 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.

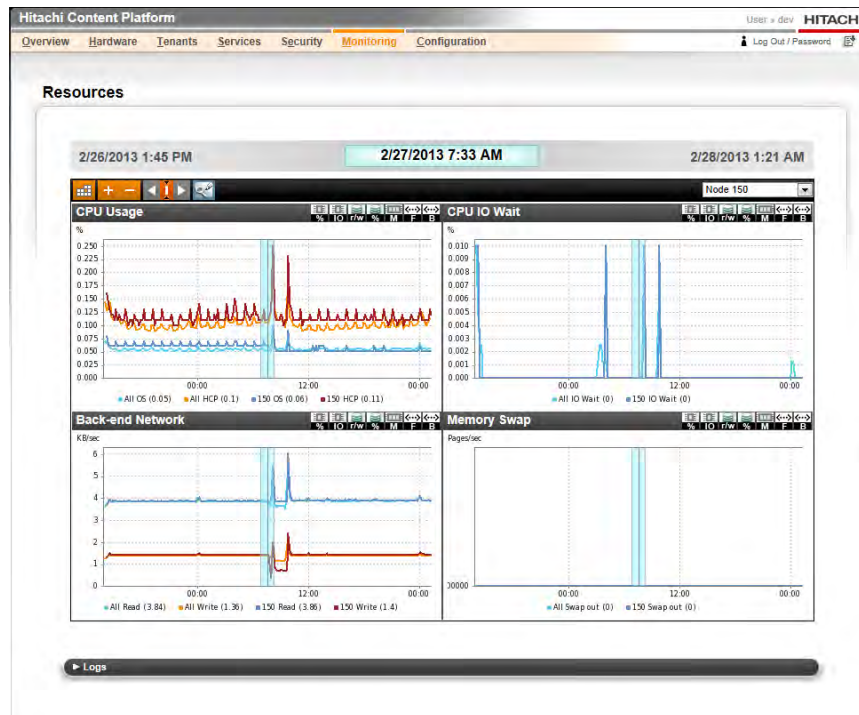
You can generate chargeback reports to track system capacity and bandwidth usage at the tenant and namespace levels.

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

For more information, see [Configuring HCP monitoring with Hitachi Remote Ops \(on page 72\)](#).

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

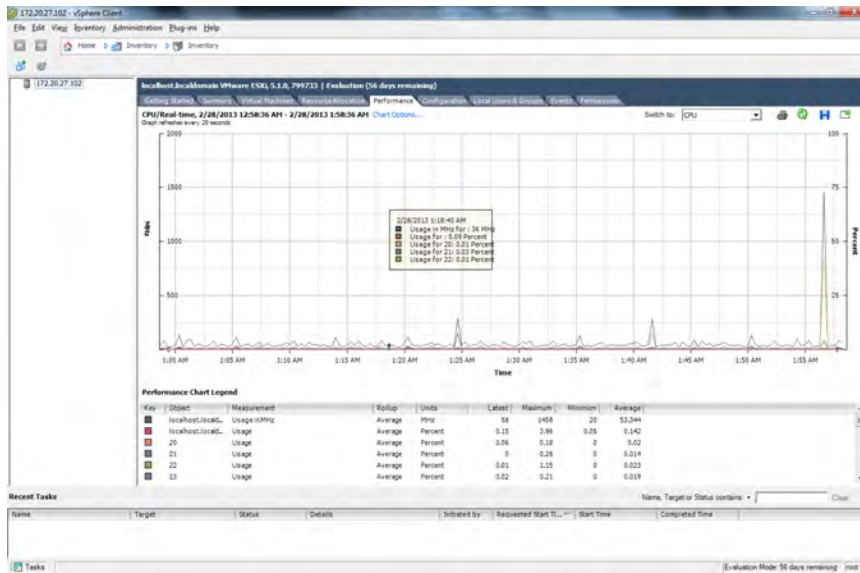
For any HCP VM node, you can run diagnostics that analyze and resolve issues with interactions between nodes and other components of the HCP environment.

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*.

VMware monitoring and performance are your responsibility. In the vSphere center, under the performance tab, clients have multiple ways to monitor resources.



For more information about VMware vSphere monitoring options, see the [VMware web site](#).

Chapter 5: Maintaining the HCP VM environment

After an HCP VM environment has been established, your site can complete environment maintenance tasks, such as moving storage node databases, adding HCP VM nodes, and recovering storage nodes.

Adding logical volumes

Procedure

1. Provision the LUNs to be added to each ESXi host in the system. See [Provisioning HCP VM storage \(on page 25\)](#) for details.
2. Add the LUNs to new datastores, one LUN per datastore. See [Adding VMFS datastores to a vSphere HA cluster \(on page 28\)](#).
3. From the VMware Sphere client, right-click the HCP VM to which you want to add capacity and select **Edit Settings**.
4. In the **Virtual Machine Properties** window, click **Add**.
5. In the **Add Hardware** window, select **Hard Disk**, then click **Next**.
6. Select **Create a new virtual disk**, then click **Next**.
7. Set the capacity to be slightly smaller than the size of the provisioned LUN. VMware adds a small amount of overhead.
8. Select **Thick Provision Eager Zeroed**.
9. Browse to the new datastore to be added, then click **Next**.
10. Select the next available SCSI disk in the **Virtual Device node** section, then click **Next**.
11. Verify the selected options, then click **Finish**.
12. In the **Virtual Machine Properties** window, verify that the new hard disk is listed, then click **OK**.
13. In the vSphere client, open the virtual machine console for the highest numbered storage node.
14. Log in as the install user.
The HCP Configuration Menu appears. Your menu will specify the HCP version that you are running.

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

Enter a selection:

```

15. 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]:

```

16. Enter **1** to add storage to the HCP System while it is online.

```

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

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

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

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

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

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

```

17. Enter **yes** to continue the procedure.
 18. 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

```

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

Typically, this shows storage added to all nodes.

(Optional) For a dedicated database volume to be configured, the volume size must be at least 50 GB and at least 1.5 times the size of the existing database for each node. All dedicated database volumes must be the same size.

If you already have a dedicated database volume, any newly added dedicated database volume must be larger than the current one.

20. 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 6 11:51:59 2020 Current status:
    node 150: 53% Complete (7/13): Running formatDrives
Fri Mar 6 11:52:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 27% complete
    (0 / 1 volumes)
Fri Mar 6 11:53:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 35% complete
    (0 / 1 volumes)
Fri Mar 6 11:54:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 42% complete
    (0 / 1 volumes)
Fri Mar 6 11:55:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 99% complete
    (0 / 1 volumes)
Fri Mar 6 11:56:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 100% complete
    (1 / 1 volumes)
Fri Mar 6 11:56:25 2020 Current status:
    node 150: 76% Complete (10/13): Running create_volume_config
Fri Mar 6 11:56:30 2020 Current status:
    node 150: 84% Complete (11/13): Running start_new_volumes
Fri Mar 6 11:56:46 2020 Current status:
    node 150: 92% Complete (12/13): Running
    sync_new_local_volumes
Fri Mar 6 11:57:01 2020 Current status:
    node 150: 100% Complete: Storage addition complete
Fri Mar 6 11:57:22 2020 Current status:
    node 150: 100% Complete: Starting new volumes
Fri Mar 6 11:57:27 2020 Current status:
    node 150: 100% Complete: Starting new volumes

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

```

- To select a dedicated volume for the database, enter `yes`. Then complete the following substeps:
 - a. 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 three previous substeps for each node in the system.
21. After you have selected the dedicated database volumes for each node, HCP Setup confirms your selections and then asks if you want to continue the procedure. Enter **yes** to continue the procedure.

```

Do you want to select a dedicated volume for database?
[Default: no]: yes
Do you want to select a new dedicated PG LUN for node 001?
[Default: no]: yes
Enter a selection for node 001 [1, 2]: 2
You chose: "2. /dev/sde at 2:0:0:4 (1TB)", is this correct?
[Default: yes]: yes
Do you want to select a new dedicated PG LUN for node 002?
[Default: no]: yes
Enter a selection for node 002: 2
You chose: "2. /dev/sde at 2:0:0:4 (1TB)", is this correct?
[Default: yes]: yes
Do you want to select a new dedicated PG LUN for node 003?
[Default: no]: yes
Do you want to select a new dedicated PG LUN for node 004?
[Default: no]: yes
This will add the new volumes and move database to the following dedicated
volumes. Do you want to continue?
    node 001: 2:0:0:4 (1TB)
    node 002: 2:0:0:4 (1TB)
    node 003: 2:0:0:3 (500GB)
    node 004: 2:0:0:3 (500GB)
[Default: no]: yes

```

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

```

Syncing install password to all nodes.
Updating EULA
Syncing date to all nodes.
Syncing HCP package to all nodes
Starting to poll nodes for progress
Fri Mar  6 11:51:59 2020 Current status:
    node 150: 53% Complete (7/13): Running formatDrives
Fri Mar  6 11:52:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 27% complete
    (0 / 1 volumes)
Fri Mar  6 11:53:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 35% complete
    (0 / 1 volumes)
Fri Mar  6 11:54:15 2020 Current status:
    node 150: 53% Complete (7/13): Formatting 42% complete
    (0 / 1 volumes)
Fri Mar  6 11:55:15 2020 Current status:

```

```

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

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

```

22. When the formatting is complete, press **Enter** to continue.
23. 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 confirm your entry, 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)
-----
              | pgdata | /RIS/archive33 | /RIS/archive94
              | pgidx  | (450M/1006)    | (1806/2006)
              | pgxlog | /RIS/archive34 | /RIS/archive95
              |      | (450M/1006)    | (1806/2006)
              | pgxlog | /RIS/archive33 | /RIS/archive94
              |      | (350M/1006)    | (1506/2006)

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

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

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

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

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

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

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

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

9. 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 confirmation prompt, enter `y` or `yes` to confirm your entry `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] 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.

- 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:

```

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

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

Adding HCP VM nodes

Before you begin

- Add new ESXi hosts or find existing ESXi hosts that can support an HCP node. For more information, see about creating ESXi hosts, see [Configuring the HCP VM environment \(on page 23\)](#).
- Unpack and upload the ISO files to the selected ESXi hosts. For more information, see [Unpacking and uploading the ISO zip file \(on page 32\)](#).
- Create the new virtual machine. For more information, see [Creating and configuring the new virtual machine \(on page 32\)](#).
- Configure the HCP VM network on the newly deployed HCP VM nodes. For more information, see [Installing the Appliance Operating System \(on page 35\)](#).
- From the highest active HCP VM node, run the Add Node service procedure. For more information, see *Installing and Maintaining an HCP System*.

Procedure

- From the **HCP Configuration** menu, enter `4` to run the HCP Setup wizard.
- 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** appears.

```

HCP Setup: Membership Update Menu
=====

```

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

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

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

[q] Return to Configuration menu

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

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

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

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

Note: Control-C cancels input.

Enter yes or no.
[Default: no]:
```

4. In response to the confirmation 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 confirmation 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) =
Allow Data at Rest Encryption = No
Customer Support Contact Information = United States:
(800) 446-0744. Outside the United states: (858) 547-4526
Storage configuration = internal
Gateway Router IPv4 Address =
Encrypt Data at Rest on Primary Storage = No
Time Settings Compliance Mode = No
HCP System Serial Number = 00001
MOE Index-only volumes = No
Enable DNS = Yes
Chassis = None
Enable Replication on This System = Yes
Multicast Network = 238.172.59.42
Time Zone = America/New_York
```

```

Current Date and Time = None
Domain Name for the System = hcp.example.com
Allow Data in Flight Encryption / SSL = Yes
Blade Servers = No
Distributor/OEM Key Access = Arizona
Time Server(s) = internal
Gateway Router Secondary IPv6 Address = None
Gateway Router IPv6 Address = None
Spindown Volumes = No
HCP Storage Nodes: 1

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 confirmation prompt, enter `y` or `yes`.
 - c. Exit the wizard and contact your authorized HCP service provider 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 must be at least 50 GB. All dedicated database volumes must be the same size.

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

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

```

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

- If the prechecks are successful, the HCP software is installed on the new nodes. For RAIN and VM systems, the software is installed on four nodes at a time. For SAIN systems, the software is installed on one cross-mapped pair of nodes at a time.

After the software installation is complete, the nodes are restarted. When the node addition is complete, the **HCP Configuration Menu** appears.

- 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 before the node addition processing is complete, contact your authorized HCP service provider for help.

7. From the **HCP Configuration Menu**, enter `q` to log out of the install shell.

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

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 confirmation 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 confirmation 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 the Configuration Menu

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

5. From the **Node Recovery Menu**, take one of these actions:

- To recover selected storage nodes, enter 1. Then, follow the on-screen instructions to identify the nodes you want to recover. Make sure to use the back-end IP address to identify each node.



Note: If you choose to use a range of IP addresses to identify the nodes, make sure 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]:
```

- Enter `y` or `yes` to delete the database while recovering the OS or `n` or `no` to recover the OS without deleting the database.
- In response to the confirmation prompt, enter `y` or `yes` to confirm your entry or `n` or `no` to try again.
- (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 restarts all the nodes that it recovered and displays this message:

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

7. If the node that you are logged in to is one of the recovered nodes, the SSH or console session ends automatically when HCP restarts the node. If this is not the case, in response to the prompt to continue, press `Enter`.

The **HCP Service Menu** appears.



Important: If HCP Setup exits at any time before the OS recovery processing is complete, contact your authorized HCP service provider. 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 confirmation 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 confirmation prompt, enter `y` or `yes` to confirm your entry or `n` or `no` to try again.

Chapter 6: 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:
 - a. Select the **Enable SNMP at `snmp.hcp-domain-name`** option.

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

1. Log in to Hitachi Remote Ops. See [Step 1: Log in to Hitachi Remote Ops \(on page 73\)](#)
2. Set the Hitachi Remote Ops base configuration, including the email addresses to which email about error conditions should be sent. [Step 2: Set the base configuration \(on page 74\)](#)
3. Conditional: Configure transport agents for reporting error conditions to Hitachi Vantara support personnel. See [Step 3: \(Conditional\) Configure transport agents \(on page 74\)](#)
4. Identify the HCP system to be monitored. See [Step 4: Identify the HCP system \(on page 75\)](#)

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. Make sure to change the default password as soon as possible.
If Hitachi Remote Ops is already in use at your site for monitoring other devices, this password might 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 ID for your site, 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 email about the error condition alerts.

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 appears.
2. In the **Device Monitoring** section:
 - a. In the **Site ID** field, type your Hitachi Vantara site ID. If you do not know your site ID, contact your authorized HCP service provider.
 - b. (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:
 - a. In the **eMail Server** field, type the fully qualified hostname or a valid IP address of the email server through which you want Hitachi Remote Ops to send email about error conditions.
 - b. In the **Local Interface** field, select the Ethernet interface that has connectivity to the specified email server. This is the interface on the Hitachi Remote Ops server.
 - c. In the **User List** field, type a comma-separated list of the email addresses to which Hitachi Remote Ops should send email about error conditions.
 - d. In the **Sender's Email Address** field, type an email address to be used in the From line of each email.

For some email servers, the value in the From line must 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

A Hitachi Remote Ops transport agent transfers notifications of error conditions to a target location for access by your authorized Hitachi Vantara service provider. 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 appears in the list of transport agents. A set of configuration fields appears 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, then click **Refresh**.
2. Take one of these actions:
 - If the **Summary** page does not display any tables, an **Add a device link** appears. 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 appears.
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 site ID. If you do not know your site 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.
14. Ensure that **Enabled** is selected.
15. Ensure that **Trace** is cleared.
16. Click **Add**.
When the procedure is successful, a message appears indicating that the HCP system has been added.



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

Appendix A: Changing the HCP VM network adapters

This appendix describes how to change the supported network adapters on your HCP VMs.

About network adapters

Starting with HCP v7.2.1, all newly installed HCP VMs are automatically configured to use VMXNET3 adapters.

If you are upgrading from a previous HCP release, your configuration might include e1000 adapters. Update your ESXi configuration to use VMXNET3 adapters, because HCP support for e1000 adapters is deprecated.

VMXNET3 adapters support both 1- and 10 gigabit network configurations, whereas e1000 network adapters support only 1 gigabit network configurations.

Best practice is to use VMXNET3 network adapters in all new and upgraded configuration.

Disabling LRO on the ESXi host for VMXNET3

To switch your HCP VMs to use the VMXNET3 network adapter, you must disable large receive offload (LRO) in the guest operating system to prevent potential TCP performance degradation.

Procedure

1. Log in to the vSphere Client.
2. Click a server that hosts ESXi for your HCP VMs.
3. Click the **Configure** tab.
4. In the **System** menu, click **Advanced System Settings**.
5. In the inventory tree, click **Edit**.
6. Scroll down to the following parameters and change their parameter field from 1 to 0:
 - Net.Vmxnet2HwLRO
 - Net.Vmxnet2SwLRO
 - Net.Vmxnet3HwLRO
 - Net.Vmxnet3SwLRO
 - Net.VmxnetSwLROSL
7. Click **OK**.

- From the vSphere client, restart the server by right-clicking it, then selecting **Reboot**.

Changing the HCP VM network adapter

You can configure an HCP VM to use VMXNET3 by removing the existing e1000 network adapter and replacing it with VMXNET3.

Step 1: Power off the HCP VM

Procedure

- Open the VMware vSphere Client.
- Right-click the HCP VM for which you want to replace the network adapters, then hover over **Power**, and then click **Shut Down Guest OS**.

Step 2: Remove the previous network adapters

Procedure

- From the vSphere Client, right-click one of the powered-off HCP VMs.
- Click **Edit Settings**.
The **Virtual Machine Properties** window appears. The number of existing network adapters varies depending on whether the HCP VM is using e1000 or VMXNET3:
 - If the HCP VM is using e1000, remove four network adapters.
 - If the HCP VM is using VMXNET3, remove two network adapters.
- Click the **Virtual Hardware** tab.
- Select a network adapter, then click **Remove**.
- Repeat step 4 until all of the network adapters are removed.
- Click **OK**.

Step 3: Configure the front-end network adapters

Before you begin

If you are using VMXNET3 adapters, the first VMXNET3 adapter must be connected to the front-end network, and the second VMXNET3 adapter must be connected to the back-end network.

Procedure

- From the vSphere eClient, right-click one of the powered-off HCP VMs, then click **Edit Settings**.
The **Virtual Machine Properties** window appears.
- On the **Hardware** tab, click **Add**.
The **Add Hardware** window opens.

3. Select **Ethernet Adapter**, then click **Next**.
4. On the **Network Connection** page, in the **Adapter Type** pane, click the **Type** list, then select **VMXNET3**.
5. In the **Network connection** pane, select **Named network with specified label**.
6. Select **Front-End Network**, then click **Next**.
7. On the **Verification** page, click **Finish**.
8. Back in the **Virtual Machine Properties** window, click **OK**.

Step 4: Configure the back-end network adapters

Before you begin

The first VMXNET3 network adapter must be connected to the front-end network, and the second VMXNET3 network adapter must be connected to the back-end network.

Procedure

1. From the vSphere eClient, right-click the powered-off HCP VM, then click **Edit Settings**.
The **Virtual Machine Properties** window appears.
2. In the **Edit Settings** window **Hardware** tab, click **Add**.
The **Add Hardware** window opens.
3. Select **Ethernet Adapter** then click **Next**.
4. On the **Network Connection** page, in the **Adapter Type** pane, click the **Type** list, then select **VMXNET3**.
5. In the **Network connection** pane, select **Named network with specified label**.
6. Select **Back-End Network**, then click **Next**.
7. On the **Verification** page, click **Finish**.
8. Back in the **Virtual Machine Properties** window, click **OK**.

Step 5: Power on the HCP VM

Procedure

1. From your vSphere eClient, right-click the newly configured HCP VM.
2. Hover your cursor over **Power**, then click **Power On**.

Result

When the HCP VM is powered on, you have successfully changed its network adapters. If you have multiple HCP VM nodes that need to be reconfigured, repeat the changing network adapter procedure for each one.

Appendix B: Modifying the DRS settings

Procedure

1. Open the vSphere Client.
2. In the left navigation bar, select the datacenter.
3. Click the **Getting Started** tab.
4. Click **Create a cluster**.
5. In the **VMware Cluster Wizard**, select **Turn On vSphere HA** and **Turn On vSphere DRS**.



Important: Enable this feature only if you fully understand its functionality.

6. Click **OK**.
7. For **DRS automation level**, select **Manual** to specify where VM guests should reside.
8. For **Power management**, select **Off**.
9. Select **Enable Host Monitoring** and keep the default settings.
10. Set **VM Monitoring** to **Disabled**.
11. Select **Disable EVC**.
12. Select where you want to store your `swapfile` location.
13. Review your settings, then click **OK**.
14. In the left navigation bar, right-click a cluster, then click **Edit Settings**.
15. In the **Settings** window, on the left navigation bar, click **DRS Groups Manager**.
16. Create a group and add the virtual machines that you want to keep on the server.
17. Create one virtual machine DRS group for each host.
18. In the **Host DRS Groups** section, click **Add** and place one host from the cluster in each group. Then click **Next**.
19. In the **Settings** window, in the left navigation bar, click **Rules**.
20. Create a new rule to match each VM group to a host group, and set the type of rule to **Virtual Machines to Hosts**.
21. Select **Should run on hosts in group**, then click **OK**.

This setting creates a rule that lets VMs run on other hosts in the cluster in the event of a failure.



Important: If you select **Must Run on Hosts in Group**, HA will not start the server on another host in the cluster, which defeats the purpose of HA.

22. Set an alarm to alert you if a failure occurs:
 - a. Right-click the cluster and go to **Alarm > Add Alarm**.

- b. In the **Alarm Settings** window, name your alarm and set **Monitor** to **Virtual Machines**.
- c. Select **Monitor for specific event occurring on this object**.
- d. Click the **Triggers** tab.
- e. Select **VM is violating a DRS VM Host affinity rule**.
- f. Set the status to either **Warning** or **Alert**.
- g. In the **Trigger Conditions** section, from the **Argument of VM** list, select a name.
- h. Set the value for each VM you want to monitor.
- i. Add one argument for each VM.
- j. Set the actions you want the system to take.

Appendix C: Configuring the HCP VM small instance

If you are deploying an HCP VM system as a small-instance system, you must change the CPU count and RAM for each node before powering on the HCP VM nodes.

Procedure

1. In vSphere eClient, right-click the HCP VM node, then click **Edit Settings**.
2. In the **Virtual Machine Properties** window, select the **Hardware** tab.
3. From the **Hardware** list, select **Memory** and then, in the **Memory Configuration** pane, adjust the allocation to 16 GB.
4. From the **Hardware** list, select **CPUs** and adjust the **Number of virtual sockets** and **Number of cores per sockets** so that the **Total number of cores** equals 4.
5. Click **OK**.

The **Virtual Machine Properties** window closes.

Appendix D: Managing failover

The HCP VM vSphere HA cluster does not automatically move a failed-over HCP VM node back to its original ESXi host after the host is available. Instead, the HCP VM system administrator must manually shut down the HCP VM nodes that need to be moved to another ESXi host.

Alternatively, the vCenter administrator can shut down the HCP VM nodes from the vCenter management console. Then the administrator can manually move the HCP VM nodes onto the preferred ESXi host and power on the nodes. After an HCP VM node starts, it rejoins the HCP VM system

After you shut down an HCP VM node and trying to move the VM to another ESXi host, you might see an error message. You can ignore the message.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact