Hitachi Content Platform

8.1

Deploying an HCP-VM System on ESXi

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

MK-94HCP010-01
June 2018

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 Corporation 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 Corporation.

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

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

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

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

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

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

Contents

Preface

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

Intended audience

This book is intended for the people responsible for deploying an HCP-VM system on a VMware vSphere® environment at a customer site. It assumes you have experience with computer networking, creating virtual machines, familiarity with VMware products and concepts, and a basic understanding of HCP systems.

Product version

This book applies to release 8.1 of HCP.

Release notes

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

https://knowledge.hitachivantara.com/Documents

Related documents

The following documents contain additional information about Hitachi Content Platform:

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

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

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

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

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

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

- *Deploying an HCP-VM System on KVM* — This book contains all the information you need to install and configure an HCP-VM system. The book also includes requirements and guidelines for configuring the KVM environment in which the system is installed.

- *Third-Party Licenses and Copyrights* — This book contains copyright and license information for third-party software distributed with or embedded in HCP.

- *HCP-DM Third-Party Licenses and Copyrights* — This book contains copyright and license information for third-party software distributed with or embedded in HCP Data Migrator.

- *Installing an HCP RAIN System - Final On-site Setup* — This book contains instructions for deploying an assembled and configured HCP RAIN system at a customer site. It explains how to make the necessary physical connections and reconfigure the system for the customer computing environment. The book also provides instructions for assembling the components of an HCP RAIN system that was ordered without a rack and for configuring Hi-Track Monitor to monitor the nodes in an HCP system.

- *Installing an HCP SAIN System - Final On-site Setup* — This book contains instructions for deploying an assembled and configured single-rack HCP SAIN system at a customer site. It explains how to make the necessary physical connections and reconfigure the system for the customer computing environment. It also contains instructions for configuring Hi-Track® Monitor to monitor the nodes in an HCP system.

Accessing product documentation

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

Getting help

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

Hitachi Vantara Community 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.

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

Comments

Please send us your comments on this document:

HCPDocumentationFeedback@HitachiVantara.com

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

Thank you!

1

HCP system overview

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

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

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

HCP-VM system components and architecture

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

The figure below shows the architecture of an HCP virtual machine (HCP-VM) system running on running on VMware infrastructure.

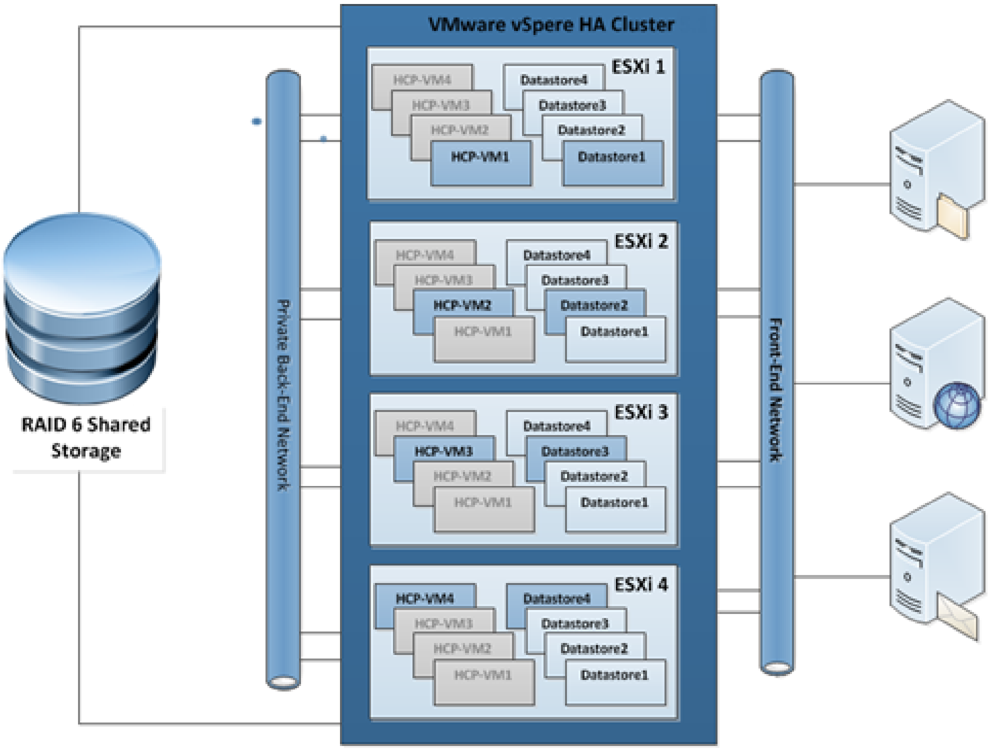

Host platform

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

Compute

An HCP-VM node must have at least eight virtual CPUs and 32 gigabytes 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're deploying an HCP-VM small-instance configuration, each HCP-VM node must have at least four virtual CPUs and 16 gigabytes of allocated RAM.

Storage

HCP-VM storage infrastructure is highly available and fault tolerant. It is recommended for the physical servers that the ESXi hosts run on to be connected to shared **SAN** storage with **RAID6** protection or Hitachi NAS (**HNAS**).

The HCP SAN storage needs to have at least two paths to each Logical Unit Number (**LUN**) and each LUN needs to 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 G10 system 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 recommended RAID6, shared SAN storage configuration, and HNAS datastores, HCP-VM also supports the following for storage configuration:

- Shared SAN arrays with virtual volumes created from Hitachi Dynamic Provisioning (DP) pools. This configuration does not support thin provisioning. It is recommended to spread datastores across multiple DP Pools to avoid resource contention and single points of failure.
- 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.
- Other devices like HNAS that export NFS v3 shares, which are mounted and used as NFS datastore.

  - It is required to use thick, eager zero when formatting NFS datastores, so additional ESXi plug-ins may be required from your vendor. Hitachi provides a VAAI plug-in that enables this functionality on the HNAS platform.

  - It is recommended to not have multiple datastores on the same file system or the same underlying disk due to performance and availability considerations.

  - Follow the vendors 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.

If you deviate from the recommended 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 on supported storage vendors and devices, see the applicable VMware documentation.

HCP network connectivity

HCP-VM network connectivity is provided to the HCP guest OS by VMware VMXNET3 or e1000 vNICs, VMware vSwitches, and dvSwitches. It is recommended that the vNICs connect to a single vSwitch for Back-end connectivity and a single vSwitch for Front-end connectivity. For 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. For e1000 vNICs, the Back-end vSwitch must be configured to provide access to two vmNICs and the Front-end vSwitch must be configured to provide access to a different set of two vmNICs. The vmNICs are setup for NIC teaming for failover by default.

Tip: NIC Teams are multiple physical network adapters sharing a single vSwitch and the physical network. NIC teams provide passive failover if there is a hardware failure or network outage. In some configurations they can increase performance by distributing the traffic across physical network adapters.

Front-end network

The HCP front-end network is used for client and management access. For HCP front-end networks, it is recommended that the ESXi host create two virtual Network Interface Cards (**NIC**) on a second pair of physical NICs. Having two physical NICs dedicated to HCP ensures redundancy and consistent performance.

Back-end network

HCP private back-end network is used for internode communication and data transfer. The ESXi host has two vmNICs which directly map to two physical NICs (pNICs) 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 1GbE and dedicated to HCP.

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

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

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

Management port network

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

Storage network

Hitachi Vantara recommends that the VMkernel network be set up in a private network or with a unique VLAN ID that provides 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 HCP-VM system needs to be connected to Hi-Track monitor. For more information on Hi-Track monitor, see Chapter 6: "Configuring HCP monitoring with Hi-Track Monitor" on page 79.

Hi-Track monitor does not support non-Hitachi hardware. To monitor hardware supplied by other vendors, you need to use third-party monitoring tools.

HCP software

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

Because HCP-VM software is not bound to hardware, the software does not support zero copy failover and hardware cannot be monitored in the system management console.

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 search nodes

HCP search has reached end of service life, therefore HCP search nodes are not available for HCP-VM systems. As with physical HCP systems, this functionality is provided by Hitachi HDDS Enterprise search products.

HCP-VM node failover (vCenter and vSphere HA)

If you wish to set up automatic failover in the event of an ESXi host failure, HCP-VM requires an instance of the VMware vCenter server to be available in the customer environment for enabling HCP-VM node failover. Failover functionality is provided by a vSphere HA cluster.

A vSphere High Availability (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.

Each ESXi host participating in an HCP-VM system will be configured to be part of a single vSphere HA cluster in vCenter. This enables high availability in cases where one or more servers or ESXi hosts fail. When the master host detects a server or ESXi host failure, it can restart the HCP-VM node that was running on the server or ESXi host that failed on other healthy ESXi hosts in the cluster.

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

The HCP-VM vSphere HA cluster will not be configured to automatically move the failed-over HCP-VM node back to its original ESXi host once the server or ESXi host is available. The HCP-VM system administrator will manually shutdown the HCP-VM node, and the vCenter administrator will manually move the HCP-VM node onto the preferred ESXi host and power on the HCP-VM node. Once the node boots, it will re-join 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. In this case, 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. Just like 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 can either be a pre-existing server in the customer environment, or can be allocated as part of the HCP-VM HA cluster of ESXi hosts. It is recommended (but not required) that the vCenter server be separate from the HCP-VM HA cluster. The vCenter server can consume a fair amount of resources on the ESXi host which could be utilized 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 is to 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 two terabytes of active and HCP S Series storage. The basic storage license also provides two terabytes of extended storage. If you need additional storage, please contact your Hitachi Vantara sales representative.

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

Configuration guidelines for the HCP-VM environment

This chapter describes the requirements and recommendations for the successful installation and operation of an HCP-VM system.

Supported VMware versions

HCP-VM supports multiple versions of VMware. For a list of supported VMware versions, see *HCP 8.1 Release Notes*.

VMware supported functionality

HCP-VM supports the following VMware functionality:

- vSphere HA cluster
- The VMware tools package is included in the HCP OS with HCP-VM. This lets the HCP-VM node shutdown from the vCenter management console. Pausing live migration, and other functionality enabled by the inclusion of the tools package are **not** currently supported.
- **DRS** may be used in a manual capacity to assist with VM to host affinity as described in Appendix B: "Changing the DRS settings" on page 93.
- Other failover capabilities provided by VMware such as vMotion, storage vMotion, DRS and FT are **not** supported by this version of HCP-VM.

HCP-VM does not support software used for VM replication.

The following HCP features are specific to the physical HCP appliances (HCP RAIN system and HCP SAIN system) and are not applicable to HCP-VM through alternate means:

- **Autonomic Tech Refresh**: Provides the capability of migrating a VM to a different host, this allows for server refresh. The raw storage layer is obscured from HCP in the VMware environment; any storage refresh would need to be handled at the VMware layer.

- **Zero Copy Failover**: VMware HA replaces this capability by restarting an HCP guest VM on a running ESXi host after it is lost due to an ESXi host failure. This ZCF-like storage availability is provided by shared SAN storage.

- **Spindown**, **IDX** (indexing) only: 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 via their preferred mechanism.

- **VLAN tagging**: VMware's active-active NIC Teaming is designed for load balancing and redundancy. Both physical NIC's must be configured with the same VLAN tagging. Also, VMware vSwitch is a layer 3 switch and will not route traffic out 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 the appropriate physical NIC.

Prerequisites and recommendations

HCP-VM systems can be configured in two ways: standard and small instance. In order to deploy a standard HCP-VM system, you need:

- A shared SAN storage, RAID6 (Recommended) system

- A minimum of 3.66 TB of usable storage space

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

- A minimum of four HCP-VM nodes

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

Note: Do not commit more than 256GB of RAM an HCP-VM node. Over committing RAM can slow HCP-VM system performance.

- NFS datastores: Recommended Volume Size:
  - As discussed in VMware NFS Best Practice: "The following statement appears in the VMware Configuration Maximums Guide: "Contact your storage array vendor or NFS server vendor for information about the maximum NFS volume size." When creating this paper, we asked a number of our storage partners if there was a volume size that worked well. All partners said that there was no performance gain or degradation depending on the volume size and that customers might build NFS volumes of any size, so long as it was below the array vendor's supported maximum. This can be up in the hundreds of terabytes, but the consensus is that the majority of NFS datastores are in the tens of terabytes in terms of size. The datastores sizes vary greatly from customer to customer."
- Two physical NICs on each ESXi host in the vSphere HA cluster dedicated to the HCP-VM back-end network
- Two physical NICs for the VMware management network for vSphere HA and HCP-VM front-end network
- Two port fibre channel HBA cards (or VMware compatible IO device) for shared SAN storage connectivity (when applicable)
- ESXi requires a minimum of 2 GB of physical RAM. VMware recommends providing at least 8 GB of RAM to take full advantage of ESXi features and run virtual machines in typical production environments.

HCP-VM small instance configuration has the same prerequisites and recommendations as the standard configurations for 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 MQE querying or object and directory counts above published maximums.

HCP-VM system limits

An HCP-VM system standard configuration supports the following maximum values:

- 40 HCP-VM nodes
- 59 data LUNs on each HCP-VM node (ESXi guest OS limitation)
- Max open VMDK storage per host (ESXi Limitation)
    - 5.0 update 2: 60 TB
    - 5.1 update 2: 60 TB
    - 5.5: 128 TB
    - 6.0: 128 TB
- HCP-VM supports 2 TB VMDK for 5.0 and 5.1
- HCP-VM supports 16 TB VMDK for 5.5 and 6.0

For more information on HCP supported limits, see *HCP 8.1 Release Notes*.

An HCP-VM system small instance configuration supports the following maximum values:

- 16 HCP-VM nodes
- 59 data LUNs on each HCP-VM node

HCP-VM availability considerations

An HCP repository considered in a state of continuous availability if there is one HCP-VM node per ESXi host and if one more than half of the HCP-VM nodes are healthy and running.

If 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.

If your HCP system is in a state of continuous availability, the HCP system can survive a single ESXi host failure without affecting HCP functionality.

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 ESXi hosts CPU, RAM or disk can cause HCP system instability.

A vSphere HA cluster is recommended, but not required.

3

Configuring the HCP-VM environment

This section provides the information and steps required to provision the VMware environment to be ready for an HCP-VM deployment.

ESXi considerations

A customer may want to deploy the HCP-VM system on existing ESXi hosts in their environment. Before attempting to do this, make sure the hosts meet the minimum requirements for compute and memory cataloged in "Prerequisites and recommendations" on page 10.

Depending on the versions of ESXi, vCenter, and/or vSphere that the customer is using, various steps and terminology can appear differently in the GUI. This document is most applicable to version 6.5 but can be used as a reference guide for previous versions. For documentation specifically relating to ESXi versions 6.0 and earlier, see *Deploying an HCP-VM on ESXi* (MK-94HCP010-00) from version 8.0 of HCP. For a list of supported VMware versions, see *HCP 8.1 Release Notes*. For VMware functionality that HCP-VM supports and does not support, see "VMware supported functionality" on page 9.

When provisioning storage for use with HCP-VM, be sure to review and follow the ESXi Storage Guide (for version 6.5: vSphere Storage for ESXi 6.5 and vCenter Server 6.5) as well as the relevant storage vendor's VMware best practices guide. For more information on provisioning storage for use with HCP-VM, see "Provisioning HCP-VM storage" on page 19.

All ESXi hosts that will contain an HCP-VM node must have Network Time Protocol (NTP) enabled. This is done with the Time Configuration option in the vSphere client on each individual ESXi host.

Important: NTP must be enabled for each ESXi host individually.

Enabling NTP for the ESXi hosts

To configure ESXi hosts for NTP:

1. Access your vSphere client.
2. In your vSphere client, select the ESXi host for which you want to enable NTP.
3. Click on **Configure**.
4. Under **System**, click on **Time Configuration**.
5. Click on **Edit**.
6. In the window that appears, you can specify if you want to manually configure the date and time for the host or if you want to enable the NTP client. If you want the NTP service to synchronize the date and time of the host with an NTP server, select the option that enables the NTP client.
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 on **OK** and then start or restart the NTP service.
10. Repeat the procedure with the same time server for all ESXi hosts that will have an HCP-VM node.

Tip: Write down the NTP server used in your ESXi hosts so you can use it for the HCP-VM installation.

Configuring a vSphere HA cluster for HCP-VM (Recommended)

A 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.

Creating a datacenter

To configure a vSphere HA cluster for HCP-VM, you first need to create a datacenter. To create a datacenter:

1. Access the vSphere client.
2. In the left side navigation bar, right-click on your server and select **New Datacenter**.
3. In the window that appears, enter a name for your HCP-VM datacenter. Here is a good example name: `HCP-VM_center_1`.
4. Click on **OK**.

Adding a cluster

After you have created a datacenter, you need to add a cluster. To add a cluster:

1. In the left side navigation bar, right-click on the datacenter you created and select **New Cluster**.
2. In the window that appears, enter a name for the cluster. Here is a good example name: `hcp-vm-cluster-1`.
3. Select the checkbox to turn on vSphere HA.

Important: Do **not** click on the option to turn on DRS.

Note: DRS can be turned on later to define VM affinity to a particular host or group of hosts. This function does not provide further automation of failover. The settings described merely assist with keeping VMs on a particular host, and alert if the rule cannot be followed. For details on the settings required, see "Appendix B: Changing the DRS settings" on page 93.

4. Click on **OK**.

Adding ESXi hosts

After you have added a cluster, you can add ESXi hosts. To add ESXi hosts to the cluster:

1. In the left side navigation bar, right-click on the cluster you created and select **Add Host**.

2. In the **Add Host Wizard**, enter the ESXi host connection information.

3. Enter the ESXi host Username and Password.

4. Review the host summary.

5. Enter the license information for the ESXi host if it does not have any assigned.

6. Optionally, select a lockdown mode if you want to prevent remote users from logging in directly.

Note: The decision to implement a lockdown mode should be made by the customer.

7. Review your choices then click on **Finish** to add the ESXi host to the vSphere HA cluster.

8. Repeat the above procedural steps on adding ESXi hosts for all of the other ESXi hosts in the system.

Note:

- The number of ESXi hosts cannot exceed 32 (vSphere 5.0/5.1/5.5 HA cluster limitation).
- If the number of hosts exceeds 32, a second vSphere HA cluster needs to be created with the same settings in the same instance of vCenter.
- The ESXi hosts should be balanced between the two clusters.
- At this point, all hosts could have an alert that there aren't enough heartbeat datastores.
    - This can be verified by clicking on the host, selecting the **Summary** tab, and observing the **Configuration Issues** at the top of the page.

Provisioning HCP-VM storage

When provisioning storage for use with HCP-VM, be sure to review and follow the ESXi Storage Guide (for version 6.5: vSphere Storage for ESXi 6.5 and vCenter Server 6.5) as well as the relevant storage vendor's VMware best practices guide.

Its possible to provision HCP-VMs in a local storage or shared SAN storage configuration. Local storage is not recommended due to its increased data availability risk. For that reason, it is recommended to set your Data Protection Level (**DPL**) to two on a local storage configuration. For more information on DPL, see *Administering HCP*.

The following are guidelines for provisioning shared SAN storage for use with HCP-VM with the recommended configuration:

- Datastores used for HCP-VM nodes must be backed by shared RAID6 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, it is recommended to 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 diagram below illustrates a sample SAN layout for VMDK and RDM. The number of storage controller ports dedicated to an HCP-VM system is dependent on the capabilities of the storage array. For Hitachi Vantara mid-range storage, the best practice is to spread host access across all cores.

Consult the storage vendor documentation for sizing and configuration options.

Fibre Channel Connectivity

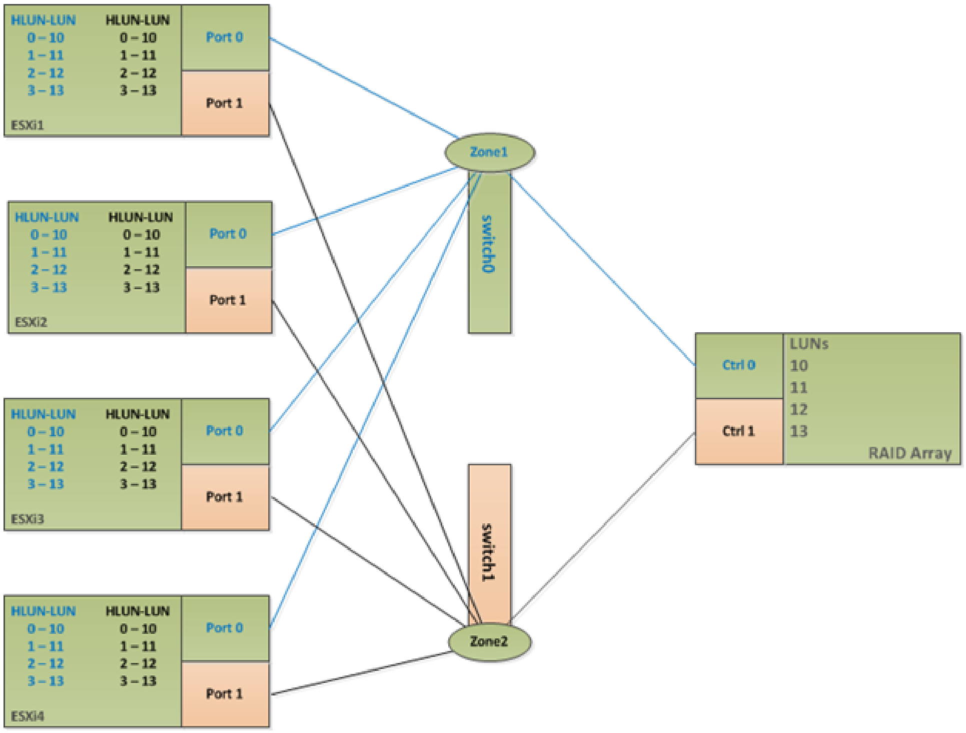

FC Switch 1, HCP-VM path 1

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 charts are sample HostGroup / 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 booting from the LUN should be granted access.

Array path 1

Host Group Name	Hosts	HLUN	ArrayLUN	VMware datastore
HCP_VM_cluster_1_path_1	ESXi-1 ESXi-2 ESXi-3 ESXi-4	1	10	hcp-vm_cluster-1_node_1_datastore_1
		2	11	hcp-vm_cluster-1_node_2_datastore_1
		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

(Continued)

HCP_VM_ cluster_1_path_ 2	ESXi-1 ESXi-2 ESXi-3 ESXi-4	1	10	hcp-vm_cluster-1_node_1_ datastore_1
		2	11	hcp-vm_cluster-1_node_2_ datastore_1
		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

It is recommended to have 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.

Here is a visual depiction of the cluster layout.

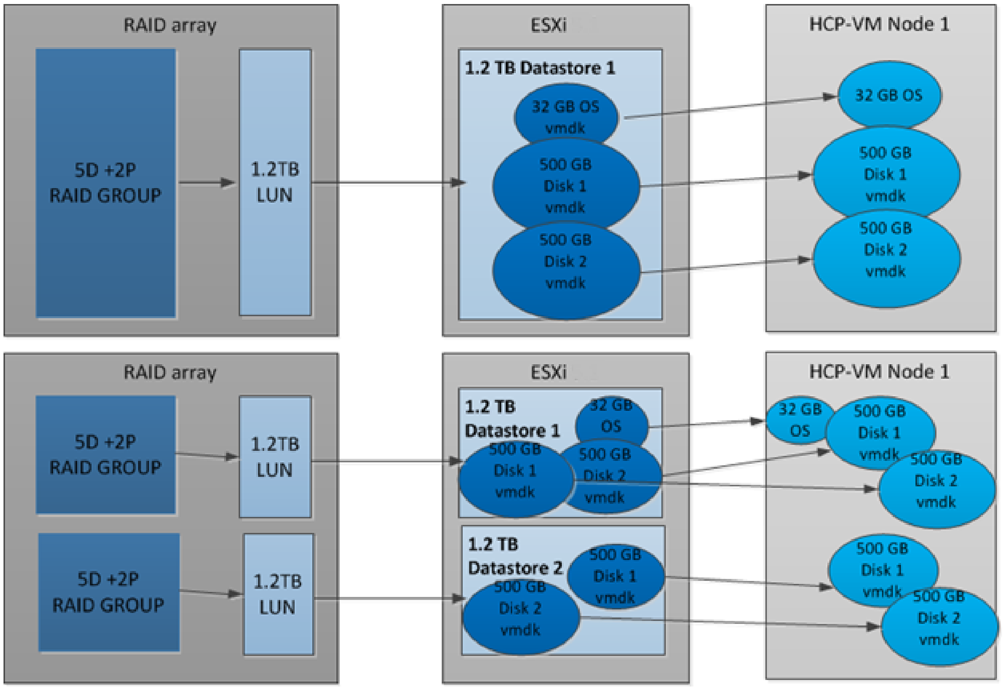

To add VMFS datastores to vSphere HA clusters:

1. Access your vSphere Client.
2. In the **Datastores** section, click on the option to create a new datastore.
3. To specify the datastore type, select **VMFS**.
4. Enter a meaningful name for the datastore. A good example name is: `hcp-vm_cluster_1_node_1_datastore_1`.
5. Select the VMFS version for the datastore (VMFS 6 is recommended).
6. Specify partition configuration details for the datastore.
7. Review your choices then click on **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 a variety of different ways. To achieve the best performance, follow these recommendations for configuring HNAS in a VMware vSphere environment:

- In general, a 4 KB file system block size is recommended. 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 utilized 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) LU\LDEV per RAID group consuming all space (if possible).

Creating an NFS datastore

To set up an NFS datastore, follow these steps:

1. Access your vSphere Client.
2. In the **Datastores** section, click on the option to create a new datastore.
3. To specify the datastore type, select **NFS**.
4. Enter a meaningful name for the datastore. A good example name is: `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 require access to the datastore.
8. Review your choices then click on **Finish** to create the NFS datastore.

Important: Ensure that you mount datastores with the same volume label on all vSphere ESXi hosts within VMware high availability (HA) environments.

Configuring networking

Networks should be configured for particular switches in the system before the ISO file is deployed. Make sure to review *Administering HCP* for the latest information on Network Administration in HCP.

To configure networking for front-end and back-end switching:

Configuring networking for front-end switching

1. From the vSphere Client, select an ESXi host.
2. Click on **Configure**.

3. In the **Networking** section, click on **Virtual Switches**.

4. Click on **Add Host Networking**.

5. Select **Virtual Machine Port Group for a Standard Switch** then click on **Next**.

6. Click on **Select an existing standard switch**, then click on **Browse**.

7. Select the default **vSwitch0**, then click on **OK**. Then click on **Next**.

8. For the **Network label**, enter **Front-end Network** then click on **Next**.

9. Review your settings then click on **Finish**.

Configuring networking for back-end switching

1. Click on **Add Host Networking**.

2. Select **Virtual Machine Port Group for a Standard Switch** then click on **Next**.

3. Click on **New standard switch**, then click on **Next**.

4. Click on the **Add adapters** icon.

5. Select your configured back-end network adapter, then click on **OK**. Then click on **Next**.

6. For the **Network label**, enter **Back-end Network** then click on **Next**.

7. Review your settings then click on **Finish**.

Optionally, you can configure networking for management port switching. To configure networking for management port switching:

Configuring networking for management port switching

1. Click on **Add Host Networking**.

2. Select **Virtual Machine Port Group for a Standard Switch** then click on **Next**.

3. Click on **New standard switch**, then click on **Next**.

4. Click on the **Add adapters** icon.

5. Select your configured management port network adapter, then click on **OK**. Then click on **Next**.

6. For the **Network label**, enter **Management Port Network** then click on **Next**.

7. Review your settings then click on **Finish**.

4

Creating the HCP-VM system

For general installation recommendations, prior to 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 and upload the ISO file into a datastore. To unpack the ISO.zip and upload the ISO file into a datastore:

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

Creating the new virtual machine

Before you create the new virtual machine, make sure you have reviewed the configuration guidelines in Chapter 2: "Configuration guidelines for the HCP-VM environment" on page 9.

The following procedural steps all utilize the **New virtual machine** wizard of your vSphere client. You need to perform these steps for all of the nodes in the HCP system. For more information on applicable versions and considerations that can apply to these procedural steps, see "ESXi considerations" on page 15.

Creating the virtual machine

To begin the process for creating a new virtual machine in vSphere:

1. Log in to your vSphere client and connect to the vCenter server where you configured the vSphere HA cluster for HCP-VM.

   Once you have logged in to your vSphere client, you should see the datacenters, clusters, and ESXi hosts that were previously added to vCenter.

2. In the navigation bar, right-click on the ESXi host to target for the deployment and select the option to create a new virtual machine.

   The **New virtual machine** wizard appears.

3. In the **New virtual machine** wizard, select **Create a new virtual machine** then click on **Next**.

4. Enter a name for the new virtual machine and select a datacenter, then click on **Next**.

5. Select a host, then click on **Next**.

6. Select the datastore, then click on **Next**.

7. Select the compatibility for the virtual machine, then click on **Next**.

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

8. For the **Guest OS family**, select **Linux**.

9. For the **Guest OS version**, select **Red Hat Enterprise Linux 6 (64-bit)**, then click on **Next**.

10. Specify the following information:

    ○ For **CPU**, select **8**.

    ○ For **Cores per socket**, select **4**.

- For **Memory**, specify at least **32 GB**.
- For **Hard disk 1** (this is the OS disk), specify at least **32 GB**.
- For **Disk provisioning**, select **Thick provisioned, eager zeroed**.
- Specify **SCSI controller 0** and **SCSI (0:0)**.

11. To create additional VM data disks (you need to create at least two additional hard disks or three if you want a dedicated database volume), do one of the following:

    - To create additional VMDK hard disks:

        1. In the **New device** dropdown menu, select **New Hard Disk** then click on **Add**.
        2. For the first additional hard disk, select **SCSI (0:1)**. For the second additional hard disk, select **SCSI (0:2)**, and so on.
        3. Specify a size of at least **500 GB** for a user data volume or at least **50 GB** for a dedicated database volume.
        4. For **Disk provisioning**, select **Thick provisioned, eager zeroed**.

    - To create additional RDM hard disks:

        1. In the **New device** dropdown menu, select **SCSI Controller** then click on **Add**.

           Using a separate SCSI controller for the RDM hard disk is recommended for better performance.

        2. In the **New device** dropdown menu, select **RDM Disk** then click on **Add**.
        3. For the first additional hard disk, select **SCSI (1:0)**. For the second additional hard disk, select **SCSI (1:1)**, and so on.
        4. Specify a size of at least **500 GB**.
        5. For **Disk provisioning**, select **Thick provisioned, eager zeroed**.

Configuring the network adapters

You need to configure at least two network adapters. The first is for the front-end network and the second is for the back-end network. These are the networks that you have defined in "Configuring networking" on page 25.

Before you add the second network adapter, you need to configure the first network adapter for the front-end network. To configure the network adapters:

1. In the **New virtual machine** wizard, click on the **New Network** dropdown menu, then click on the front-end network you have defined in "Configuring networking" on page 25.

2. Set the network adapter for the front-end network by specifying the following information:

   ○ Set the first network adapter to the **Front-end Network**.

   ○ Select the **Connect At Power On** option.

   ○ Set the **Adapter Type** to **VMXNET 3**.

   ○ Set the **MAC Address** to **Automatic**.

3. After you have configured the front-end network, create a network adapter for the back-end network. To create a network adapter for the back-end network, click on the **New device** dropdown menu and select **Network** then click on **Add**. Then specify the following information:

   ○ Set the second network adapter to the **Back-end Network**.

   ○ Select the **Connect At Power On** option.

   ○ Set the **Adapter Type** to **VMXNET 3**.

   ○ Set the **MAC Address** to **Automatic**.

4. Optionally, to create a network adapter for the HCP management port network, click on the **New device** dropdown menu and select **Network** then click on **Add**. Then specify the following information:

   ○ Set the third network adapter to the **Management Port Network**.

   ○ Select the **Connect At Power On** option.

- Set the **Adapter Type** to **VMXNET 3**.
- Set the **MAC Address** to **Automatic**.

5. Expand **New CD/DVD Drive**.
6. Select **Datastore ISO file** from the dropdown menu.
7. Next to **CD/DVD Media**, click on **Browse**.
8. Select the ISO file that you have unpacked and uploaded in "Unpacking and uploading the ISO zip file" on page 29, then click on **OK**.
9. Select the **Connect At Power On** option.
10. Click on **Next**.
11. Review your settings then click on **Finish** to create your new virtual machine.

Installing the Appliance Operating System

After deploying the ISO file, the following steps need to be performed for **all** HCP-VM nodes in the vSphere cluster. They must be done in this order:

1. Power on the first node.
2. Follow the configuration instructions below.
3. Repeat for the next node in the HCP-VM system.

Note: Before continuing with this procedure, you will need the front-end IP addresses, network mask, default gateway and back-end IP addresses from the network administrator at the customer site. All back-end IP addresses must be on the same subnet. For easier installations and support, request the last octet of the front-end and back-end be sequential.

To configure the HCP-VM network:

1. Access the vSphere Client.
2. In the left side navigation bar, right-click on the lowest numbered node and click on **Open Console**.

   The installation program prompts for the installation mode.

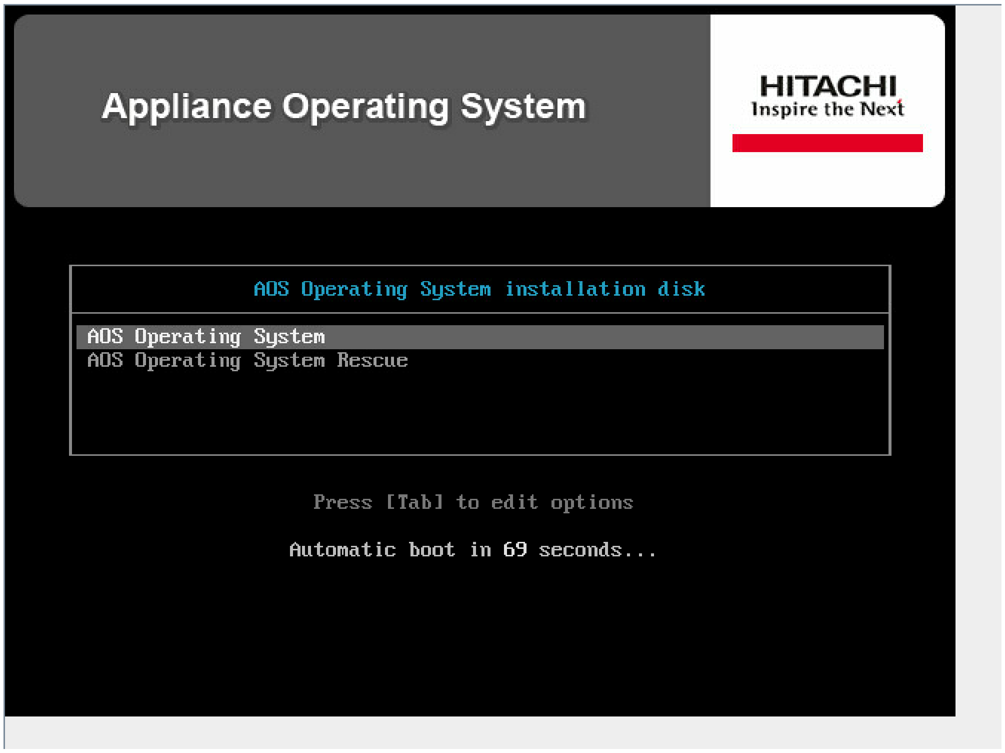

```
AOS Operating System installation disk

AOS Operating System
AOS Operating System Rescue

Press [Tab] to edit options

Automatic boot in 69 seconds...
```

3. 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.

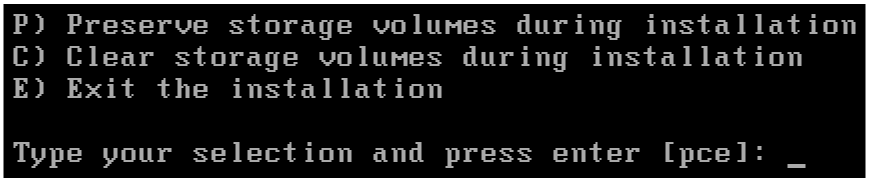

```
P) Preserve storage volumes during installation
C) Clear storage volumes during installation
E) Exit the installation

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

4. Enter *c* to clear existing storage volumes.

5. In response to the confirming prompt, enter *y*.

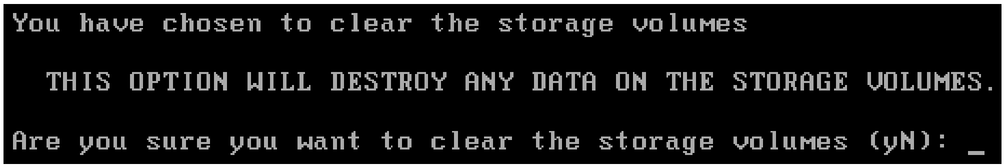

```
You have chosen to clear the storage volumes

  THIS OPTION WILL DESTROY ANY DATA ON THE STORAGE VOLUMES.

Are you sure you want to clear the storage volumes (yN): _
```

6. When prompted, enter the front-end network IP mode for the 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*.

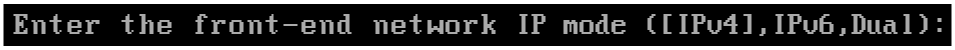

```
Enter the front-end network IP mode ([IPv4],IPv6,Dual):
```

7. 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 don't want to provide a VLAN ID.

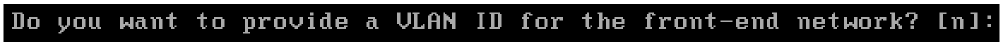

```
Do you want to provide a VLAN ID for the front-end network? [n]:
```

8. If you entered *y* in response to the previous step, when prompted, enter the VLAN ID for the [hcp_system] network. Valid values are integers in the range 0001 through 4,094. The VLAN ID you specify can include leading zeroes but cannot be more than four digits long.

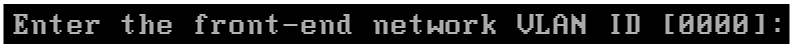

```
Enter the front-end network VLAN ID [0000]:
```

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

9. 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.

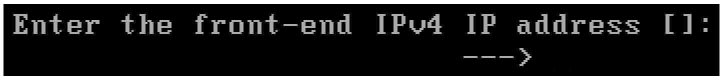

```
Enter the front-end IPv4 IP address []:
                         --->
```

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

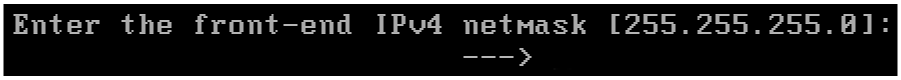

```
Enter the front-end IPv4 netmask [255.255.255.0]:
                         --->
```

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

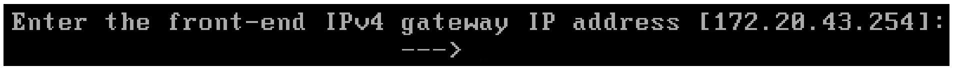

```
Enter the front-end IPv4 gateway IP address [172.20.43.254]:
                         --->
```

10. Optionally, if you entered *IPv6* or *Dual*, respond to the additional configuration prompts.

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

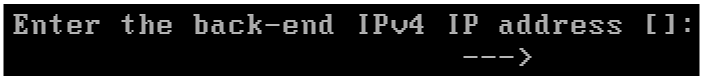

```
Enter the back-end IPv4 IP address []:
                                  --->
```

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

12. In response to the confirming prompt:

    - 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 again with the front-end network bonding mode.

When the installation is complete, the node reboots automatically. If the node does not reboot automatically, the OS installation failed. In this case, please contact your authorized HCP service provider for help.

Once the node reboots and shows the AOS login screen on the console, the node is ready for the HCP software installation.

13. Complete these procedural steps on each of the HCP nodes.

Installing the HCP software

The HCP install 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 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 utilize DARE features, please contact your authorized HCP service provider before performing the software installation.

To install the HCP software:

1. Access the vSphere client.

2. In the left side navigation bar, select the node with the highest last octet in its back-end IP address.

3. Right-click on the VM and click on **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 *hcpinsta11* (The last two characters are the number one).

6. Press Enter to display the HCP 8.1 Configuration Menu.

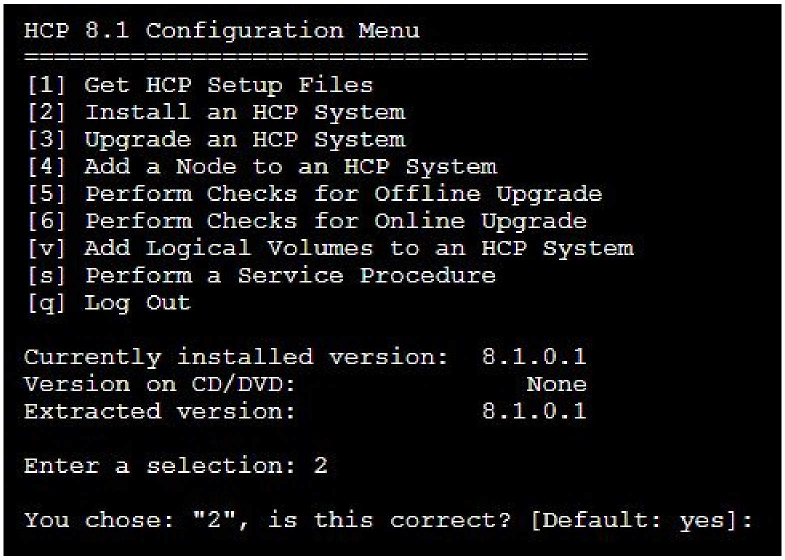

```
HCP 8.1 Configuration Menu
=====================================
[1] Get HCP Setup Files
[2] Install an HCP System
[3] Upgrade an HCP System
[4] Add a Node to an HCP System
[5] Perform Checks for Offline Upgrade
[6] Perform Checks for Online Upgrade
[v] Add Logical Volumes to an HCP System
[s] Perform a Service Procedure
[q] Log Out

Currently installed version:  8.1.0.1
Version on CD/DVD:               None
Extracted version:            8.1.0.1

Enter a selection: 2

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

7. Enter **2**.

8. In response to the confirming prompt, press Enter.

When you press Enter, the HCP Setup wizard **New Install** menu appears.

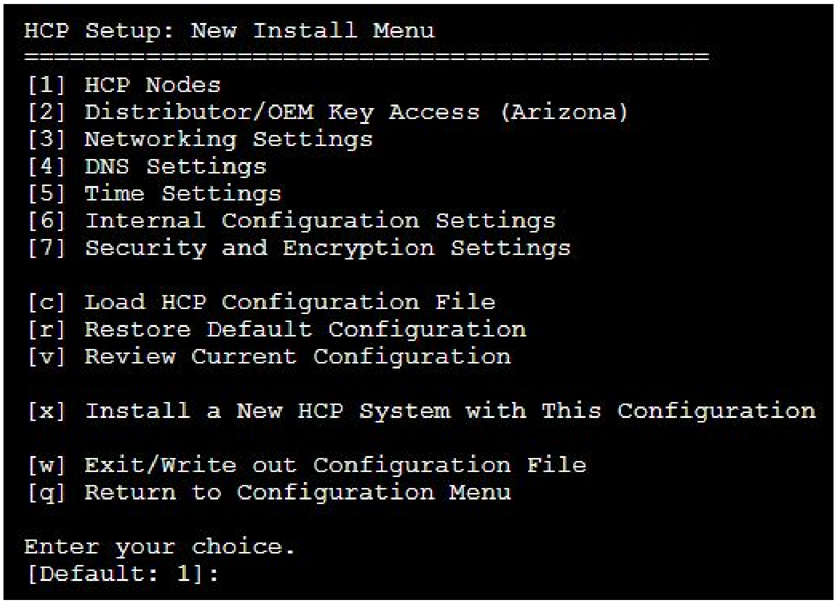

```
HCP Setup: New Install Menu
============================================
[1] HCP Nodes
[2] Distributor/OEM Key Access (Arizona)
[3] Networking Settings
[4] DNS Settings
[5] Time Settings
[6] Internal Configuration Settings
[7] Security and Encryption Settings

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

[x] Install a New HCP System with This Configuration

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

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

Step 1: Identify the nodes in the HCP system

To identify the nodes in the HCP system:

1. From the HCP Setup wizard **New Install** menu, enter **1**.

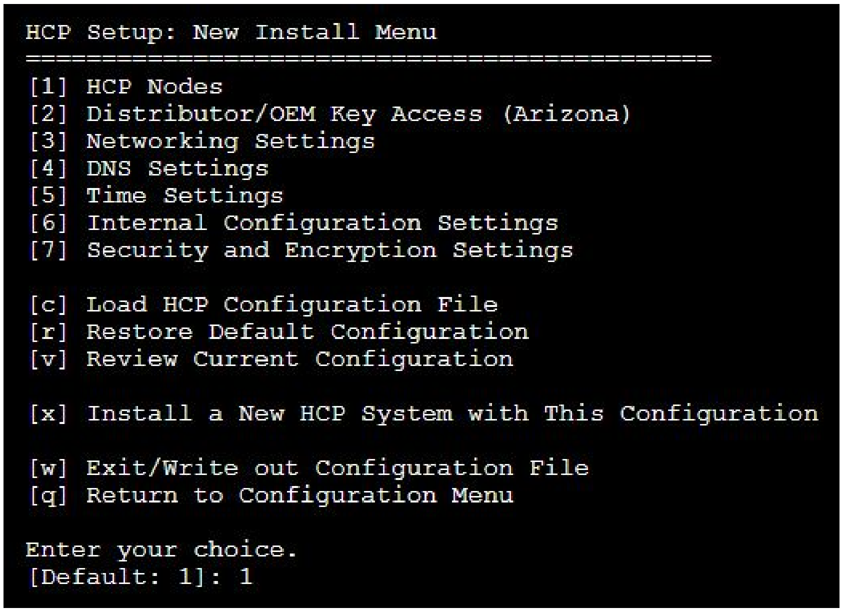

```
HCP Setup: New Install Menu
============================================
[1] HCP Nodes
[2] Distributor/OEM Key Access (Arizona)
[3] Networking Settings
[4] DNS Settings
[5] Time Settings
[6] Internal Configuration Settings
[7] Security and Encryption Settings

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

[x] Install a New HCP System with This Configuration

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

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

When you enter **1**, the **HCP Nodes** menu appears.

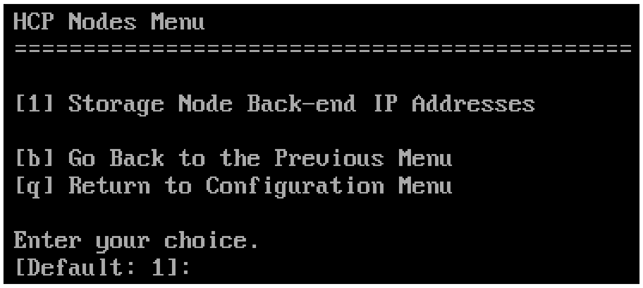

```
HCP Nodes Menu
==============================================

[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 presents a default value that’s 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.

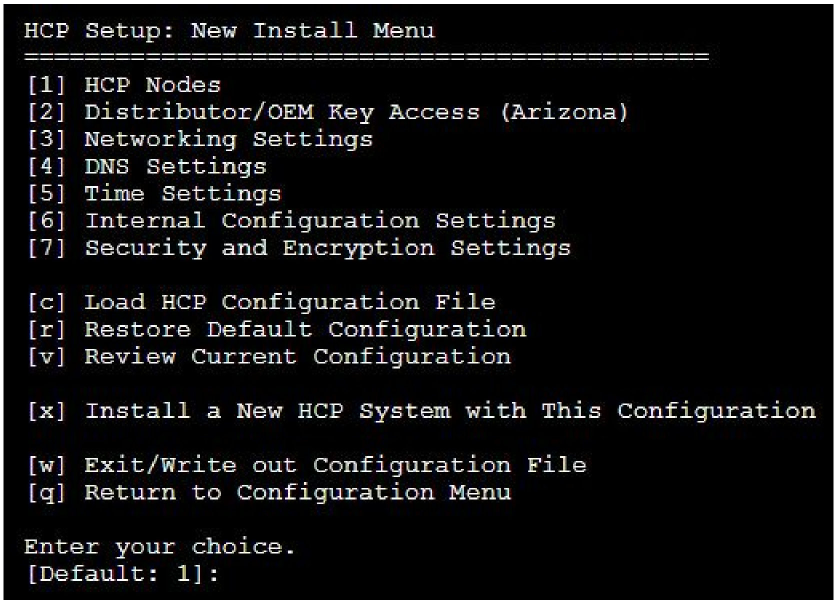

```
HCP Setup: New Install Menu
=============================================
[1] HCP Nodes
[2] Distributor/OEM Key Access (Arizona)
[3] Networking Settings
[4] DNS Settings
[5] Time Settings
[6] Internal Configuration Settings
[7] Security and Encryption Settings

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

[x] Install a New HCP System with This Configuration

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

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

Step 2: Configure the HCP-VM system

From the **New Install** menu, you can execute the additional options for configuring the HCP system. Each option either opens a lower-level menu with configuration options, or leads directly to a configuration option.

To configure the HCP system:

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

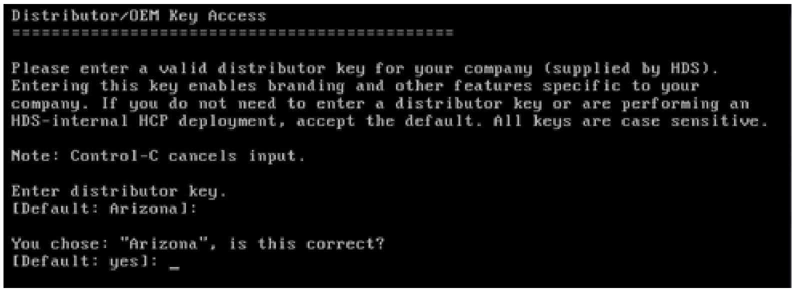

```
Distributor/OEM Key Access
============================================

Please enter a valid distributor key for your company (supplied by HDS).
Entering this key enables branding and other features specific to your
company. If you do not need to enter a distributor key or are performing an
HDS-internal HCP deployment, accept the default. All keys are case sensitive.

Note: Control-C cancels input.

Enter distributor key.
[Default: Arizona]:

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

2. Change the distributor key.

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

3. Enter *y* or *yes* to confirm the change and return to the **New Install** menu.

4. From the **New Install** menu, enter **3** to configure the networking options.

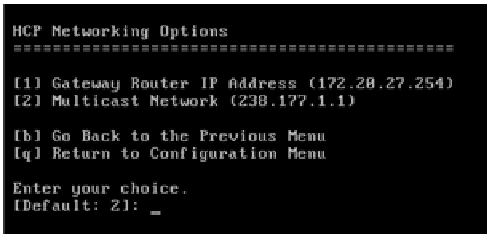

```
HCP Networking Options
============================================

[1] Gateway Router IP Address (172.20.27.254)
[2] Multicast Network (238.177.1.1)

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

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

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

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

7. Enter **b** to return to the **New Install** menu.

8. From the **New Install** menu, enter **4** to configure the DNS options.

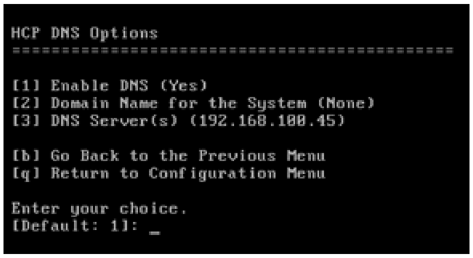

```
HCP DNS Options
============================================

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

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

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

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

10. Enter the system domain name.

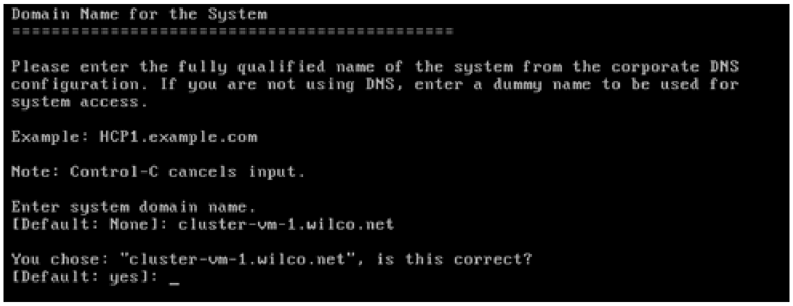

```
Domain Name for the System
==========================================

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

Example: HCP1.example.com

Note: Control-C cancels input.

Enter system domain name.
[Default: None]: cluster-vm-1.wilco.net

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

11. If **Option 1: Enable DNS** is not set to yes, change it to yes.

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

13. Enter **b** to return to the **New Install** menu.

14. From the **New Install** menu, enter **5** to configure the time settings.

15. Enter **1** and set the time configuration to a time server. Use the same time server that has been configured for all ESXi hosts in the HCP-VM system.

    This was set up in "Enabling NTP for the ESXi hosts" on page 16.

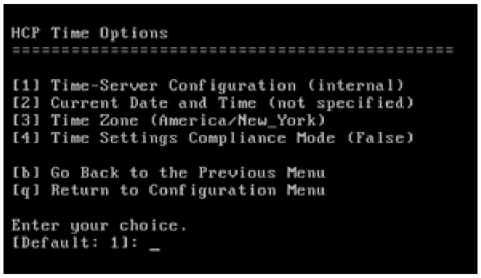

```
HCP Time Options
==========================================

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

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

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

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

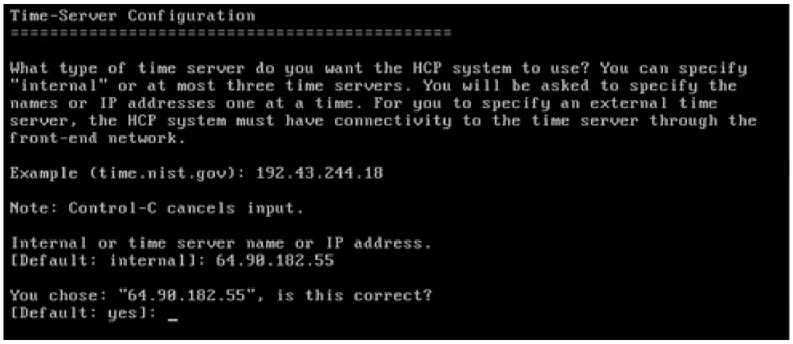

```
Time-Server Configuration
============================================

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

Example (time.nist.gov): 192.43.244.18

Note: Control-C cancels input.

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

You chose: "64.90.182.55", is this correct?
[Default: yes]: _
```

17. Enter *y* or *yes* to confirm the change and return to the **New Install** menu.

18. From the **New Install** menu, enter **6** to change the internal configuration settings.

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

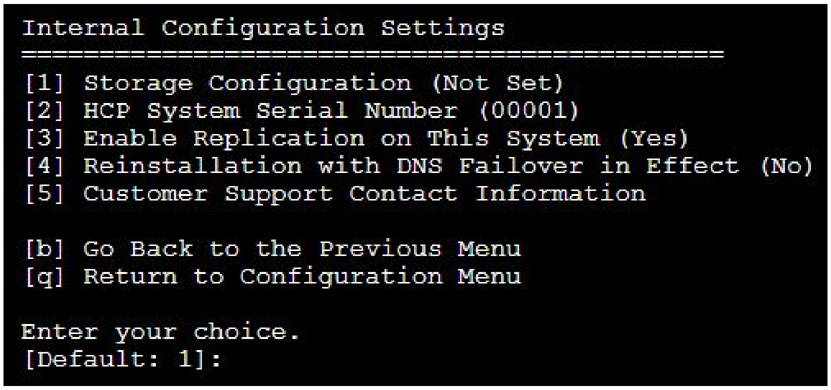

```
Internal Configuration Settings
=============================================
[1] Storage Configuration (Not Set)
[2] HCP System Serial Number (00001)
[3] Enable Replication on This System (Yes)
[4] Reinstallation with DNS Failover in Effect (No)
[5] Customer Support Contact Information

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

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

19. From the **Internal Configuration Settings** menu, enter **1** to set the storage configuration.

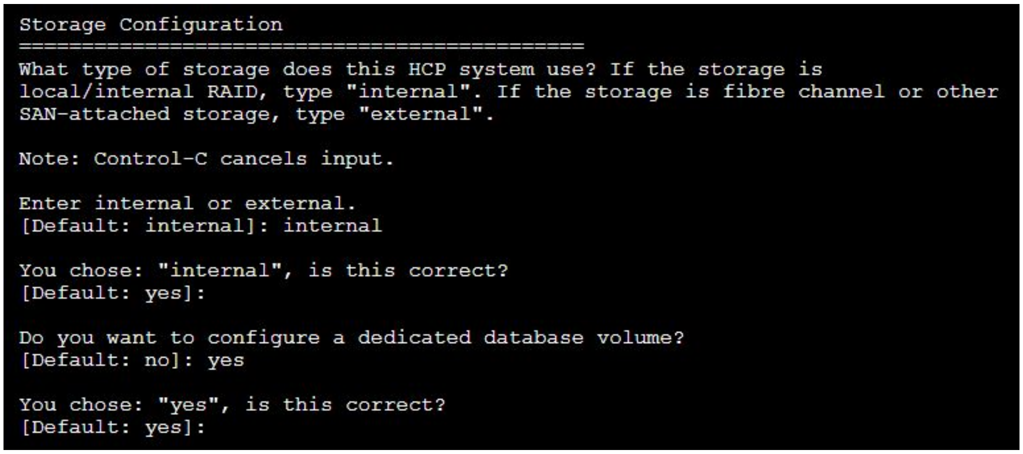

```
Storage Configuration
=============================================
What type of storage does this HCP system use? If the storage is
local/internal RAID, type "internal". If the storage is fibre channel or other
SAN-attached storage, type "external".

Note: Control-C cancels input.

Enter internal or external.
[Default: internal]: internal

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

Do you want to configure a dedicated database volume?
[Default: no]: yes

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

20. Enter *internal*.

21. Press Enter in response to the confirming prompt.

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

22. If HCP Setup asks whether you want to configure a dedicated database volume, enter *yes* if you want to configure a dedicated database volume or *no* if you do not want to configure a dedicated database volume.

23. Press Enter to confirm your choices and return to the **Internal Configuration Settings** menu.

24. From the **Internal Configuration Settings** menu, enter **2** to set the serial number for the HCP system.

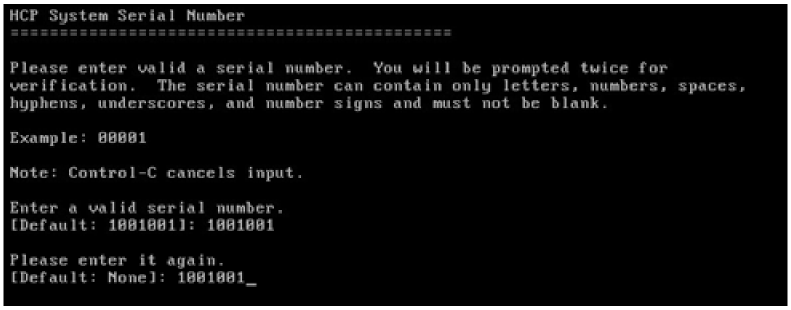

```
HCP System Serial Number
=============================================

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

Example: 00001

Note: Control-C cancels input.

Enter a valid serial number.
[Default: 1001001]: 1001001

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

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

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

Important: The HCP system serial number is required to license the system. Omitting the serial number will cause the system to report that you are in violation of your license agreement.

27. From the **Internal Configuration Settings** menu, enter **3** to configure whether replication will be enabled.

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

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

28. From the **Internal Configuration Settings** menu, enter **4** to configure whether reinstallation with DNS failover will be enabled.

29. From the **Internal Configuration Settings** menu, enter **5** to set contact information. To specify no contact information, hit **space**.

30. Enter **b** to return to the **New Install** menu.

Step 3: Execute the installation

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

To execute the HCP software installation:

1. From the **New Install** menu, enter **x**.

If you have installed as the install user, the wizard informs you that data-in-flight encryption is enabled, then asks if you are sure that it is legal to ship a system with data-in-flight encryption enabled to the country in which you are deploying the system.

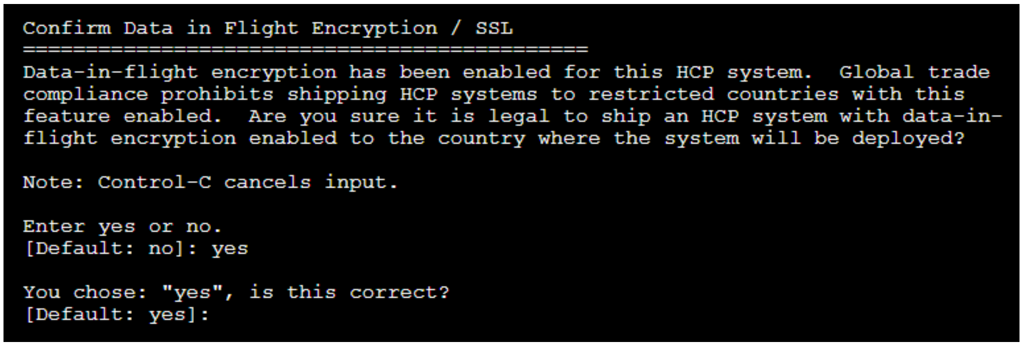

```
Confirm Data in Flight Encryption / SSL
=============================================
Data-in-flight encryption has been enabled for this HCP system.  Global trade
compliance prohibits shipping HCP systems to restricted countries with this
feature enabled.  Are you sure it is legal to ship an HCP 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.

After you press Enter, the wizard displays the configuration confirmation.

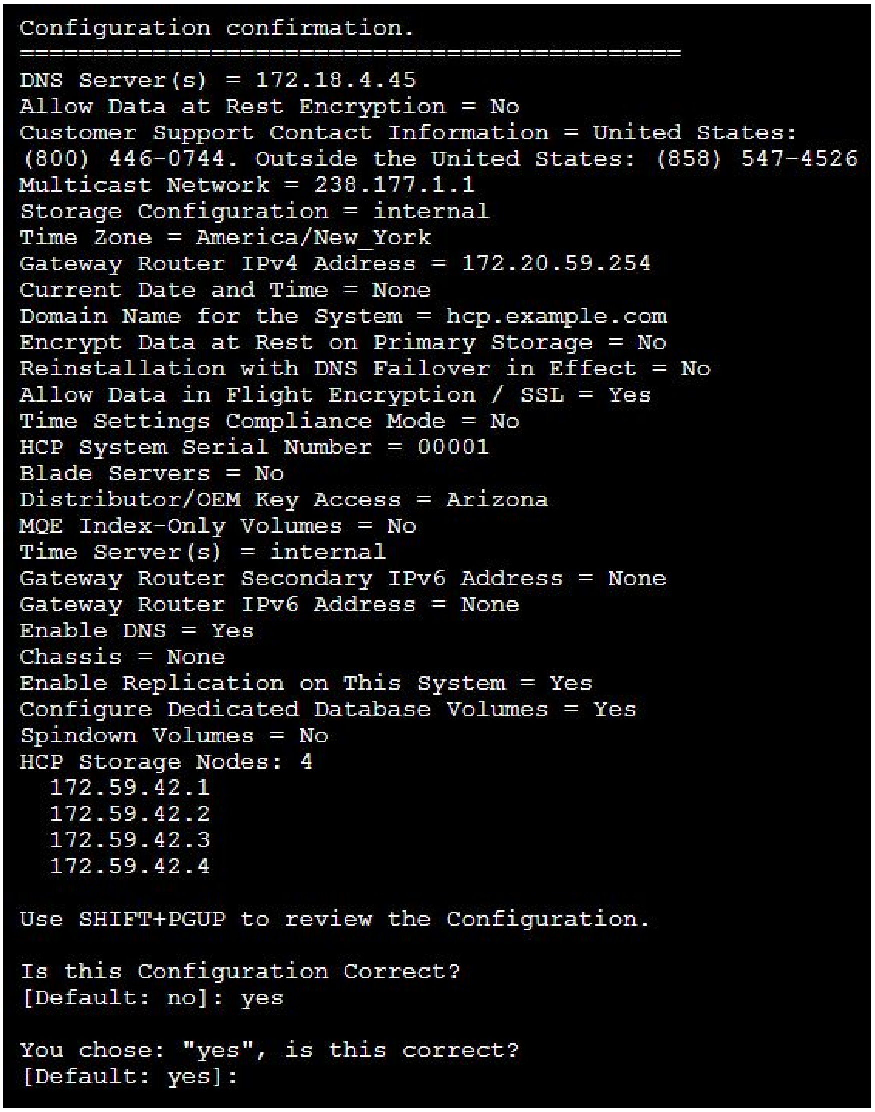

```
Configuration confirmation.
=============================================
DNS Server(s) = 172.18.4.45
Allow Data at Rest Encryption = No
Customer Support Contact Information = United States:
(800) 446-0744. Outside the United States: (858) 547-4526
Multicast Network = 238.177.1.1
Storage Configuration = internal
Time Zone = America/New_York
Gateway Router IPv4 Address = 172.20.59.254
Current Date and Time = None
Domain Name for the System = hcp.example.com
Encrypt Data at Rest on Primary Storage = No
Reinstallation with DNS Failover in Effect = No
Allow Data in Flight Encryption / SSL = Yes
Time Settings Compliance Mode = No
HCP System Serial Number = 00001
Blade Servers = No
Distributor/OEM Key Access = Arizona
MQE Index-Only Volumes = No
Time Server(s) = internal
Gateway Router Secondary IPv6 Address = None
Gateway Router IPv6 Address = None
Enable DNS = Yes
Chassis = None
Enable Replication on This System = Yes
Configure Dedicated Database Volumes = Yes
Spindown Volumes = No
HCP Storage Nodes: 4
  172.59.42.1
  172.59.42.2
  172.59.42.3
  172.59.42.4

Use SHIFT+PGUP to review the Configuration.

Is this Configuration Correct?
[Default: no]: yes

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

4. Review the configuration.

5. Perform one of the following:

   - If the configuration is not correct:

     1. Enter *n* or *no*.

     2. In response to the confirmation prompt, enter *y* or *yes*.

     3. Correct the configuration information.

- If the configuration is correct:

    1. Enter *y* or *yes*.

    2. In response to the confirmation prompt, enter *y* or *yes*.

After you have confirmed that the configuration information is correct, HCP Setup performs a set of installation prechecks.

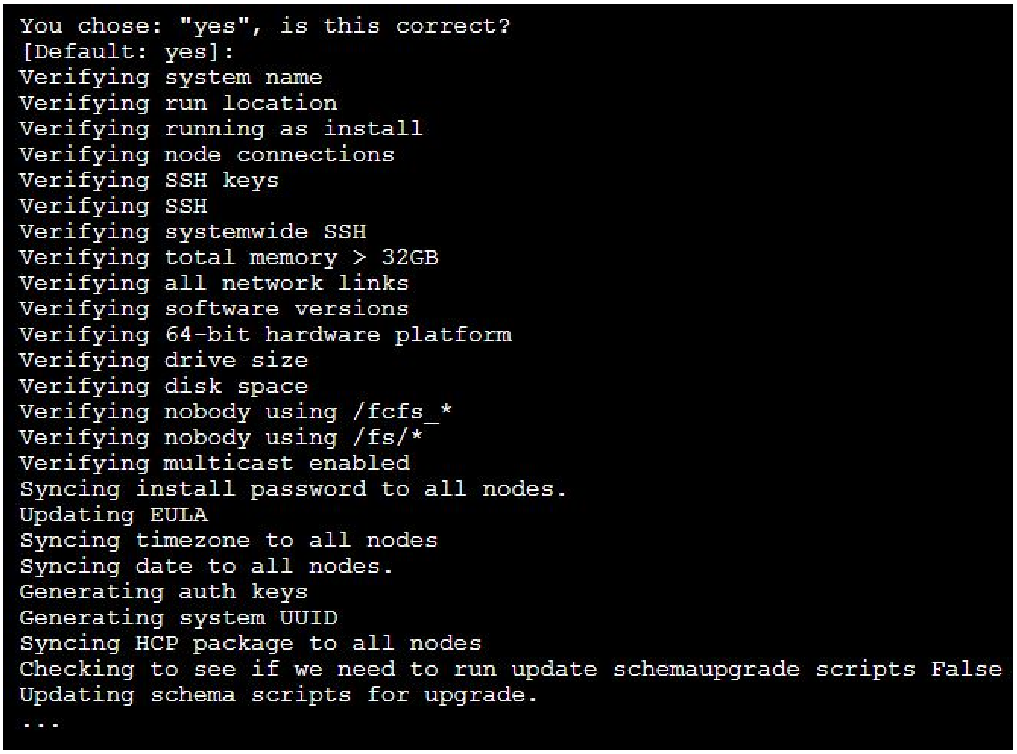

```
You chose: "yes", is this correct?
[Default: yes]:
Verifying system name
Verifying run location
Verifying running as install
Verifying node connections
Verifying SSH keys
Verifying SSH
Verifying systemwide SSH
Verifying total memory > 32GB
Verifying all network links
Verifying software versions
Verifying 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.
...
```

Only if you have previously selected that you want to configure dedicated database volumes:

1. When prompted, select the dedicated database volume for the first node.

2. Press Enter to confirm your selection.

3. Repeat the above two steps for each node in the system.

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

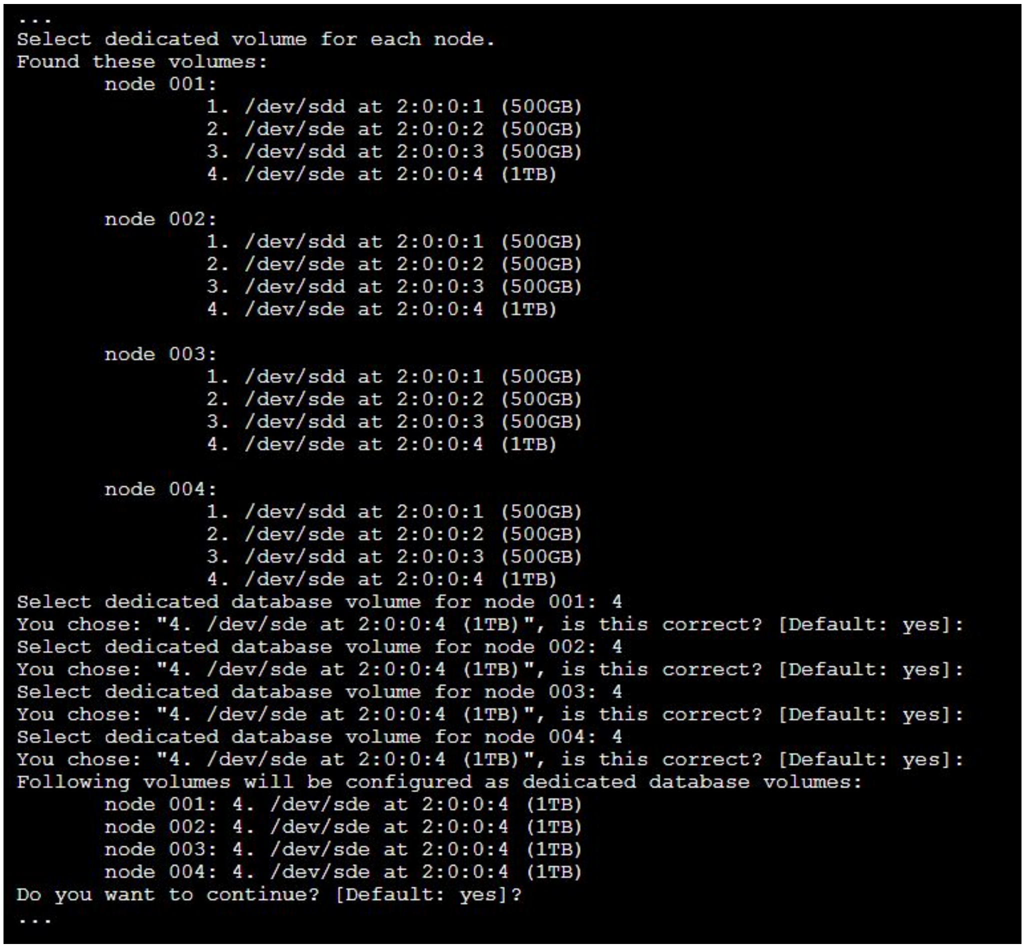

```
...
Select dedicated volume for each node.
Found these volumes:
       node 001:
               1. /dev/sdd at 2:0:0:1 (500GB)
               2. /dev/sde at 2:0:0:2 (500GB)
               3. /dev/sdd at 2:0:0:3 (500GB)
               4. /dev/sde at 2:0:0:4 (1TB)

       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]?
...
```

4. Press Enter to continue the installation.

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.

Important: If you enabled encryption in the system configuration, HCP Setup displays the encryption key after doing some initial setup. It then prompts you to enter the key. Before entering the encryption key, write it down on paper.

After initial set up, HCP displays and asks you to enter the encryption key. Once you enter it, HCP Setup completes the installation. You do not get a second chance to see the encryption key, and the key is not stored for later retrieval.

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

If HCP Setup exits at any time before the installation processing is complete, make a note of all error messages and then contact your authorized HCP service provider for help.

After the installation is complete, the HCP-VM nodes reboot, and, instead of the Operating System login prompt you should see an **hcp-node-<nodeNumber>** prompt.

After the reboot, you can also check the runlevel of the node by pressing Alt+F5 when inside the console.

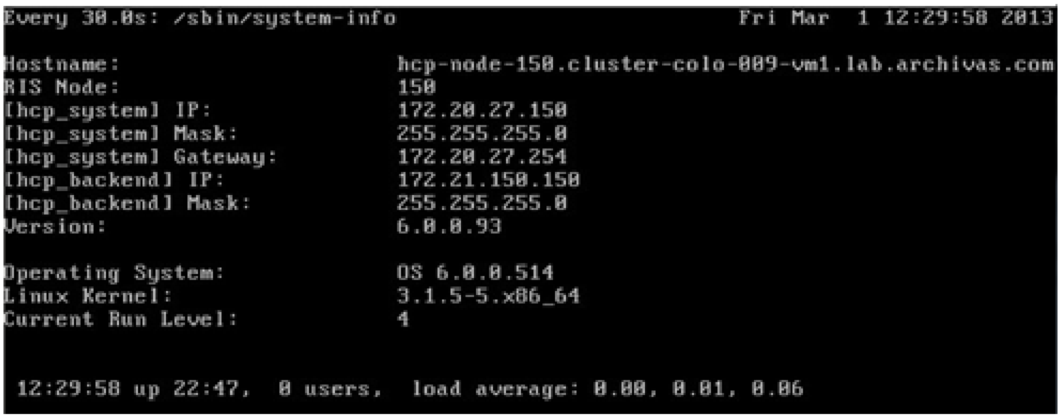

```
Every 30.0s: /sbin/system-info                                Fri Mar  1 12:29:58 2013

Hostname:                         hcp-node-150.cluster-colo-009-vm1.lab.archivas.com
RIS Node:                         150
[hcp_system] IP:                  172.20.27.150
[hcp_system] Mask:                255.255.255.0
[hcp_system] Gateway:             172.20.27.254
[hcp_backend] IP:                 172.21.150.150
[hcp_backend] Mask:               255.255.255.0
Version:                          6.0.0.93

Operating System:                 OS 6.0.0.514
Linux Kernel:                     3.1.5-5.x86_64
Current Run Level:                4


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

Step 4: Verify the HCP software installation

Access the HCP System Management Console to verify that the HCP software installed correctly.

To verify the HCP system installation:

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

   - 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`

`node-ip-address` is the Front-end IP address of any storage node in the HCP system.

Note: If you enter *http* instead of *https* in the URL, the browser returns an error. Enter the URL again, this time using *https*.

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 doesn't open, contact your authorized HCP service provider for help.

3. Check the serial number on the login page. If the serial number is incorrect, contact your authorized HCP service provider for help.

4. Log into the System Management Console with the following username and password:

   - Username: *security*
   - Password: *Chang3Me!*

   Once you login, the Console displays either the **Change Password** page or the **Hardware** page.

   If the Console displays the **Hardware** page, it means the nodes are still starting HCP. This process can take several minutes. When more than half the nodes have completed their startup process, the Console automatically displays the **Change Password** page.

   If the **Hardware** page remains displayed after several minutes, please contact your authorized HCP service provider for help.

5. On the **Change Password** page:

   a. In the **Existing Password** field, enter *Chang3Me!*.

b. In the **New Password** field, enter a new password.

c. In the **Confirm New Password** field, type your new password again.

d. Click on **Update Password**.

A valid password must contain any UTF-8 characters, including white space. The minimum length is six characters. The maximum is 64 characters. A password must include at least one character from two of these three groups: alphabetic, numeric, and other. For example:

- Valid password: *P@sswOrd*
- Invalid password: *password*

6. In the top-level menu, click on **Hardware**.

7. On the **Hardware** page, make sure the nodes have the:

- Node status is **Available**.
- Status of each logical volume is **Available**.

Tip: To see the status of a logical volume, hover over the volume icon.

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

If any nodes have a status other than **Available**, or if any logical volumes for available nodes have a status other than **Available** or **Spun down**, please contact your authorized HCP service provider for help. Also contact your service provider if the number of logical volume icons for each node does not match the expected number of logical volumes for the node.

8. Do either of the following steps:

a. Set additional configuration options, as described in "Setting additional configuration options" on the next page. You can set additional configuration options only if the installation was successful.

b. Log out of the System Management Console and close the browser window to ensure that no one can return to the Console without logging in.

Setting additional configuration options

After verifying that the HCP system was correctly installed, you can set additional configuration options. For example, you can enable the management port network, enable syslog logging, or disable ping.

To set additional configuration options:

1. Log into the HCP System Management Console as the security user (if you are not logged in already).

2. Create a new user account with the administrator role.

   Alternatively, you can add the administrator role to the security user account and then skip step 3 below.

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.

   For information on creating user accounts and performing system configuration activities, see *Administering HCP*.

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 needs to be monitored outside of the virtual machine environment for failures.

Hitachi servers and network components that are part of the HCP-VM system can be connected to HiTrack for monitoring. For more information on HiTrack, see Chapter 6: "Configuring HCP monitoring with Hi-Track Monitor" on page 79

Non-Hitachi equipment should be monitored using the vendor or customer equivalent of Hi-Track.

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

HCP IPMI monitoring and Hitachi array monitoring is not available for HCP-VMs.

Software monitoring

HCP maintains a system log which logs all events that happen within the system. 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/or email addresses. Additionally, you can use SNMP to view and, when allowed, change HCP system settings.

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

The HCP Software application's health can be monitored with HiTrack. For more information on HiTrack, see Chapter 6: "Configuring HCP monitoring with Hi-Track Monitor" on page 79

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.

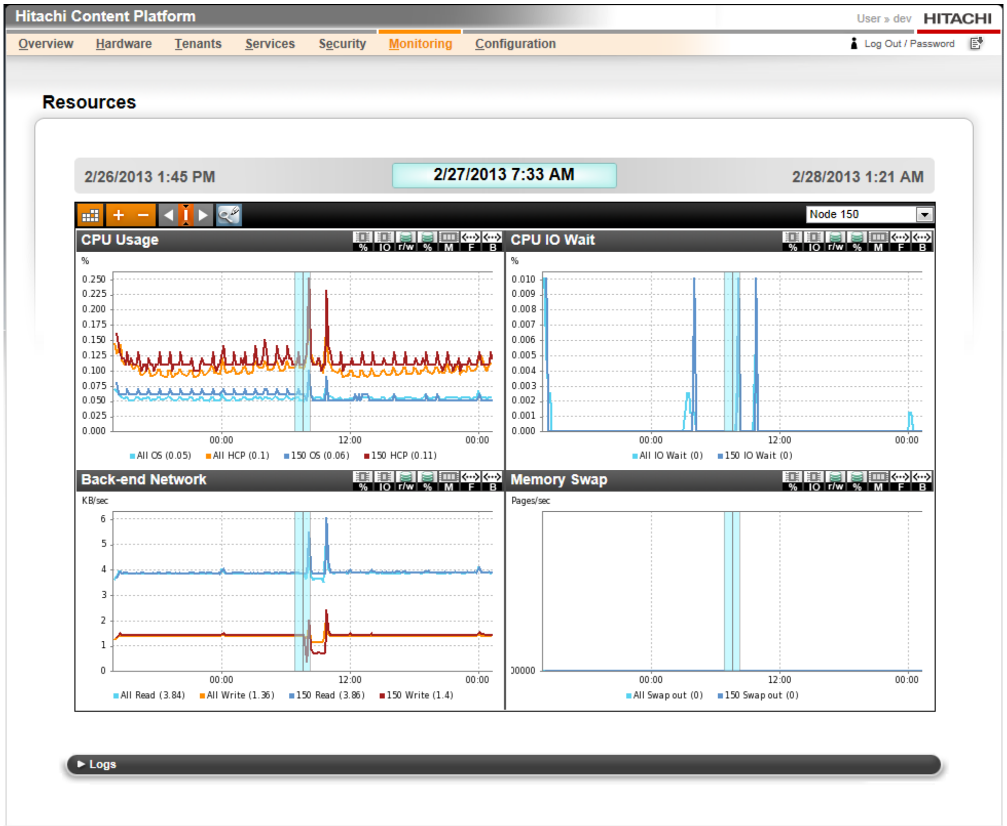

HCP-VM diagnostic menu

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. The diagnostics are available through the system 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 facilities, see the *HCP System Management* help.

VMware Monitoring and Performance is the responsibility of the customer. In the vSphere center, under the performance tab, clients have multiple ways to monitor resources.

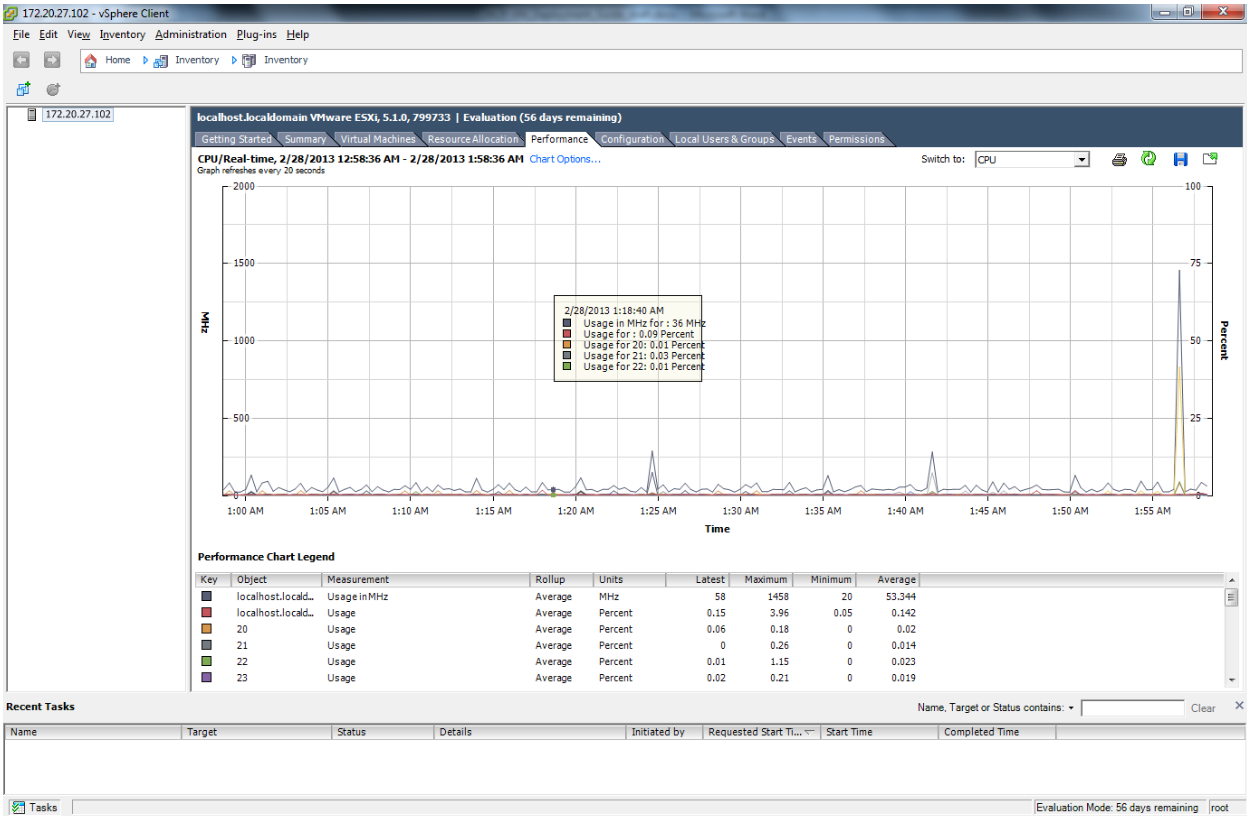

localhost.localdomain VMware ESXi, 5.1.0, 799733 | Evaluation (56 days remaining)

CPU/Real-time, 2/28/2013 12:58:36 AM - 2/28/2013 1:58:36 AM

Graph refreshes every 20 seconds

Switch to: CPU

2/28/2013 1:18:40 AM
■ Usage in MHz for : 36 MHz
■ Usage for : 0.09 Percent
■ Usage for 20: 0.01 Percent
■ Usage for 21: 0.03 Percent
■ Usage for 22: 0.01 Percent

Performance Chart Legend

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	localhost.locald...	Usage in MHz	Average	MHz	58	1458	20	53.344
🟥	localhost.locald...	Usage	Average	Percent	0.15	3.96	0.05	0.142
🟧	20	Usage	Average	Percent	0.06	0.18	0	0.02
■	21	Usage	Average	Percent	0	0.26	0	0.014
🟩	22	Usage	Average	Percent	0.01	1.15	0	0.023
🟪	23	Usage	Average	Percent	0.02	0.21	0	0.019

Name	Target	Status	Details	Initiated by	Requested Start Ti...	Start Time	Completed Time

Evaluation Mode: 56 days remaining

For more details on monitoring options, refer to the VMware Monitoring and Performance guide which can be found here: http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-51-monitoring-performance-guide.pdf

5

Maintenance procedures

This chapter describes and provides procedural steps for keeping your HCP-VM system running at an optimal performance level.

Adding logical volumes

To add logical volumes, follow these steps:

1. As described in "Provisioning HCP-VM storage" on page 19, provision the LUNs to be added to each ESXi host in the system.

2. As described in "Adding VMFS datastores to a vSphere HA cluster" on page 22, add the LUNs to new datastores.

Important: This must be one LUN per datastore.

3. From the vSphere client, right-click on the HCP-VM to which capacity should be added and select **Edit Settings**.

4. In the Virtual Machine Properties window that opens, click on **Add**.

5. In the **Add Hardware** window, select **Hard Disk**.

6. Click on **Next**.

7. Select **Create a new virtual disk**.

8. Click on **Next**.

9. Set the capacity to be slightly smaller than the size of the LUN that was provisioned (VMware adds a small amount of overhead).

In the following example, the size of the LUN provisioned was 1.2 TB.

10. Select **Thick Provision Eager Zeroed**.

11. Browse to the new datastore that will be added then click on **Next**.

12. Select the next available SCSI disk in the Virtual Device node section then click on **Next**.

13. Verify the options selected then click on **Finish**.

14. Back in the **Virtual Machine Properties** window, verify that the new Hard Disk is listed then click on **OK**.

15. In the vSphere client, open the virtual machine console for the highest-numbered storage node.

16. Log in as the install user.

    The **HCP 8.1 Configuration Menu** appears.

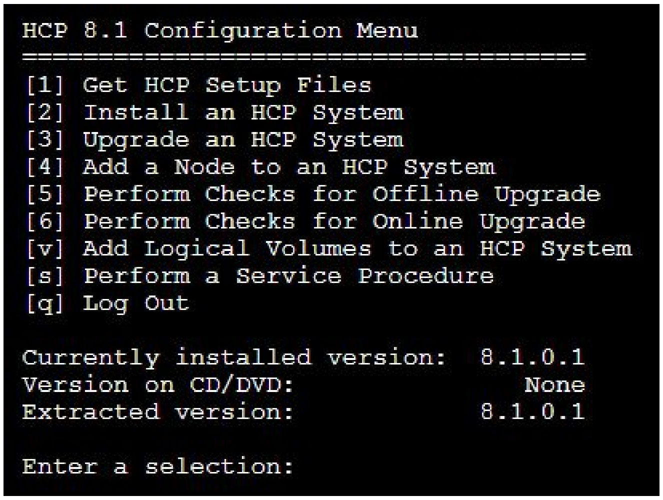

```
HCP 8.1 Configuration Menu
=====================================
[1] Get HCP Setup Files
[2] Install an HCP System
[3] Upgrade an HCP System
[4] Add a Node to an HCP System
[5] Perform Checks for Offline Upgrade
[6] Perform Checks for Online Upgrade
[v] Add Logical Volumes to an HCP System
[s] Perform a Service Procedure
[q] Log Out

Currently installed version:  8.1.0.1
Version on CD/DVD:               None
Extracted version:            8.1.0.1

Enter a selection:
```

17. Enter *v* to add logical volumes to an HCP system.

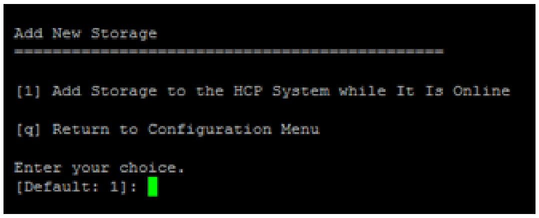

```
Add New Storage
=============================================

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

[q] Return to Configuration Menu

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

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

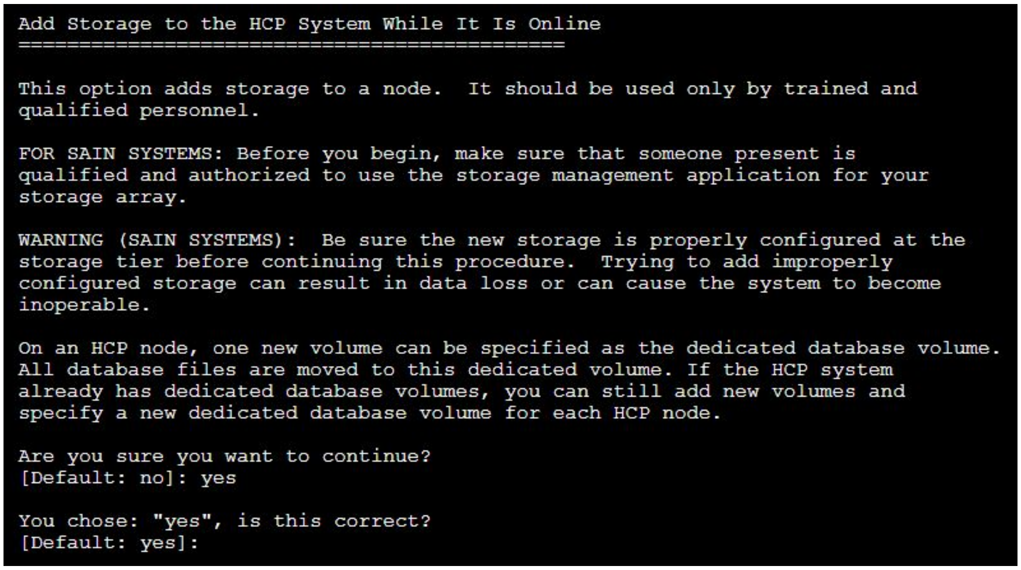

```
Add Storage to the HCP System While It Is Online
=============================================

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

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

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

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

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

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

19. Enter *yes* to continue the procedure.

20. Press Enter to confirm.

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

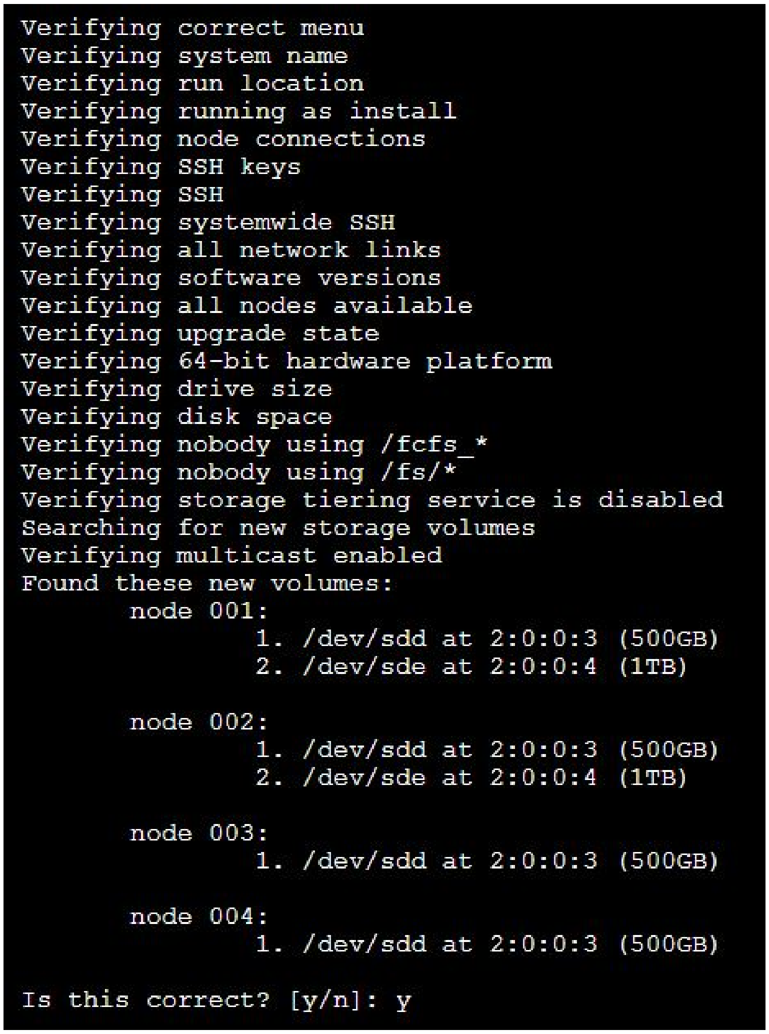

```
Verifying correct menu
Verifying system name
Verifying run location
Verifying running as install
Verifying node connections
Verifying SSH keys
Verifying SSH
Verifying systemwide SSH
Verifying all network links
Verifying software versions
Verifying all nodes available
Verifying upgrade state
Verifying 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
```

21. Enter *y* to verify the new volumes that were found. Typically this would show storage added to all nodes.

    Optionally, if you want to configure a dedicated database volume, the volume size needs to be at least 50 GB and at least 1.5 times the size of the existing database for each node. All dedicated database volumes need to be the same size. If you already have a dedicated database volume, any newly-added dedicated database volume needs to be larger than the current one.

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

    ○ If you do not want to select a dedicated volume for the database, enter *no*. HCP Setup formats and adds the new volumes. During this process, HCP Setup reports on its progress.

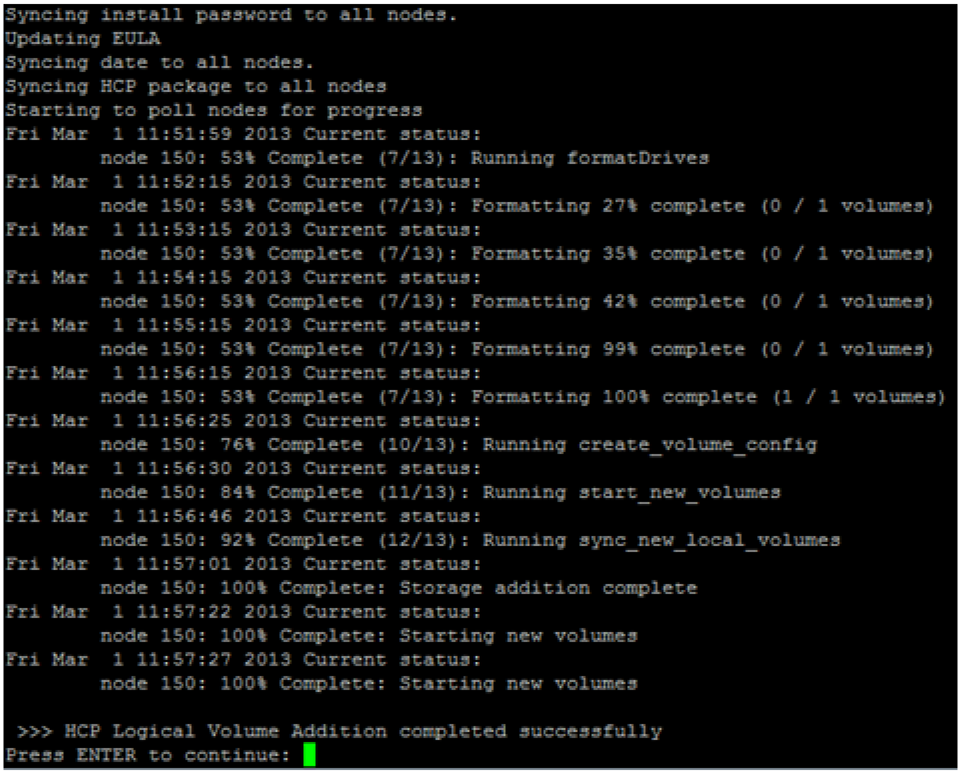

```
Syncing install password to all nodes.
Updating EULA
Syncing date to all nodes.
Syncing HCP package to all nodes
Starting to poll nodes for progress
Fri Mar  1 11:51:59 2013 Current status:
        node 150: 53% Complete (7/13): Running formatDrives
Fri Mar  1 11:52:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 27% complete (0 / 1 volumes)
Fri Mar  1 11:53:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 35% complete (0 / 1 volumes)
Fri Mar  1 11:54:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 42% complete (0 / 1 volumes)
Fri Mar  1 11:55:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 99% complete (0 / 1 volumes)
Fri Mar  1 11:56:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 100% complete (1 / 1 volumes)
Fri Mar  1 11:56:25 2013 Current status:
        node 150: 76% Complete (10/13): Running create_volume_config
Fri Mar  1 11:56:30 2013 Current status:
        node 150: 84% Complete (11/13): Running start_new_volumes
Fri Mar  1 11:56:46 2013 Current status:
        node 150: 92% Complete (12/13): Running sync_new_local_volumes
Fri Mar  1 11:57:01 2013 Current status:
        node 150: 100% Complete: Storage addition complete
Fri Mar  1 11:57:22 2013 Current status:
        node 150: 100% Complete: Starting new volumes
Fri Mar  1 11:57:27 2013 Current status:
        node 150: 100% Complete: Starting new volumes

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

- If you want to select a dedicated volume for the database, enter *yes*, then:

  1. If you want to select a dedicated database volume for the first node, enter *yes*.

  2. If you entered *yes*, select the dedicated database volume for the first node.

  3. Press Enter to confirm your selection.

  4. Repeat the above three steps for each node in the system.

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

  5. Enter *yes* to continue the procedure.

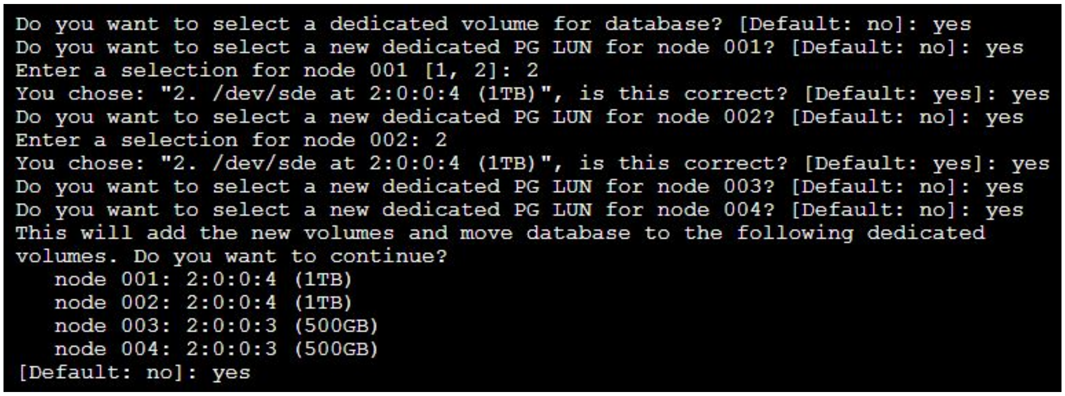

```
Do you want to select a dedicated volume for database? [Default: no]: yes
Do you want to select a new dedicated PG LUN for node 001? [Default: no]: yes
Enter a selection for node 001 [1, 2]: 2
You chose: "2. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default: yes]: yes
Do you want to select a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

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

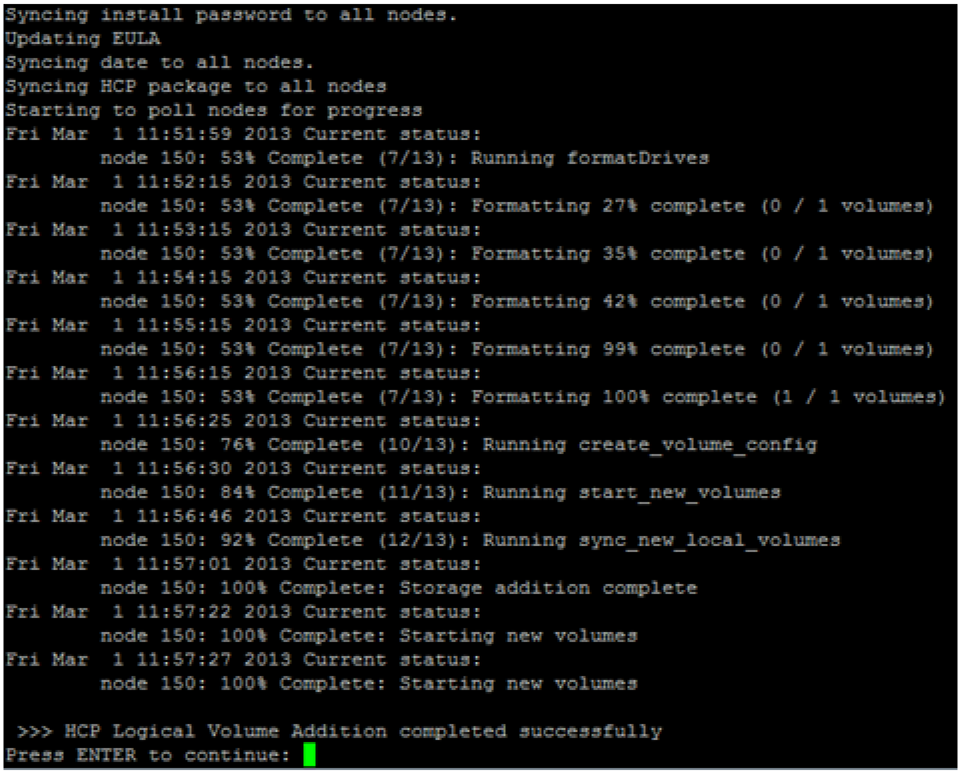

```
Syncing install password to all nodes.
Updating EULA
Syncing date to all nodes.
Syncing HCP package to all nodes
Starting to poll nodes for progress
Fri Mar  1 11:51:59 2013 Current status:
        node 150: 53% Complete (7/13): Running formatDrives
Fri Mar  1 11:52:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 27% complete (0 / 1 volumes)
Fri Mar  1 11:53:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 35% complete (0 / 1 volumes)
Fri Mar  1 11:54:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 42% complete (0 / 1 volumes)
Fri Mar  1 11:55:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 99% complete (0 / 1 volumes)
Fri Mar  1 11:56:15 2013 Current status:
        node 150: 53% Complete (7/13): Formatting 100% complete (1 / 1 volumes)
Fri Mar  1 11:56:25 2013 Current status:
        node 150: 76% Complete (10/13): Running create_volume_config
Fri Mar  1 11:56:30 2013 Current status:
        node 150: 84% Complete (11/13): Running start_new_volumes
Fri Mar  1 11:56:46 2013 Current status:
        node 150: 92% Complete (12/13): Running sync_new_local_volumes
Fri Mar  1 11:57:01 2013 Current status:
        node 150: 100% Complete: Storage addition complete
Fri Mar  1 11:57:22 2013 Current status:
        node 150: 100% Complete: Starting new volumes
Fri Mar  1 11:57:27 2013 Current status:
        node 150: 100% Complete: Starting new volumes

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

23. When the formatting is complete, press Enter to continue.

24. Log in to the HCP System Management Console to verify the newly added volumes.

Moving storage node databases to optimal volumes

To move the HCP database to optimal volumes:

1. From the **HCP 8.1 Configuration** menu, enter **s** to display the **HCP Service** menu.

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

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

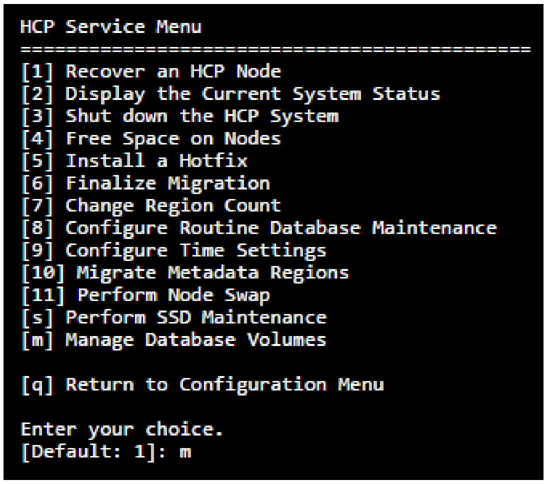

```
HCP Service Menu
=============================================
[1] Recover an HCP Node
[2] Display the Current System Status
[3] Shut down the HCP System
[4] Free Space on Nodes
[5] Install a Hotfix
[6] Finalize Migration
[7] Change Region Count
[8] Configure Routine Database Maintenance
[9] Configure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

[q] Return to Configuration Menu

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

3. From the **HCP Service** menu, enter **m**.

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

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

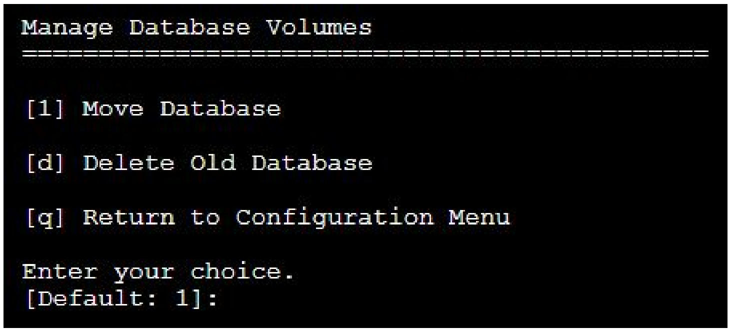

```
Manage Database Volumes
=============================================

 [1] Move Database

 [d] Delete Old Database

 [q] Return to Configuration Menu

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

5. From the **Manage Database Volumes** menu, enter **1**.

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

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

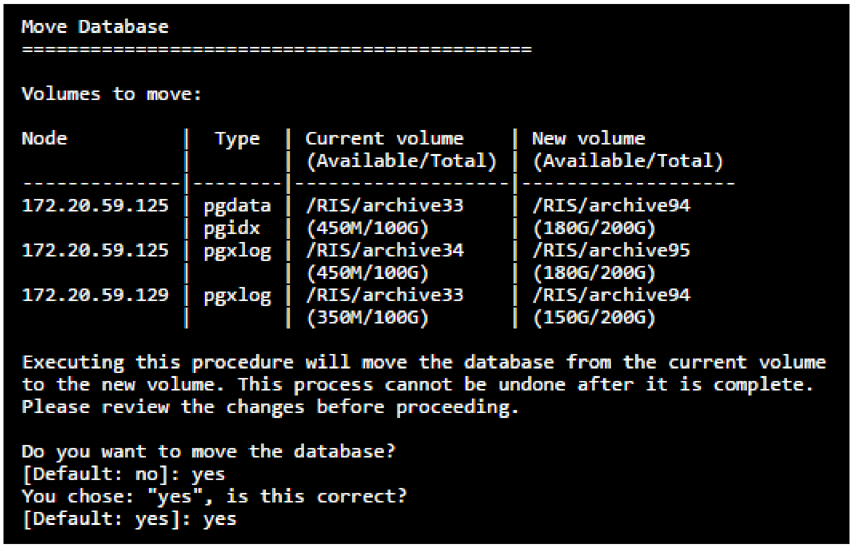

```
Move Database
============================================

Volumes to move:

Node           |  Type   | Current volume     | New volume
               |         | (Available/Total)  | (Available/Total)
---------------|---------|--------------------|-------------------
172.20.59.125  | pgdata  | /RIS/archive33     | /RIS/archive94
               | pgidx   | (450M/100G)        | (180G/200G)
172.20.59.125  | pgxlog  | /RIS/archive34     | /RIS/archive95
               |         | (450M/100G)        | (180G/200G)
172.20.59.129  | pgxlog  | /RIS/archive33     | /RIS/archive94
               |         | (350M/100G)        | (150G/200G)

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

Do you want to move the database?
[Default: no]: yes
You chose: "yes", is this correct?
[Default: yes]: yes
```

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

7. When you have reviewed the configuration, enter *y* or *yes* to confirm the move or *n* or *no* to try again.

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

   Once the procedure is initiated, the progress of the database move appears and details the current status of the HCP Service Procedure.

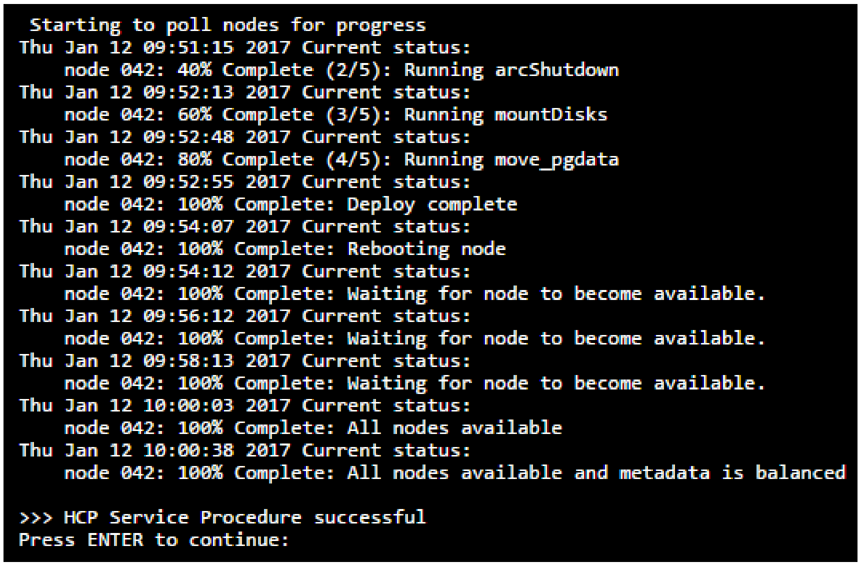

```
 Starting to poll nodes for progress
Thu Jan 12 09:51:15 2017 Current status:
    node 042: 40% Complete (2/5): Running arcShutdown
Thu Jan 12 09:52:13 2017 Current status:
    node 042: 60% Complete (3/5): Running mountDisks
Thu Jan 12 09:52:48 2017 Current status:
    node 042: 80% Complete (4/5): Running move_pgdata
Thu Jan 12 09:52:55 2017 Current status:
    node 042: 100% Complete: Deploy complete
Thu Jan 12 09:54:07 2017 Current status:
    node 042: 100% Complete: Rebooting node
Thu Jan 12 09:54:12 2017 Current status:
    node 042: 100% Complete: Waiting for node to become available.
Thu Jan 12 09:56:12 2017 Current status:
    node 042: 100% Complete: Waiting for node to become available.
Thu Jan 12 09:58:13 2017 Current status:
    node 042: 100% Complete: Waiting for node to become available.
Thu Jan 12 10:00:03 2017 Current status:
    node 042: 100% Complete: All nodes available
Thu Jan 12 10:00:38 2017 Current status:
    node 042: 100% Complete: All nodes available and metadata is balanced

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

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

Deleting databases from older database volumes

To delete a database from an older database volume:

1. From the **HCP 8.1 Configuration** menu, enter **s** to display the **HCP Service** menu.

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

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

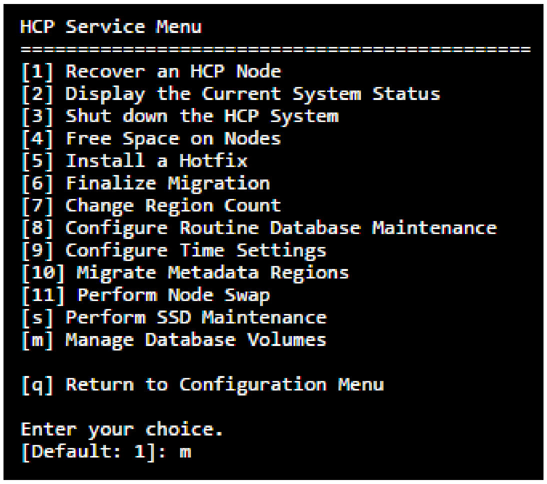

```
HCP Service Menu
=============================================
[1] Recover an HCP Node
[2] Display the Current System Status
[3] Shut down the HCP System
[4] Free Space on Nodes
[5] Install a Hotfix
[6] Finalize Migration
[7] Change Region Count
[8] Configure Routine Database Maintenance
[9] Configure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

[q] Return to Configuration Menu

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

3. From the **HCP Service** menu, enter **m**.

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

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

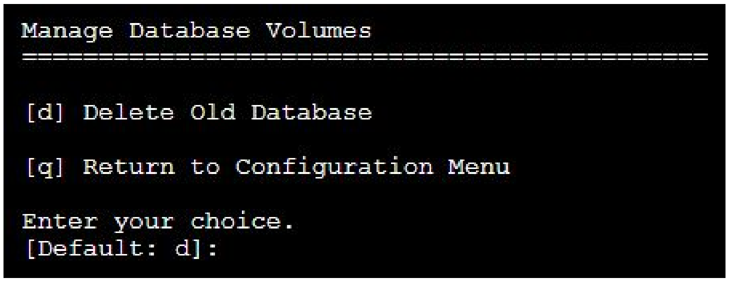

```
Manage Database Volumes
============================================

[d] Delete Old Database

[q] Return to Configuration Menu

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

5. From the **Manage Database Volumes** menu, enter **d**. You can delete the database from the older database volume only if you have completed the database move procedure.

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

   When you enter *y* or *yes*, HCP Setup displays the **Delete Old Database** menu.

   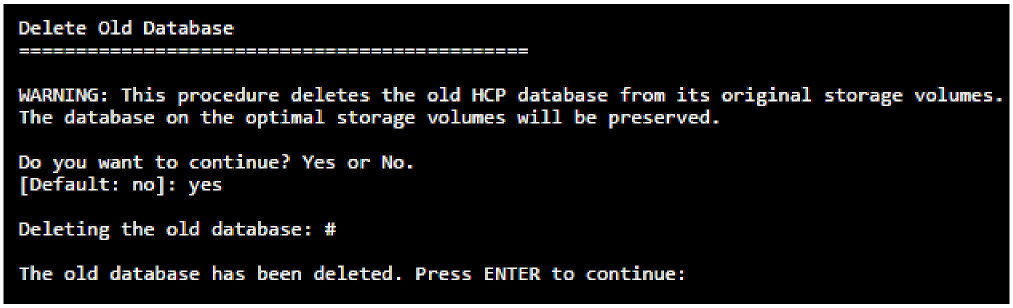

   ```
   Delete Old Database
   =============================================

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

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

   Deleting the old database: #

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

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

   Once the procedure is complete, press Enter to return to the **HCP Service** menu.

Adding HCP-VM nodes

The process for adding HCP-VM nodes is:

1. Add new ESXi hosts or find existing ESXi hosts that can support an HCP node. For more information about creating ESXi hosts, see Chapter 3: "Configuring the HCP-VM environment" on page 15.

2. Unpacking and uploading the ISO files to the selected ESXi hosts. For more information about this process, see "Unpacking and uploading the ISO zip file" on page 29.

3. Creating the new virtual machine. For more information about this process, see "Creating the new virtual machine" on page 29.

4. Configuring the HCP-VM network on the newly deployed HCP-VM nodes. For more information about configuring network information, see "Installing the Appliance Operating System" on page 33.

5. From the highest active HCP-VM node, run the add node service procedure. For more information on this and other procedures, refer to the *Installing and Maintaining an HCP System* manual.

   To add nodes:

   1. From the **HCP 8.1 Configuration** menu, enter **4** to run the HCP Setup wizard.

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

      When you enter *y* or *yes*, the **Membership Update** menu appears.

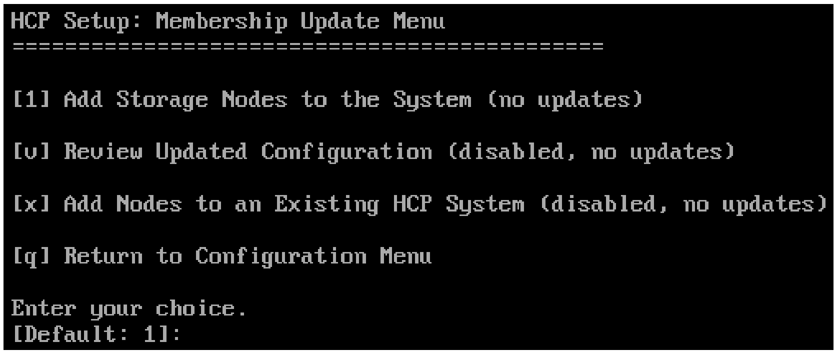

```
HCP Setup: Membership Update Menu
=============================================

[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.

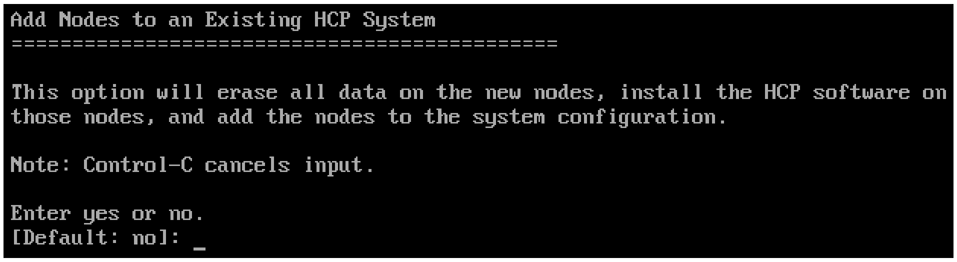

```
Add Nodes to an Existing HCP System
============================================

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 confirming prompt, enter *y* or *yes* to confirm that you want to perform the procedure or *n* or *no* to back out.

      The wizard prompts again for confirmation.

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

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

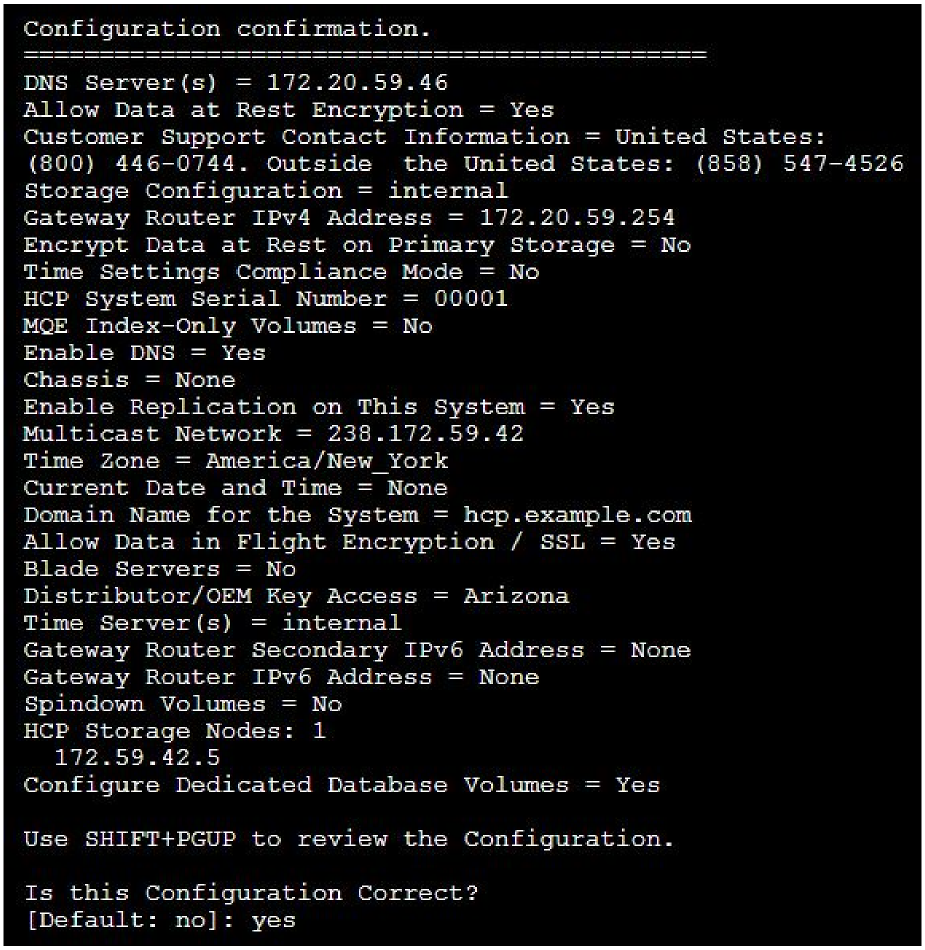

```
Configuration confirmation.
=============================================
DNS Server(s) = 172.20.59.46
Allow Data at Rest Encryption = Yes
Customer Support Contact Information = United States:
(800) 446-0744. Outside  the United States: (858) 547-4526
Storage Configuration = internal
Gateway Router IPv4 Address = 172.20.59.254
Encrypt Data at Rest on Primary Storage = No
Time Settings Compliance Mode = No
HCP System Serial Number = 00001
MQE Index-Only Volumes = No
Enable DNS = Yes
Chassis = None
Enable Replication on This System = Yes
Multicast Network = 238.172.59.42
Time Zone = America/New_York
Current Date and Time = None
Domain Name for the System = hcp.example.com
Allow Data in Flight Encryption / SSL = Yes
Blade Servers = No
Distributor/OEM Key Access = Arizona
Time Server(s) = internal
Gateway Router Secondary IPv6 Address = None
Gateway Router IPv6 Address = None
Spindown Volumes = No
HCP Storage Nodes: 1
  172.59.42.5
Configure Dedicated Database Volumes = Yes

Use SHIFT+PGUP to review the Configuration.

Is this Configuration Correct?
[Default: no]: yes
```

6. Review the configuration.

7. Take one of these actions:

   - If the configuration is correct:

     1. Enter *y* or *yes*.

     2. In response to the confirming prompt, enter *y* or *yes*.

     After you have confirmed that the configuration information is correct, HCP Setup performs a set of prechecks.

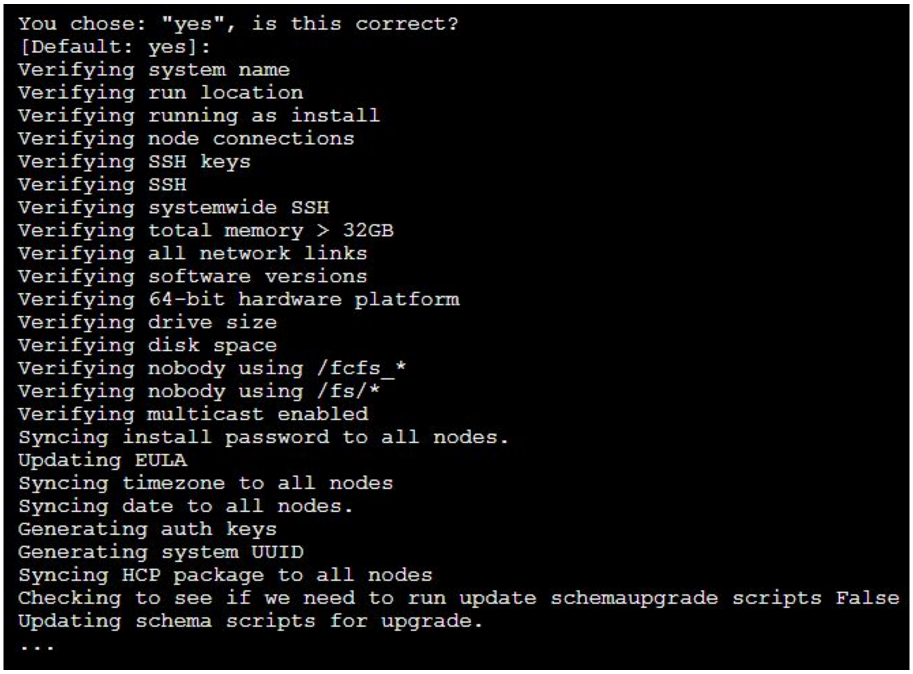

```
You chose: "yes", is this correct?
[Default: yes]:
Verifying system name
Verifying run location
Verifying running as install
Verifying node connections
Verifying SSH keys
Verifying SSH
Verifying systemwide SSH
Verifying total memory > 32GB
Verifying all network links
Verifying software versions
Verifying 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.
...
```

Only if your current HCP system has dedicated database volumes, each newly-added node needs to have at least three volumes. Also, the dedicated database volume size needs to be at least 50 GB. All dedicated database volumes need to be the same size. To select dedicated database volumes:

1. When prompted, select the dedicated database volume for the first node.

2. Press Enter to confirm your selection.

3. Repeat the above two steps for each node in the system.

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

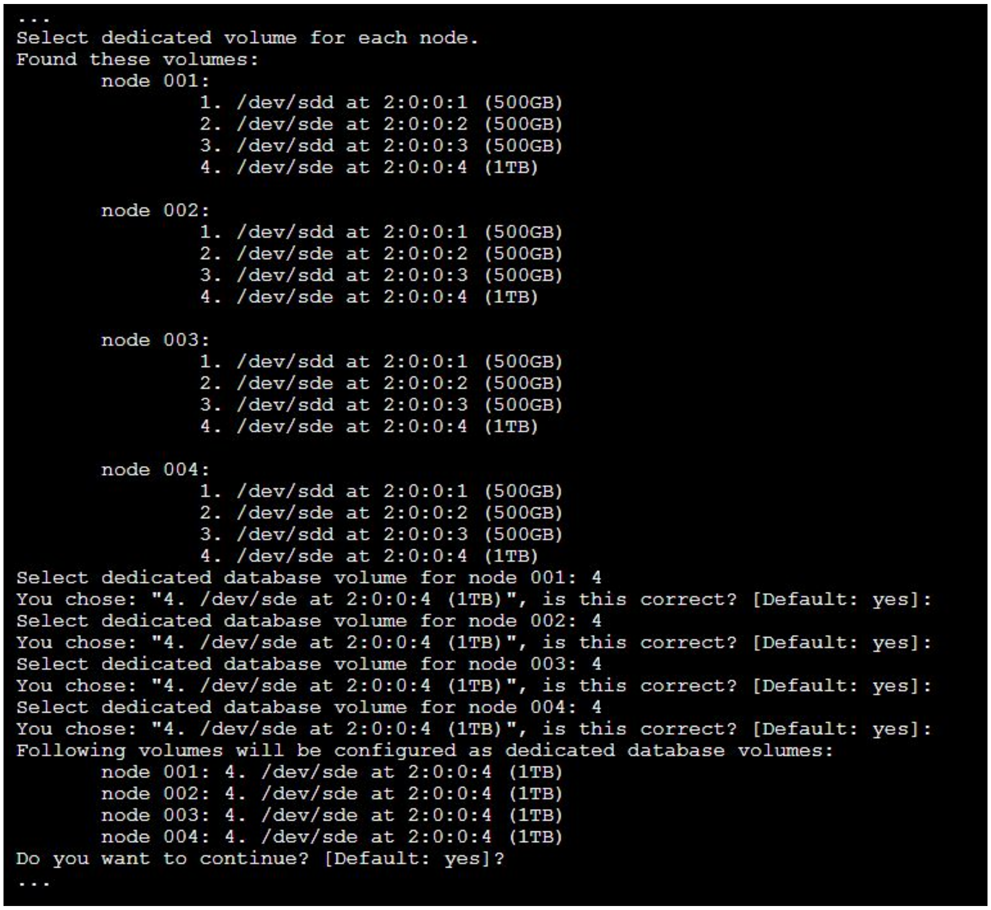

```
...
Select dedicated volume for each node.
Found these volumes:
       node 001:
               1. /dev/sdd at 2:0:0:1 (500GB)
               2. /dev/sde at 2:0:0:2 (500GB)
               3. /dev/sdd at 2:0:0:3 (500GB)
               4. /dev/sde at 2:0:0:4 (1TB)

       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]?
...
```

4. Press Enter to continue the procedure.

   If the prechecks are successful, the HCP software is installed on the new nodes, working on four nodes at a time for RAIN and VM systems and on one cross-mapped pair of nodes at a time for SAIN systems. After installing the software on a set of nodes, HCP Setup reboots those nodes.

   When the node addition is complete, the **HCP 8.1 Configuration** menu reappears.

   If any of the prechecks fail, HCP Setup exits. In this case, fix the problem and then start the node addition procedure again.

   If HCP Setup exits at any time before the node addition processing is complete, please contact your HCP support center for help.

– If the configuration is incorrect:

   1. Enter *n* or *no*.

   2. In response to the confirming prompt, enter *y* or *yes*.

   3. Exit the wizard and contact your HCP support center for help.

8. From the **HCP 8.1 Configuration** menu, enter **q** to log out of the install shell.

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

Recovering storage nodes

The following sections provide procedural steps on how to recover storage nodes when preserving storage volumes and how to recover storage nodes when clearing storage volumes. For more information on these and other procedures, refer to the *Installing and Maintaining an HCP System* manual.

Recovering storage nodes (preserving storage volumes)

To recover storage nodes when preserving storage volumes:

1. From the **HCP 8.1 Configuration** menu, enter **s** to display the **HCP Service** menu.

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

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

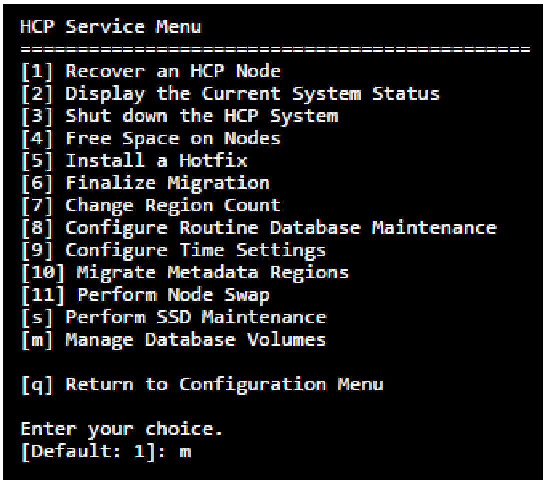

```
HCP Service Menu
=============================================
[1] Recover an HCP Node
[2] Display the Current System Status
[3] Shut down the HCP System
[4] Free Space on Nodes
[5] Install a Hotfix
[6] Finalize Migration
[7] Change Region Count
[8] Configure Routine Database Maintenance
[9] Configure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

[q] Return to Configuration Menu

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

3. From the **HCP Service** menu, enter **1** for recovery operations.

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

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

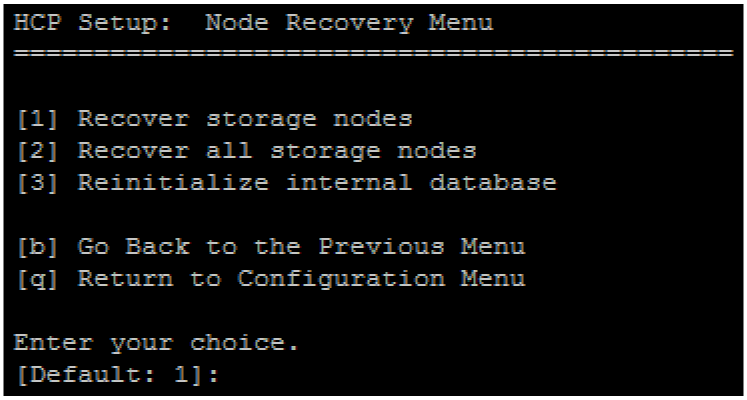

```
HCP Setup:  Node Recovery Menu
=============================================


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

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

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

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

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

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

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

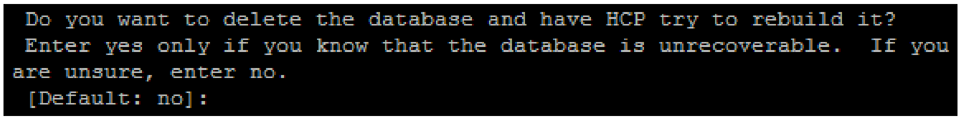

```
 Do you want to delete the database and have HCP try to rebuild it?
 Enter yes only if you know that the database is unrecoverable.  If you
are unsure, enter no.
 [Default: no]:
```

In response:

1. Enter *y* or *yes* to delete the database while recovering the OS or *n* or *no* to recover the OS without deleting the database.

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

3. Optionally, if you are performing the OS recovery on an HCP G10 Node with Attached Storage or on an HCP system with dedicated database volumes, HCP Setup displays the following prompt. If you receive this prompt, enter *y* or *yes* to format only the internal database drive or enter *n* or *no* to keep the internal database drive in its original state.

   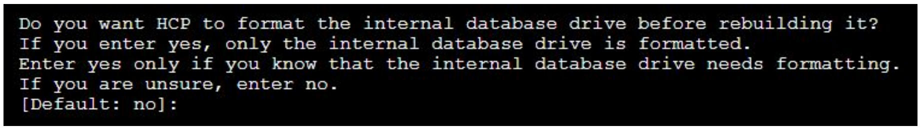

   ```
   Do you want HCP to format the internal database drive before rebuilding it?
   If you enter yes, only the internal database drive is formatted.
   Enter yes only if you know that the internal database drive needs formatting.
   If you are unsure, enter no.
   [Default: no]:
   ```

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

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 back.

- To recover all storage nodes, enter **2**.

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

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

   HCP Setup performs a series of prechecks and, if they are successful, recovers the OS on the selected nodes or all nodes, as applicable. If any of the prechecks fail, HCP Setup exits. In this case, fix the problem and then start the OS recovery procedure again.

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

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

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

   The **HCP Service** menu reappears.

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

8. From the **HCP Service** menu, enter **q** to return to the **HCP 8.1 Configuration** menu.

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

10. From the **HCP 8.1 Configuration** menu, enter **q** to log out of the install shell.

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

Recovering storage nodes (clearing storage volumes)

To recover storage nodes when clearing storage volumes:

1. From the **HCP 8.1 Configuration** menu, enter **s** to display the **HCP Service** menu.

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

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

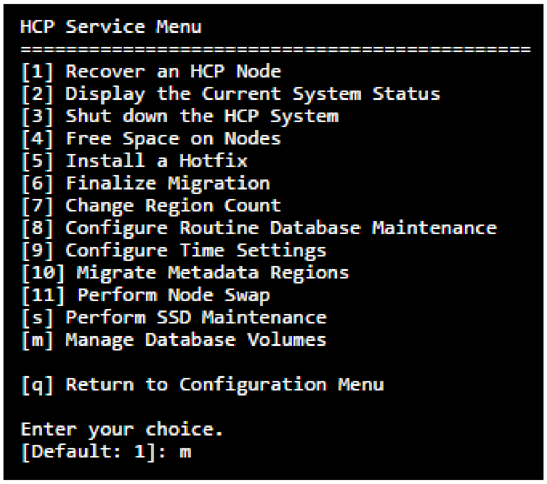

```
HCP Service Menu
=============================================
[1] Recover an HCP Node
[2] Display the Current System Status
[3] Shut down the HCP System
[4] Free Space on Nodes
[5] Install a Hotfix
[6] Finalize Migration
[7] Change Region Count
[8] Configure Routine Database Maintenance
[9] Configure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

[q] Return to Configuration Menu

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

3. From the **HCP Service** menu, enter **1** for recovery operations.

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

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

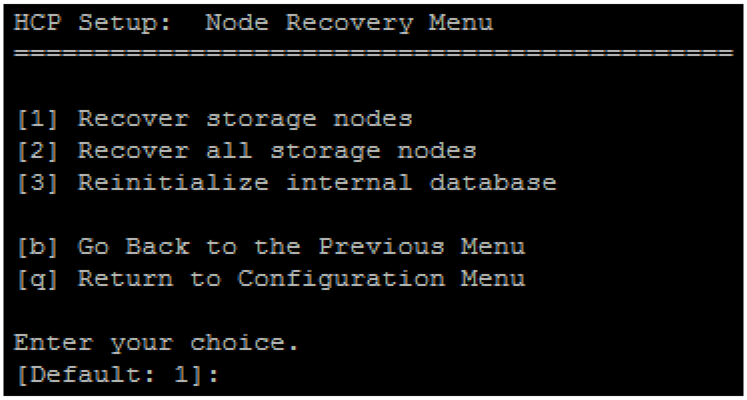

```
HCP Setup:  Node Recovery Menu
=============================================

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

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

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

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

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

Optionally, if you set FORCE_FORMAT to 1, when you enter *y* or *yes*, HCP Setup displays the **FORCE_FORMAT** prompt.

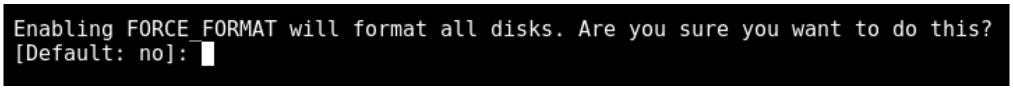

```
Enabling FORCE_FORMAT will format all disks. Are you sure you want to do this?
[Default: no]:
```

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

6. Enter *y* or *yes* to allow the system to format all disks.

   Then follow the on-screen instructions to identify the node containing the logical volumes you want to recover.

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

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

   After you have entered the unique key, HCP Setup performs a set of prechecks.

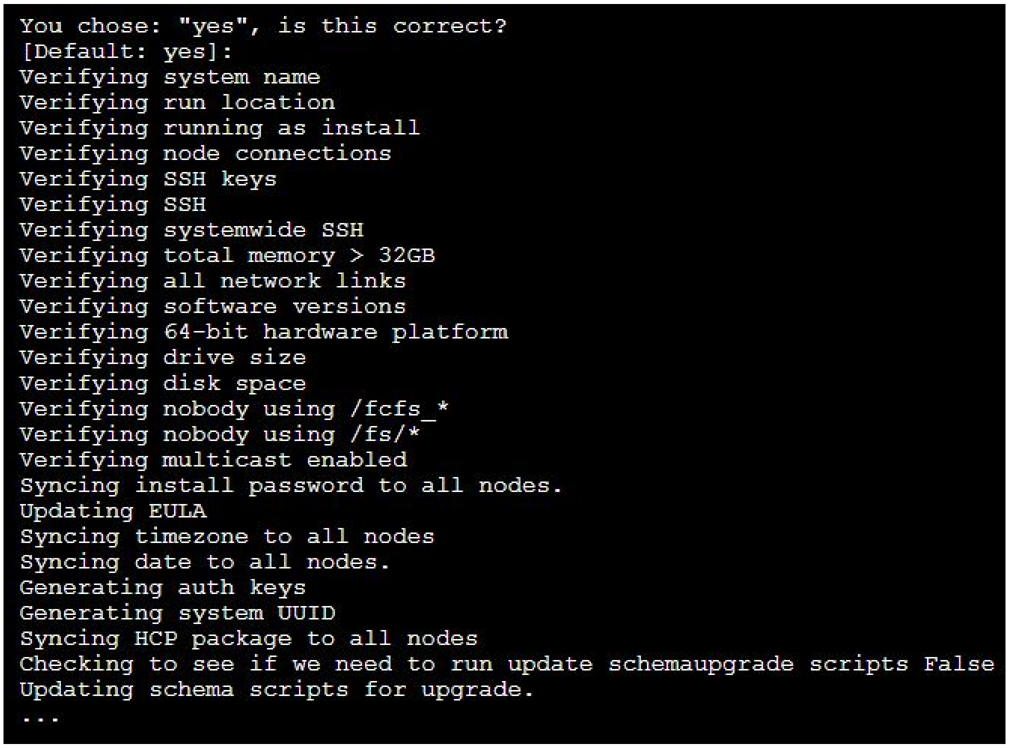

```
You chose: "yes", is this correct?
[Default: yes]:
Verifying system name
Verifying run location
Verifying running as install
Verifying node connections
Verifying SSH keys
Verifying SSH
Verifying systemwide SSH
Verifying total memory > 32GB
Verifying all network links
Verifying software versions
Verifying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

Only if your current HCP system has dedicated database volumes:

1. When prompted, select the dedicated database volume for the first node.

2. Press Enter to confirm your selection.

3. Repeat the above two steps for each node in the system.

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

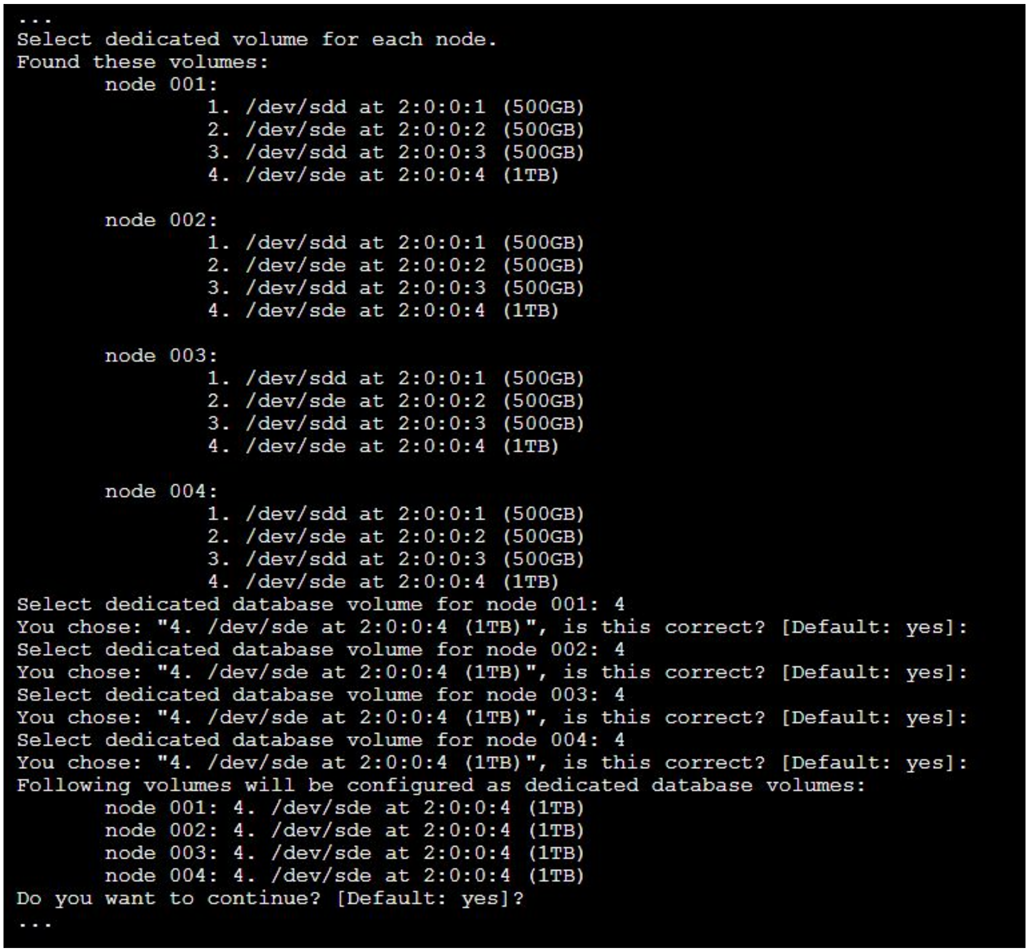

```
...
Select dedicated volume for each node.
Found these volumes:
       node 001:
               1. /dev/sdd at 2:0:0:1 (500GB)
               2. /dev/sde at 2:0:0:2 (500GB)
               3. /dev/sdd at 2:0:0:3 (500GB)
               4. /dev/sde at 2:0:0:4 (1TB)

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

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

       node 004:
               1. /dev/sdd at 2:0:0:1 (500GB)
               2. /dev/sde at 2:0:0:2 (500GB)
               3. /dev/sdd at 2:0:0:3 (500GB)
               4. /dev/sde at 2:0:0:4 (1TB)
Select dedicated database volume for node 001: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default: yes]:
Select dedicated database volume for node 002: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default: yes]:
Select dedicated database volume for node 003: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default: yes]:
Select dedicated database volume for node 004: 4
You chose: "4. /dev/sde at 2:0:0:4 (1TB)", is this correct? [Default: yes]:
Following volumes will be configured as dedicated database volumes:
       node 001: 4. /dev/sde at 2:0:0:4 (1TB)
       node 002: 4. /dev/sde at 2:0:0:4 (1TB)
       node 003: 4. /dev/sde at 2:0:0:4 (1TB)
       node 004: 4. /dev/sde at 2:0:0:4 (1TB)
Do you want to continue? [Default: yes]?
...
```

4. Press Enter to continue the procedure.

   If the prechecks are successful, HCP Setup recovers all the logical volumes on the selected nodes or all nodes, as applicable. If any of the prechecks fail, HCP Setup exits. In this case, fix the problem and then start the node recovery procedure again.

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

Important: If HCP Setup exits at any time before the node recovery processing is complete, please contact your HCP support center for help. Do *not* try the node recovery again.

8. From the **HCP Service** menu, enter **q** to return to the **HCP 8.1 Configuration** menu.

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

10. From the **HCP 8.1 Configuration** menu, enter **q** to log out of the install shell.

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

Adding a management port network

To add a management port network to an existing HCP-VM configuration:

1. On each node in the vSphere cluster, configure a network to be used as a management port network. For information on how to configure networking for management port switching, see "Configuring networking" on page 25.

2. Perform the following steps on each virtual machine in the HCP-VM cluster:

    a. From the vSphere web client, select the HCP virtual machine.

    b. Click on the **Configure** tab.

    c. Click on **Edit**.

    d. In the **New device** dropdown menu, select **Network** then click on **Add**.

    e. When the new network device appears, select **Management Port Network**.

    f. Click on **OK**.

3. Configure the management port network from the System Management Console. For information on how to configure the management port network using the System Management Console, see *Administering HCP*.

6

Configuring HCP monitoring with Hi-Track Monitor

Hi-Track Monitor is a Hitachi Vantara product that enables remote monitoring of the nodes in an HCP-VM system. With Hi-Track Monitor, you can view the status of these components in a web browser. You can also configure Hi-Track Monitor to notify you by email of error conditions as they occur. Additionally, you can configure Hi-Track Monitor to report error conditions to Hitachi Vantara support personnel.

Hi-Track Monitor is for monitoring and error notification purposes only. It does not allow any changes to be made to the system.

Hi-Track Monitor 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 Hi-Track servers. However, Hi-Track support for IPv6 network connections varies based on the Hi-Track server operating system. For information on requirements for Hi-Track servers that support IPv6 networks, see the applicable Hi-Track documentation.

This chapter explains how to set up monitoring of HCP nodes with Hi-Track Monitor.

This chapter assumes that Hi-Track Monitor is already installed and running according to the documentation that comes with the product.

Enabling SNMP in HCP

To enable Hi-Track Monitor to work with HCP, you need to enable SNMP in the HCP System Management Console. When you enable SNMP, you can select version 1 or 2c or version 3.

By default, Hi-Track Monitor 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 Hi-Track Monitor to match what you specify in HCP. For more information on this, see the Hi-Track Monitor documentation.

To enable SNMP in HCP for use with Hi-Track Monitor:

1. Log into the HCP System Management Console using the initial user account, which has the security role.

2. In the top-level menu of the System Management Console, select **Monitoring ▶ SNMP**.

3. In the **SNMP Settings** section on the **SNMP** page:

   - Select the **Enable SNMP at snmp.*hcp-domain-name*** option.

   - Select either **Use version 1 or 2c** (recommended) or **Use version 3**.

     If you select **Use version 3**, specify a username and password in the **Username**, **Password**, and **Confirm Password** fields.

   - Optionally, in the **Community** field, type a different community name.

4. Click on the **Update Settings** button.

5. In the entry field in the **Allow** section, type the IP address that you want HCP to use to connect to the server on which Hi-Track Monitor is installed. Then click on the **Add** button.

6. Log out of the System Management Console and close the browser window.

Configuring Hi-Track Monitor

To configure Hi-Track Monitor to monitor the nodes in the HCP system, follow the steps outlined in the table below.

Step	Activity	More information
1	Log into Hi-Track Monitor.	Step 1: "Log into Hi-Track Monitor" below
2	Set the Hi-Track Monitor base configuration, including the email addresses to which email about error conditions should be sent.	Step 2: "Set the base configuration" on the next page
3	Optionally, configure transport agents for reporting error conditions to Hitachi Vantara support personnel.	Step 3 (conditional): "Configure transport agents" on page 83
4	Identify the HCP system to be monitored.	Step 4: "Identify the HCP system" on page 84

Step 1: Log into Hi-Track Monitor

To log into Hi-Track Monitor:

1. Open a web browser window.

2. In the address field, enter the URL for the Hi-Track Monitor server (using either the hostname or a valid IP address for the server) followed by the port number 6696; for example:

   ```
   http://hitrack:6696
   ```

3. In the **Select one of the following UserIds** field, select **Administrator**.

4. In the **Enter the corresponding password** field, type the case-sensitive password for the Administrator user. By default, this password is *hds*.

   If Hi-Track Monitor is already in use at your site for monitoring other devices, this password may have been changed. In this case, see your Hi-Track Monitor administrator for the current password.

5. Click on the **Logon** button.

Step 2: Set the base configuration

The Hi-Track Monitor base configuration specifies information such as the customer site ID, how frequently to scan devices, and whether to report communication errors that occur between Hi-Track Monitor and monitored devices. The base configuration also specifies the addresses to which Hi-Track Monitor should send email about error conditions.

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

To set the Hi-Track Monitor base configuration:

1. In the row of tabs at the top of the Hi-Track Monitor interface, click on **Configuration**.

   The **Base** page is displayed by default. To return to this page from another configuration page, click on **Base** in the row of tabs below **Configuration**.

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

   - In the **Site ID** field, type your Hitachi Vantara customer ID. If you don’t know your customer ID, contact your authorized HCP service provider for help.

   - Optionally, specify different values in the other fields to meet the needs of your site. For information on these fields, click on the **Help on this table’s entries** link above the fields.

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

   - In the **eMail Server** field, type the fully qualified hostname or a valid IP address of the email server through which you want Hi-Track Monitor to send email about error conditions.

   - In the **Local Interface** field, select the Ethernet interface that has connectivity to the specified email server. (This is the interface on the Hi-Track Monitor server.)

   - In the **User List** field, type a comma-separated list of the email addresses to which Hi-Track Monitor should send email about error conditions.

- ○ In the **Sender's Email Address** field, type a well-formed email address to be used in the From line of each email.

  Some email servers require that the value in the From line be an email address that is already known to the server.

4. Click on the **Submit** button.
5. Optionally, to send a test email to the specified email addresses, click on the **Test Email** button.

Step 3 (conditional): Configure transport agents

A Hi-Track Monitor transport agent transfers notifications of error conditions to a target location where Hitachi Vantara support personnel can access them. The transfer methods available are HTTPS, FTP, or dial up. For the destinations for each method, contact your authorized HCP service provider.

You can specify multiple transport agents. Hi-Track tries them in the order in which they are listed until one is successful.

To configure a transport agent:

1. In the row of tabs below **Configuration**, click on **Transport Agents**.
2. In the field below **Data Transfer Agents**, select the transfer method for the new transport agent.
3. Click on the **Create** button.

   The new transport 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 on what to specify, see the Hi-Track Monitor documentation.
5. Click on the **Submit** button.

You can change the order of multiple transport agents by moving them individually to the top of the list. To move a transport agent to the top of the list:

1. In the **Move to Top** column, select the transport agent you want to move.
2. Click on the **Submit** button.

Step 4: Identify the HCP system

To identify the HCP system to be monitored:

1. In the row of tabs at the top of the Hi-Track Monitor interface, click on **Summary**.

   The **Summary** page displays up to four tables that categorize the devices known to Hi-Track Monitor — Device Errors, Communication Errors, Devices Okay, and Not Monitored. To show or hide these tables, click in the checkboxes below the table names at the top of the page to select or deselect the tables, as applicable. Then click on the **Refresh** button.

   While no tables are shown, the page contains an **Add a device** link.

2. Take one of these actions:

   - If the **Summary** page doesn't display any tables, click on the **Add a device** link.

   - If the **Summary** page displays one or more tables, click on 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. Optionally, in the **Name** field, type a name for the HCP system. The name can be from one through 40 characters long. Special characters and spaces are allowed.

   Typically, this is the hostname of the system.

5. Optionally, in the **Location** field, type the location of the HCP system. The location can be from one through 40 characters long. Special characters and spaces are allowed.

6. Optionally, in the **Group** field, type the name of a group associated with the HCP system (for example, Finance Department). The group name can be from one through 40 characters long. Special characters and spaces are allowed.

7. In the **Site ID** field, type your Hitachi Vantara customer ID. If you don't know your customer ID, contact your authorized HCP service provider for help.

8. In the **IP Address or Name (1)** field, type a valid front-end IP address for the lowest-numbered storage node in the HCP system. In the **Local Interface** field, leave the value as **-any-**.

9. In the **IP Address or Name (2)** field, type a valid front-end IP address for the highest-numbered storage node in the HCP system. In the **Local Interface** field, leave the value as **-any-**.

10. In the **SNMP Access ID** field, select the SNMP user that corresponds to the SNMP configuration in HCP. Typically, this is **public**.

    For information on configuring SNMP in HCP, see Enabling SNMP in HCP.

11. In the **Comms Error Reporting** field, select one of these options to specify whether Hi-Track should report communication errors that occur between Hi-Track Monitor and the HCP system:

    - **Yes** — Report communication errors.

    - **No** — Don’t report communication errors.

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

    - **Default** — Use the setting in the base configuration.

12. Leave **Enabled** selected.

13. Leave **Trace** unselected.

14. Click on the **Add** button.

    If the operation is successful, the interface displays a message indicating that the HCP system has been added. Do not click on the **Add** button again. Doing so will add the system a second time.

A

Changing the HCP-VM network adapters

This appendix covers changing between the supported network adapters on your HCP-VMs.

About network adapters

HCP supports two types of network adapters: VMXNET3 (recommended) and e1000. With release 7.2.1 of HCP, all newly installed HCP-VMs are automatically configured to use VMXNET3 adapters. You can configure older HCP-VMs to also use VMXNET3 adapters or you can configure new HCP-VMs to use e1000.

e1000 network adapters only support one gigabit network configurations. VMXNET3 support both one and 10 gigabit network configurations. Especially if you have a 10 gigabit network configuration, VMXNET3 network adapters are recommended.

Disabling LRO on the ESXi host for VMXNET3

If you are using or want to switch your HCP-VMs to the VMXNET3 network adapter, you need to disable LRO in the guest Operating System to prevent potential TCP performance degradation.

To disable the LRO on the ESXi host:

1. Log in to the vSphere Client.

2. Click on a server that hosts ESXI for your HCP-VMs.

3. Click on the **Configure** tab.

4. In the **System** menu, click on **Advanced System Settings**.

5. In the inventory tree, click on **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 on **OK**.

8. From the vSphere client, reboot the server by right-clicking on it, and in the dropdown menu clicking **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. The following steps show you how to switch adapters.

Step 1: Power off the HCP-VM

Before you can switch the network adapter, you need to power off the HCP-VM. To power off the HCP-VM:

1. Open your vSphere client.

2. Right-click on the HCP-VM that needs to have its network adapter replaced.

   A drop down menu appears.

3. In the drop down menu, hover your cursor over **Power** and in the second dropdown menu that opens, click on **Shut Down Guest OS**.

Step 2: Remove the previous network adapters

Once the HCP-VM is powered off, you need to remove the previous network adapters from the HCP-VMs. To remove the previous adapters:

1. From the vSphere client, right-click on one of the powered-off HCP-VMs.

A dropdown menu appears.

2. In the drop down menu that appears, click on **Edit Settings...**.

   A **Virtual Machine Properties** window appears.

   The number of existing network adapters varies depending on whether the HCP-VM is currently using e1000 or VMXNET3. If the HCP-VM is using e1000, you need to remove four network adapters. If the HCP-VM is using VMXNET3, you need to remove two network adapters. This procedure shows an HCP-VM using e1000 and switching to VMXNET3.

3. In the **Virtual Hardware** tab of the **Edit Settings** window, select a network adapter and click on **Remove**. Repeat this step until all network adapters are removed.

4. Click on **OK**.

Step 3: Set the front-end network adapters

After you have removed the previous network adapters, you need to configure the front-end network adapters. To set the front-end network adapters:

1. From the vSphere client, right-click on one of the powered-off HCP-VMs.

   A dropdown menu appears.

2. In the drop down menu that appears, click on **Edit Settings...**.

   A **Virtual Machine Properties** window appears.

3. In the **Edit Settings** window **Hardware** tab, click on **Add**.

   An **Add Hardware** window opens.

4. In the **Add hardware** window that opens, click on **Ethernet Adapter**.

5. Click on **Next**.

6. On the **Network Connection** page in the **Adapter Type** panel **Type** dropdown field, select **VMXNET3**.

7. In the **Network connection** panel, select **Named network with specified label**.

8. Select **Front-End Network**.

9. Click on **Next**.

10. On the **Verification** page, click on **Finish**.

11. Back in the **Virtual Machine Properties** window, click on **OK**.

    The first VMXNET3 network adapter needs to be connected to the front-end network. The second VMXNET3 network adapter needs to be connected to the back-end network. For more information on connecting the second VMXNET3 network adapter to the back-end network, see Step 4: "Set the back-end network adapters" below

Step 4: Set the back-end network adapters

Once the front-end network adapters are set, you need to configure the back-end network adapters. To set the back-end network adapters:

1. From the vSphere client, right-click on the powered-off HCP-VM.

   A dropdown menu appears.

2. In the drop down menu that appears, click on **Edit Settings...**.

   A **Virtual Machine Properties** window appears.

3. In the **Edit Settings** window **Hardware** tab, click on **Add**.

   An **Add Hardware** window opens.

4. In the **Add hardware** window that opens, click on the **Ethernet Adapter**.

5. Click on **Next**.

6. On the **Network Connection** page in the **Adapter Type** panel **Type** dropdown field, select **VMXNET3**.

7. In the **Network connection** panel, select **Named network with specified label**.

8. Select **Back-End Network**.

9. Click on **Next**.

10. On the **Verification** page, click on **Finish**.

11. Back in the **Virtual Machine Properties** window, click on **OK**.

    The first VMXNET3 network adapter needs to be connected to the front-end network. The second VMXNET3 network adapter needs to be connected to the back-end network. For more information on connecting the VMXNET3 network adapter to the front-end network, see the previous step Step 3: "Set the front-end network adapters" on page 89

Step 5: Power on the HCP-VM

Once you have configured the HCP-VM network adapters, you need to power on the HCP-VM. To power on the HCP-VM:

1. From your vSphere client, right-click on the newly configured HCP-VM.

   A dropdown menu appears.

2. In the dropdown menu, hover your cursor over **Power** and in the second dropdown menu that opens click on **Power On**.

Once the HCP-VM is powered on you have successfully configured its network adapter. If you have multiple HCP-VM nodes that need to be reconfigured, repeat the changing network adapter procedure for the other HCP-VMs.

B

Changing the DRS settings

To modify the DRS settings:

1. Access the vSphere Client.
2. In the left side navigation bar, select the datacenter.
3. In the right side window, under the **Getting Started** tab, click on **Create a cluster**.
4. In **VMware Cluster Wizard**, select **Turn On vSphere HA** and **Turn On vSphere DRS**.
5. Click on **OK**.

Important: Only turn on this feature if you feel your environment will benefit from it and you fully understand its functionality.

6. Select **Manual** for the DRS automation level in order to specify where VM guests should reside.
7. Select **Off** for the Power Management.
8. Select **Enable Host Monitoring** and keep the default settings.
9. Set the **VM Monitoring** to **Disabled**.
10. Select **Disable EVC**.
11. Select where you want to store your Swapfile location.
12. Review your settings then click on **OK**.

13. In the left side navigation bar, select a Cluster and right-click it. Then click on **Edit Settings**.

14. On the left side navigation bar of the **Settings** window, click on **DRS Groups Manager**.

15. In the **DRS Groups Manager** , create a group and add the Virtual Machines that you would like to keep on a specific server.

16. Create one Virtual Machine DRS Group for each Host.

17. Click on **Add** in the Host DRS Groups section and place one host from the cluster in each group, then click on **Next**.

18. On the left side navigation bar of the **Settings** window, click on **Rules**.

19. Create a new Rule where each VM group is matched to a Host Group, and set the type of rule to be **Virtual Machines to Hosts**.

20. Select **Should run on hosts in group** then click on **OK**.

Note: You will create a rule that lets VMs run on other hosts in the event of a failure. We will also setup a rule to alert you if that failure occurs. If you select **Must Run on Hosts in Group** then HA will not bring the server up on another in the cluster in the even of Host failure defeating the purpose of HA.

Setting an alarm

To set an alarm:

1. Right-click on the Cluster and hover your cursor over **Alarm** in the submenu. Then click on **Add Alarm**.

2. In the **Alarm Settings** window, name your alarm and set the **Monitor** to **Virtual Machines**.

3. Select **Monitor for specific event occurring on this object**.

4. Go to the **Triggers** tab and select **VM is violating a DRS VM-Host affinity rule**.

5. Set the status to either warning or alert depending on how severe you think it should be.

6. Under the **Trigger Conditions** select an **Argument of VM** name.

7. Set the Value equal to each VM you want to monitor.

8. Add one argument for each VM.

9. Set the Actions you want the system to take.

C

Configuring the HCP-VM small instance

Optionally, if you are deploying an HCP-VM system as a small instance system, before powering on the HCP-VM nodes you need to change the CPU count and RAM for each node:

1. In vSphere Client right-click on the HCP-VM node and choose **Edit Settings** from the context menu.

2. Select the **Hardware** tab in the **Virtual Machine Properties** window.

3. Select **Memory** from the hardware list and adjust the allocation to 16 GB in the **Memory Configuration** pain.

4. Select **CPUs** from the hardware list and adjust the **Number of virtual sockets** and **Number of core per sockets** so that the **Total number of cores** equals 4.

5. Click on **OK** to save your changes and close the **Virtual Machine Properties** window.

D

Managing failover

The HCP-VM vSphere HA cluster does **not** automatically move the failed-over HCP-VM node back to its original ESXi host once the server or ESXi host is available. An HCP-VM system administrator needs to manually shutdown the HCP-VM node(s) that need to be moved to another ESXi host.

Alternatively, the vCenter administrator can issue a shutdown of the HCP-VM node from the vCenter management console.

The vCenter administrator will then manually move the HCP-VM node onto the preferred ESXi host, and power on the HCP-VM node. Once the HCP-VM node boots, it will re-join the HCP-VM system.

After powering down an HCP-VM node and attempting to move that VM to another ESXi host with some VMware configurations, you may see an error message that can be safely ignored.

Glossary

A

access control list (ACL)

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

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

ACL

See access control list (ACL).

Active Directory (AD)

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

AD

See Active Directory (AD).

alert

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

C

capacity

The total amount of primary storage space in HCP, excluding the space required for system overhead for all data to be stored in primary running storage and primary spindown storage, including the fixed-content data, metadata, any redundant data required to satisfy services plans, and the metadata query engine index.

CIFS

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

custom metadata

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

D

database

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

data center

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

data protection level (DPL)

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

datastore

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

domain

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

domain name system

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

DNS

See domain name system.

DPL

See data protection level (DPL).

E

ESXi

See "VMware ESXi".

H

Hitachi Content Platform (HCP)

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

HCP VM system

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

HDDS

See "Hitachi Data Discovery Suite (HDDS)"

Hitachi Data Discovery Suite (HDDS)

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

host

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

L

logical unit number (LUN)

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

logical volume

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

LUN

See "logical unit number (LUN)".

M

metadata

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

multipathing

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

N

namespace

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

network

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

network file system

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

network interface controller (NIC)

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

NFS

See network file system.

NIC

See "network interface controller (NIC)".

node

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

O

object

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

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

P

ping

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

pNIC

See "network interface controller (NIC)".

Q

query

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

R

RAIN

See redundant array of independant nodes (RAIN).

redundant array of independant nodes (RAIN)

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

replication

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

repository

The aggregate of the namespaces defined for an HCP system.

running storage

Storage on continuously spinning disks.

S

SAIN

See "SAN-attached array of independent nodes (SAIN)".

SAN-attached array of independent nodes (SAIN)

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

search console

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

search facility

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

secure shell

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

secure sockets layer

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

service

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

service plan

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

simple network management protocol (SNMP)

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

SNMP

See "simple network management protocol (SNMP)".

SNMP trap

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

spindown storage

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

SSH

See "secure shell".

SSL

See secure sockets layer.

SSL server certificate

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

storage node

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

subdomain

A subset of the computers and devices in a domain.

switch

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

syslog

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

system management console

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

T

tag

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

tagged network

A network that has a VLAN ID.

tenant

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

tenant management console

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

transaction log

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

U

unix

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

upstream DNS server

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

user account

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

user authentication

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

V

vCenter

See "VMware vCenter Server".

versioning

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

virtual local area network (VLAN)

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

virtual machine

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

VLAN

See Virtual Local Area Network (VLAN).

VLAN ID

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

vmNIC

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

VMware ESXi

The underlying operating system for the VMware vSphere product.

VMware vCenter Server

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

vNIC

See "network interface controller (NIC)".

Z

zero-copy failover

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

Index

Hitachi Vantara

Corporate Headquarters
2845 Lafayett Street
Santa Clara, CA 95050-2639 USA
www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hitachivantara.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hitachivantara.com
Asia Pacific: +852 3189 7900 or info.marketing.apac@hitachivantara.com