

Hitachi Content Platform

Release 7.3.3

Deploying an HCP-VM System

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.

© 2017 Hitachi Vantara Corporation. All rights reserved.

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

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

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara 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 HitachiVantara.com.

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 to access relevant data; 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.

EXPORT CONTROLS - Licensee will comply fully with all applicable export laws and regulations of the United States and other countries, and Licensee shall not export, or allow the export or re-export of, the Software, API, or Materials in violation of any such laws or regulations. By downloading or using the Software, API, or Materials, Licensee agrees to the foregoing and represents and warrants that Licensee is not located in, under the control of, or a national or resident of any embargoed or restricted country.

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	7
Intended audience	7
Product version	7
Related documents	7
Accessing product documentation	10
Getting help	11
Comments	11
Chapter 1: HCP system overview	13
Introduction to Hitachi Content Platform	13
HCP-VM system components and architecture	13
Host platform	14
Compute	14
Storage	14
HCP network connectivity	16
Front-end network	17
Storage network	17
Back-end network	17
Hardware monitoring and alerting	18
HCP software	18
HCP upgrades	19
HCP search nodes	19
HCP-VM node failover (vCenter and vSphere HA)	19
Storage licensing	20
Chapter 2: Configuration guidelines for HCP-VM environment	21
VMware supported versions	21
VMware supported functionality	21

Prerequisites and recommendations	22
HCP-VM system limits	24
HCP-VM availability considerations	25
Chapter 3: Configuring the HCP-VM environment	27
ESXi considerations	27
Enabling NTP for the ESXi hosts	27
Configure vSphere HA cluster for HCP-VM (Recommended)	30
Provisioning HCP-VM storage	45
Add datastores to vSphere HA cluster	49
NFS Datastores	60
Creating an NFS datastore	61
Heartbeat datastore selection	62
Preparing the ESXi network for the HCP-VM OVF deployment	63
Configuring the Storage Network (HNAS Best Practice)	64
Configuring networking for Front-end switching	65
Configure networking for Back-end switching	71
Verifying ESXi configuration on all hosts	73
Chapter 4: Creating the HCP-VM system	79
Unpacking the OVF Zip file	79
Deploying the HCP-VM OVF VDMK	79
Deploy the HCP-VM OVF RDM	91
Configuring the HCP-VM small instance	107
Configuring the HCP-VM network	108
Install HCP software	113
Step 1: Identify the nodes in the HCP system	115
Configure the HCP system	116
Execute the installation	123
Verifying the HCP installation	127
Setting additional configuration options	129
Monitoring and alerting	130
Monitoring and alerts	130
Software monitoring	130
HCP resource monitoring	131
HCP diagnostic menu	131
Chapter 5: Setting the HCP-VM network adapters	133
About network adapters	133

Disabling LRO on the ESXi host for VMXNET3	133
Setting the HCP-VM network adapter	136
Step 1: Power off the HCP-VM	136
Step 2: Remove the previous network adapters	137
Step 3: Change the guest OS	139
Step 4: Set the Front-End network adapters	140
Step 5: Set the Back-End network adapters	142
Power on the HCP-VM	145
Chapter 6: Failover management	147
Maintenance procedures	148
Adding logical volumes	148
Adding HCP-VM nodes	155
Chapter 7: Configuring HCP monitoring with Hi-Track Monitor	157
Enabling SNMP in HCP	158
Configuring Hi-Track Monitor	159
Step 1: Log into Hi-Track Monitor	159
Step 2: Set the base configuration	160
Step 3 (conditional): Configure transport agents	161
Step 4: Identify the HCP system	162
Appendix A: Configuring networking for HCP virtual network management	165
HCP networking information	165
OVF deployment information	169
Appendix B: Changing the VMDK target size	171
Appendix C: DRS settings	173
Setting an alarm	185
Glossary	189
Index	201



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 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 7.3.3 of Hitachi Content Platform.

Related documents

The following documents contain additional information about Hitachi Content Platform:

- *Administering HCP* — This book explains how to use an HCP system to monitor and manage a digital object repository. The book describes the capabilities and the hardware and software components of the system. The book presents both the concepts and instructions you need to configure the system, including creating the tenants that administer access to the repository. The book also covers the processes that maintain the integrity and security of the repository contents.

- *Managing a Tenant and Its Namespaces* — This book contains complete information for managing the HCP tenants and namespaces created in an HCP system. The book provides instructions for creating namespaces, setting up user accounts, configuring the protocols that allow access to namespaces, managing search and indexing, and downloading installation files for HCP Data Migrator. The book also explains how to work with retention classes and the privileged delete functionality.
- *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.
- *Replicating Tenants and Namespaces* — This book covers all aspects of tenant and namespace replication. Replication is the process of keeping selected tenants and namespaces in two or more HCP systems in sync with each other to ensure data availability and enable disaster recovery. The book describes how replication works, contains instructions for working with replication links and erasure coding topologies and explains how to manage and monitor the replication process.
- *HCP Management API Reference* — This book contains the information you need to use the HCP management API. This RESTful HTTP API enables you to create and manage tenants and namespaces programmatically. The book explains how to use the API to access an HCP system, specify resources, and update and retrieve resource properties.
- *Using a Namespace* — This book describes the properties of objects in HCP namespaces. This book provides instructions for using the HTTP, WebDAV, CIFS, and NFS protocols for the purpose of storing, retrieving, and deleting objects, as well as changing object metadata such as retention and shred settings. The book also explains how to manage namespace content and view namespace information in the Namespace Browser.
- *Using the HCP HS3 API* — This book contains the information you need to use the HCP HS3 API. This S3™-compatible, RESTful, HTTP-based API enables you to work with buckets and objects in HCP. The book introduces the HCP concepts you need to understand in order to use

HS3 effectively and contains instructions and examples for each of the bucket and object operations you can perform with HS3.

- *Using the HCP HSwift API* — This book contains the information you need to use the HCP HSwift API. This OpenStack Swift compatible, RESTful, HTTP-based API enables you to work with containers and objects in HCP. The book introduces the HCP concepts you need to understand in order to use HSwift effectively and contains instructions and examples for each of the container and object operations you can perform with HSwift.
- *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.
- *HCP Metadata Query API Reference* — This book describes the HCP metadata query API. This RESTful HTTP API enables you to query namespaces for objects that satisfy criteria you specify. The book explains how to construct and perform queries and describes query results. It also contains several examples, which you can use as models for your own queries.
- *Searching Namespaces* — This book describes the HCP Search Console (also called the Metadata Query Engine Console). It explains how to use the Console to search namespaces for objects that satisfy criteria you specify. It also explains how to manage and manipulate queries and search results. The book contains many examples, which you can use as models for your own searches.
- *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](http://portal.hitachivantara.com) 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: <http://portal.hitachivantara.com>

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



Note: If you purchased HCP 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 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!

HCP system overview

This chapter introduces HCP, and describes the architecture for an HCP 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 as both data and metadata. **Metadata** is responsible for describing the object. HCP distributes these objects across the storage space, and represents them as either URLs or 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 Hitachi Content Platform Virtual Machine (**HCP-VM**) system.

The illustration below shows the architecture of an HCP-VM system.



Host platform

In an HCP-VM system, each node runs in a virtual machine on an ESXi host. It is recommended that you only run one HCP-VM node on a single ESXi host.

Compute

Each HCP-VM node will have at least eight vCPUs and at least 32 GB of RAM allocated. This will enable the system to maintain performance for most client workloads and HCP system activities like encryption, scheduled service runs, and routine database maintenance. If deploying HCP-VM Small Instance configuration each HCP-VM node will have at least four vCPUs and at least 16 GB of RAM allocated.

Storage

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

Hitachi NAS (**HNAS**).

SAN storage must provide at least two paths to each Logical Unit Number (**LUN**) and each of those LUNs must be presented to each ESXi host with the exact same LUN number (**HLUN**).

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 will not be shared by HCP-VM nodes. However, HCP-VM nodes can have multiple datastores. Each datastore will be carved into one or multiple Virtual Machine Disks (**VMDK**) which are presented to the HCP OS as local disks. The HCP OS will recognize its storage as internal drives similar to HCP300 RAIN nodes. The disks will be 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 2TB minus 512bytes.

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. The recommended best practice is to spread datastores across multiple DP Pools so as to avoid resource contention and single points of failure as much as possible.
- 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 VMFS datastores.
 - 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 VMFS datastore. This configuration does not support vSphere High Availability (HA). The underlying storage in this configuration must be RAID protected.

When varying from the recommended configuration, careful consideration needs to be taken when planning the storage for an HCP-VM system. Many factors such as: 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, support windows, etc. Health of the systems must be monitored closely with any failures serviced immediately so as to ensure underlying storage does not fail.

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

Always follow the vendor's best practices for configuring their storage in a VMware environment.

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

HCP's Front-end network is used for management of the system and client access. For HCP's Front-end network, it is recommended that the ESXi host will present two vNICs on a second pair of pNICs. Best practice is to have those pNICs dedicated to HCP for greater redundancy and more consistent performance.

When two NICs are not available for the Front-end network, it is possible to operate HCP-VM with one NIC provided it has enough available bandwidth to support data traffic and management access. It should be noted that in the event of a failure of that single NIC, the HCP-VM node(s) that reside on that ESXi host will not be available to clients, but will still be available to the HCP-VM system through the Back-end network.

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.

Back-end network

HCP's private Back-end network is used for inter-node communication and data transfer between nodes.

Due to the inter-node communication, it is mandatory that the back-end network is configured to allow multicast communication between all nodes in the HCP-VM system. In most cases, it is not enough to just have multicast enabled on the switch. There will most likely be additional configuration parameters necessary to allow for the multicast traffic. Follow the switch vendor documentation to configure the network to allow multicast traffic between the HCP-VM nodes. It should be noted that it is possible to deploy an HCP-VM system via the system installation program, and not have multicast configured correctly. The processes that require multicast communication are not active until after installation. When multicast is not configured correctly, the HCP-VM system will attempt to boot to its operational runlevel, only to fall back to a lower runlevel once multicast communication fails.

For HCP's Back-end network, it is recommended that the ESXi host present two vmNICs which directly map to two physical NICs (pNICs) on the ESXi host server. It is recommended the pNICs be connected to two physical switches, on an isolated network, where pNIC-1 on all ESXi hosts connect to

the same physical switch (switch1) and pNIC-2 on all ESXi hosts connect to the same physical switch (switch2). When using two physical switches, there must be an inter-switch connection that allows connectivity for all HCP-VM network ports on one switch to all HCP-VM network ports on the second switch. The pNICs and the switches to which they are connected should be isolated from all other networks in the customer environment. The switches must be configured with spanning tree disabled, allow multicast traffic, be at least 1GbE and should be dedicated to HCP to guarantee data security and HCP reliability.

In the event that the HCP-VM back-end network travels over a public network, it is strongly recommended that the HCP-VM system reside on its own VLAN.

When two NICs are not available for the Back-end network, it is possible to operate HCP-VM with one NIC provided that it has enough available bandwidth to support data traffic and inter-node communication. It should be noted that in the event of a failure of that single NIC, the HCP-VM node (s) that reside on that ESXi host will not be available to the HCP-VM system.

Hardware monitoring and alerting

The HCP hardware based appliance has built in redundant hardware, monitoring, alerting and failover behavior that cannot be leveraged in a virtualized VMware environment. To maintain performance and data integrity, it is recommended that all underlying hardware associated with the HCP-VM system be treated as mission critical and monitored for failures. Whenever Hitachi servers, storage and networking are part of the HCP-VM system, it is recommended they be connected to HiTrack. Any non-Hitachi equipment should be closely monitored using the vendor or customer equivalent to HiTrack. 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 lengthy RAID rebuild times.

HCP software

HCP-VM provides all the same (non-hardware specific) functionality as HCP RAIN and SAIN systems. Data is RAID protected, and HCP policies and services ensure its integrity and security and optimize its space used on disk. The management and data access interfaces are the same as for RAIN and SAIN systems. A small amount of features are not available in an HCP-VM system because they are physical hardware based, so they are not practical or feasible in a virtualized environment.

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-VMs come with a storage license that provides two terabytes of active storage and a two terabytes of extended storage. If you need more storage space, please contact your Hitachi Vantara sales representative to purchase more storage license capacity.

If you upgrade HCP to version 7.1 or later, you receive an unlimited storage license that applies to both active and extended storage for one year.

For more information about storage licensing, see *Administering HCP*.

Configuration guidelines for HCP-VM environment

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

VMware supported versions

HCP-VM supports multiple versions of VMware. For more information on supported VMware versions, see the *HCP 7.3.3 Release Notes*.

VMware supported functionality

HCP-VM supports the following VMware functionality:

- vSphere HA cluster
- The VMware tools package included in the HCP OS with HCP-VM 7.0.1. 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 D.
- 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.
- Specialized HCP LUNs
 - **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 supports both a standard configuration and a small instance configuration. HCP-VM small instance configuration differs from the standard in that it requires only 4 vCPUs and only 16 GB of RAM. For the majority of use cases you will deploy your HCP-VM following the standard configuration guidelines. However, if your use case for HCP is not a resource intensive use case, you may wish to deploy your HCP-VM following the small instance configuration guidelines. An example use case for a small instance deployment is as follows:

- Up to 5 tenants
- Up to 25 namespaces
- A single active passive replication link
- Ingest duty cycle: up to 12 hours per day, 5 days per week

Other factors may impact whether the small instance deployment meets your performance requirements such as heavy MQE querying, or object / directory counts above published maximums, etc. If your small instance HCP-VM is not meeting your performance requirements it is recommended that you reconfigure vCPU and RAM according to the standard instance guidelines.

The following list is a composition of the prerequisite and recommended hardware for deploying an HCP-VM system:

- Minimum of 4 HCP-VM nodes in an HCP-VM system
- Minimum 8 vCPU allocated per HCP-VM node (allocated in OVF)
- Minimum 32GB RAM allocated per HCP-VM node (allocated in OVF)



Note: HCP does not recommend over committing RAM . Over committing RAM can degrade the performance of the HCP-VM system. If you still want to over commit RAM, see the applicable VMware documentation for best practices.

- Maximum 256GB RAM allocated per HCP-VM node (allocated in OVF)
- Shared SAN storage, RAID6 (Recommended)
- Minimum four 1.2TB LUNs allocated for default VMDK size deployment

- 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.”
- Datastores cannot be shared across HCP-VM nodes or other non-HCP-VM applications
- Two physical NICs on each ESXi host in the vSphere HA cluster dedicated for HCP-VM Back-end network (Recommended)
- Two physical NICs available for the VMware management network for vSphere HA (Recommended)
 - HCP-VM Front-end will also utilize these NICs
- Two port fibre channel HBA (or VMware compatible IO device) for shared storage connectivity (when applicable)
- ESXi requires a minimum of 2GB of physical RAM. VMware recommends providing at least 8GB 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 listed above with the following exceptions:

- Minimum 4 vCPU allocated per HCP-VM node (configured after OVF deployment)
- Minimum 16 GB RAM allocated per HCP-VM node (configured after OVF deployment)

HCP-VM system limits

The HCP-VM system is limited to the following requirements:

- 40 HCP-VM nodes
- 59 data LUNs per HCP-VM node (ESXi guest OS limitation)
- Two 500GB VMDKs minimum per HCP-VM node
- 3.66TB minimum usable per HCP-VM system
- Max open VMDK storage per host (ESXi Limitation)
 - 5.0 update 2: 60TB
 - 5.1 update 2: 60TB
 - 5.5: 128TB
 - 6.0: 128TB
- HCP-VM supports 2TB VMDK for 5.0 and 5.1
- HCP-VM supports 16TB VMDK for 5.5 and 6.0
- HCP supported limits can be found in both the customer and authorized HCP release notes.

HCP-VM Small Instance configuration has the same system limits as listed above but is limited to a maximum of 16 HCP-VM nodes instead of 40.

HCP-VM availability considerations

To ensure continuous availability of the HCP repository, $((n/2)+1)$ nodes must be running and healthy, where n represents the total number of storage nodes in the HCP system. In addition, the HCP system is only considered to be in a state of high availability if there is one HCP-VM node per ESXi host. If you have multiple HCP-VM nodes per ESXi host, you risk entering a state of metadata unavailability if any of your ESXi host fails.

The metadata unavailability state prohibits HCP namespaces from accepting write requests, including requests to store new data or change object metadata. Furthermore, the data stored in the affected nodes of your HCP system cannot be accessed until the HCP system repairs itself. The repair process can take between one and five minutes.

If your HCP system is running in accordance to HCP best practices, the HCP system can survive a single ESXi host failure without affecting HCP functionality.

It is also important to consider that:

- 1.** Zero Copy Failover is not available with HCP-VM. For a namespace with DPL 1, the loss of any single node will result in the data managed by that node being unavailable for read until that node is restored.
 - Data unavailability may be mitigated by replication to a second HCP cluster.
- 2.** ESXi hosts must not be oversubscribed on CPU, RAM or Disk because this can cause HCP system instability. It is expected that the ESXi administrator monitors resources to ensure the host is not oversubscribed.

To determine the physical sizing of the HCP-VM system, the end user must take into account HCP minimum healthy running node rules above as well as the physical limitations of the ESXi host. Using the information provided in this guide, the end user should take into account their site requirements for performance, availability, etc. HA is recommended, not required as it may not be needed, desired or possible in all environments. The best practice of a single HCP-VM node per ESXi host is recommended but not required. The end user should assess the needs of their user community and determine if they can achieve their agreed upon service levels in the event of a failure of an ESXi node or nodes without HA enabled.

Configuring the HCP-VM environment

This section will cover the steps required to provision the VMware environment to be ready for an HCP-VM deployment. These steps include the following:

- ESXi considerations
- Configuring vSphere HA cluster
- Configuring ESXi storage
- Configuring ESXi network

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 [Chapter 2: "Configuration guidelines for HCP-VM environment"](#).

All ESXi hosts that will contain an HCP-VM node must have Network Time Protocol (NTP) enabled. This is done from the vSphere client by clicking on Time Configuration under Software on the Configuration tab on each individual ESXi host.

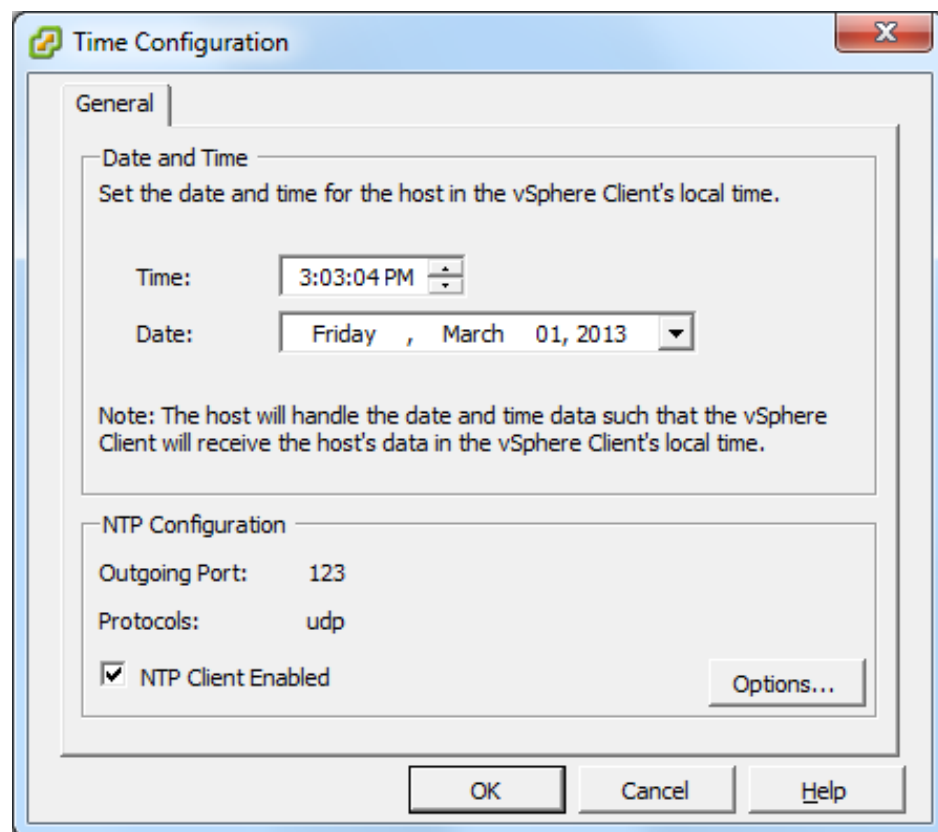


Important: NTP must be enabled for each ESXi host individually.

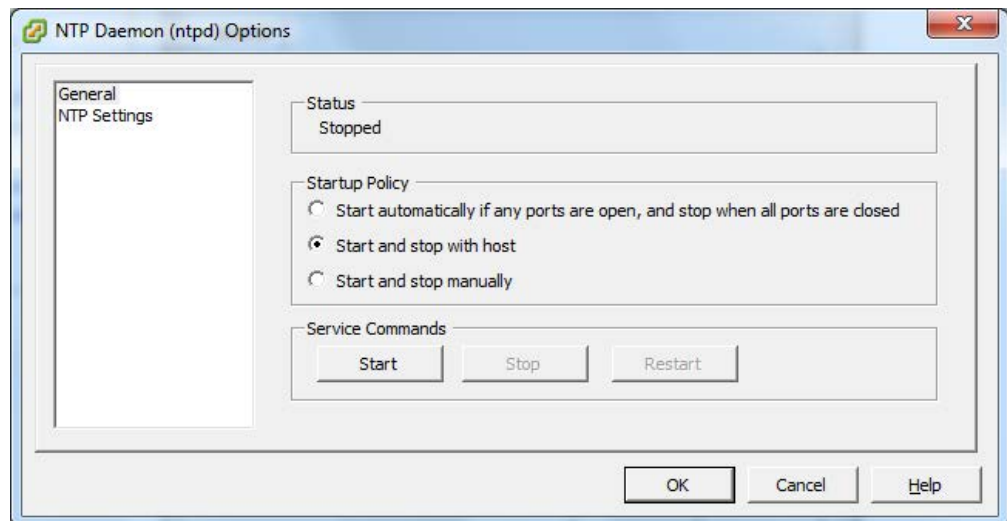
Enabling NTP for the ESXi hosts

To configure ESXi hosts for NTP:


1. Access your vSphere client.
2. In the left side navigation window, select the ESXi host for which you want to enable NTP.
3. In the right hand window, click on the **Configuration** tab.
4. Under the **Software** section in the right hand window, click **Time Configuration**.
5. In the upper right hand corner of the right hand window, click on **Properties**.
6. In the **Time Configuration** window, select the **NTP Client Enabled** check box.
7. Click on **Options**.



8. In the **NTP Daemon Options** window, select **Start and stop with host**.
9. In the left side navigation bar, click on **NTP Settings**.



10. In the **NTP Servers** section of the **NTP Daemon Options** window, click **Add** and enter the time server.
11. Select the **Restart NTP service to apply changes** checkbox.
12. Click **OK** and **OK** again in the **Time Configuration** window.



13. 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.

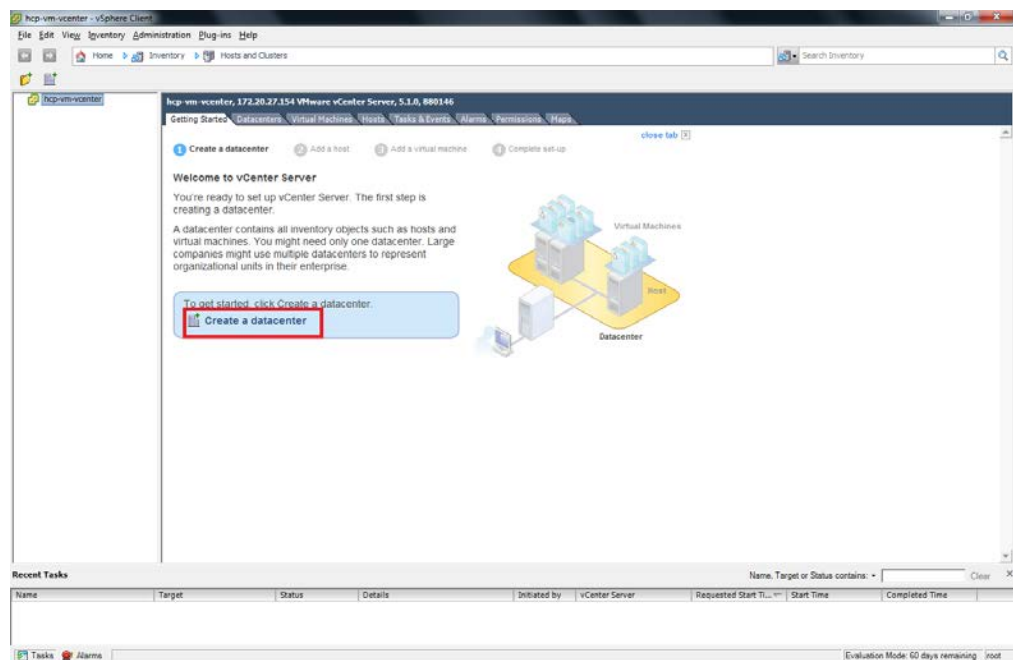
Configure vSphere HA cluster for HCP-VM (Recommended)

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.

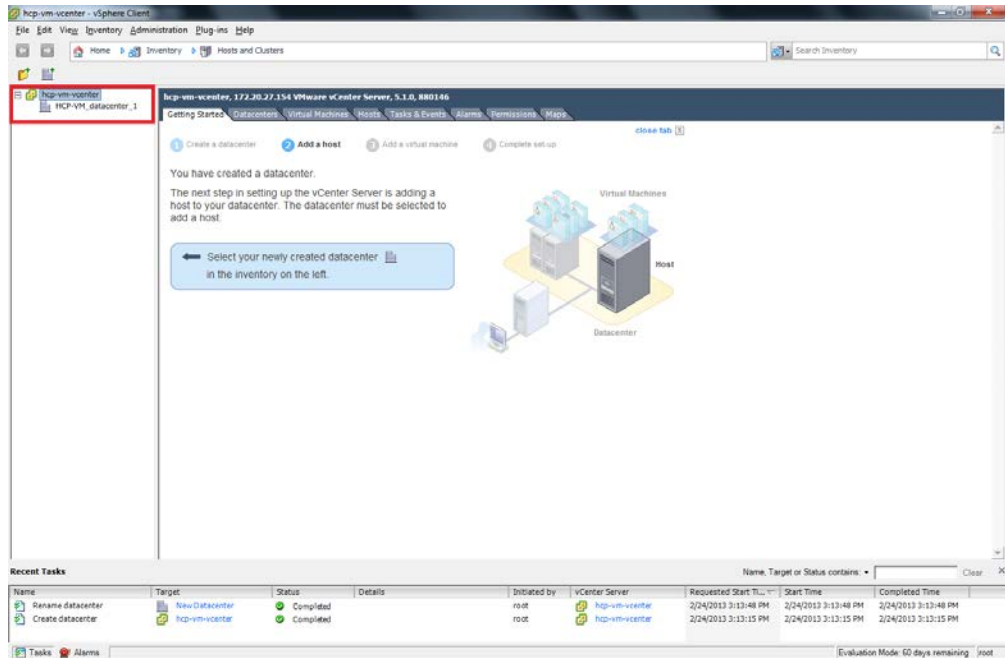
Step 1: Creating a data center

To create a datacenter:

1. Access the vSphere Client.
2. In the vSphere client, under the **Getting Started** tab, click on **Create a datacenter**.



3. In the left hand navigation bar, enter a name for your HCP-VM datacenter. Here is a good example name: HCP-VM_center_1.



Step 2: Add a cluster to the data center

To add a cluster:

1. In the **Getting Started** tab, click on **Create a cluster**. This will launch the **New Cluster Wizard**.
2. In the New Cluster Wizard, enter a name for the cluster. Here is a good example name: hcp-vm-cluster-1.

3. Select **Turn on vSphere HA**.

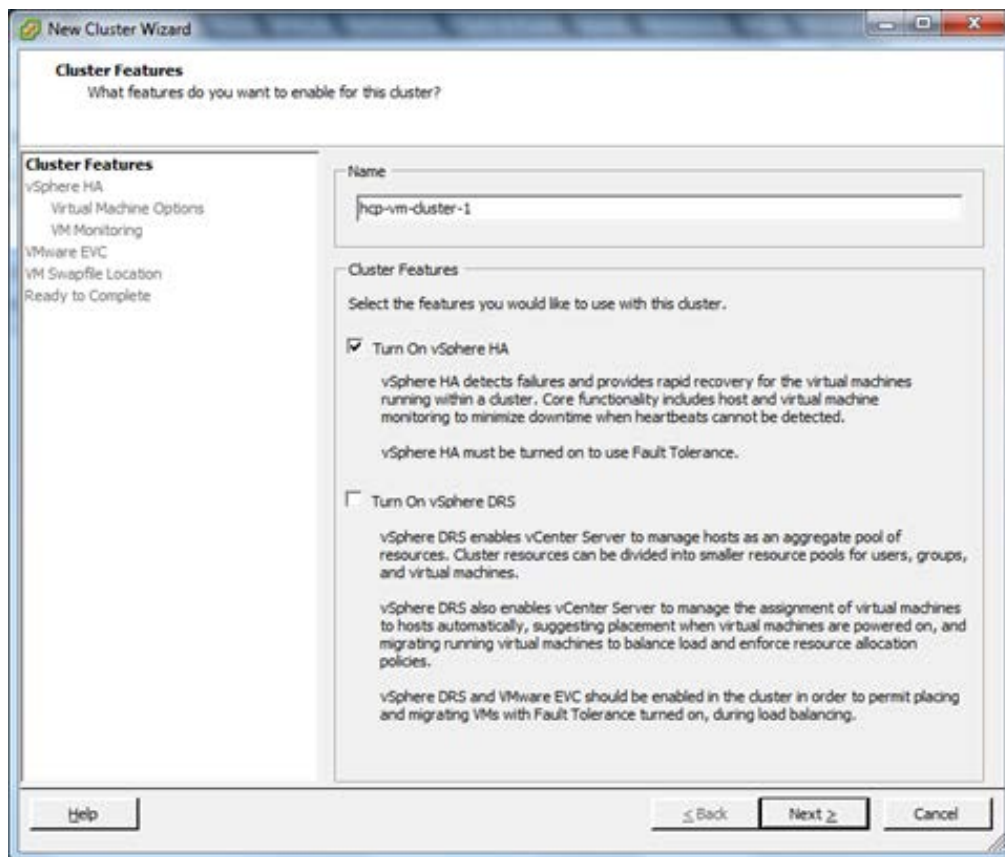


Important: Do **not** click Turn on vSphere DRS.

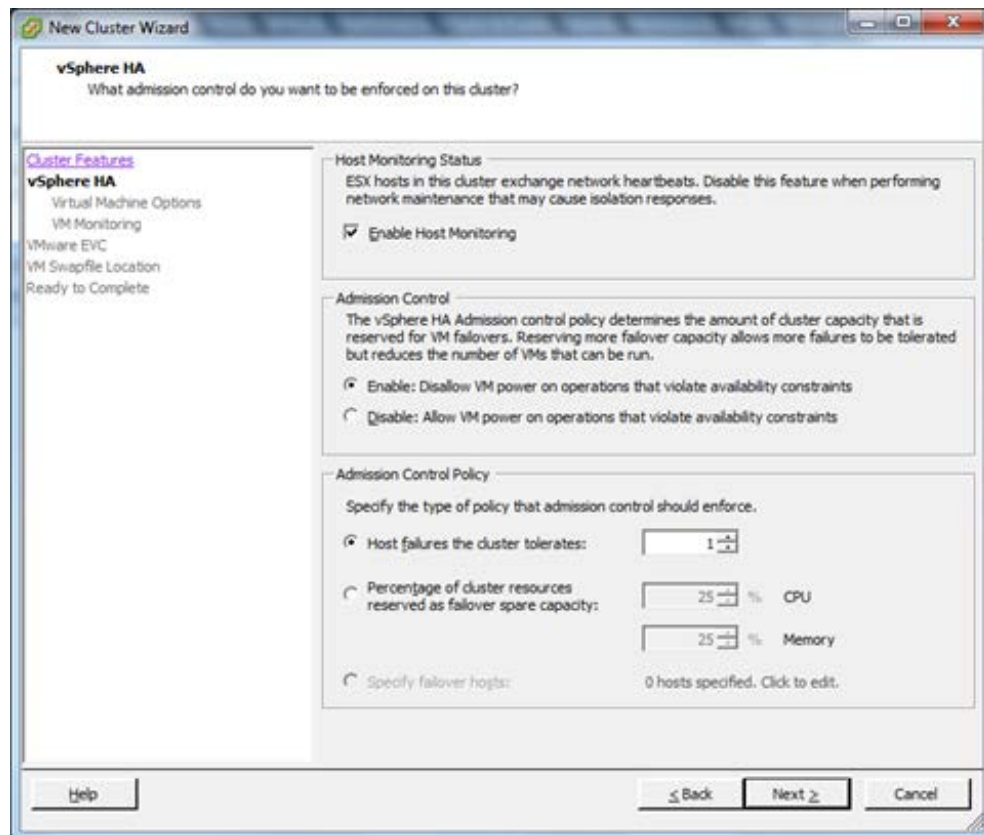


Note: DRS can be turned on later to define VM affinity to a particular host or group of hosts. This function does not provide further automation of failover. The settings described merely assist with keeping VMs on a particular host, and alert if the rule cannot be followed. See appendix D for details on the settings required.

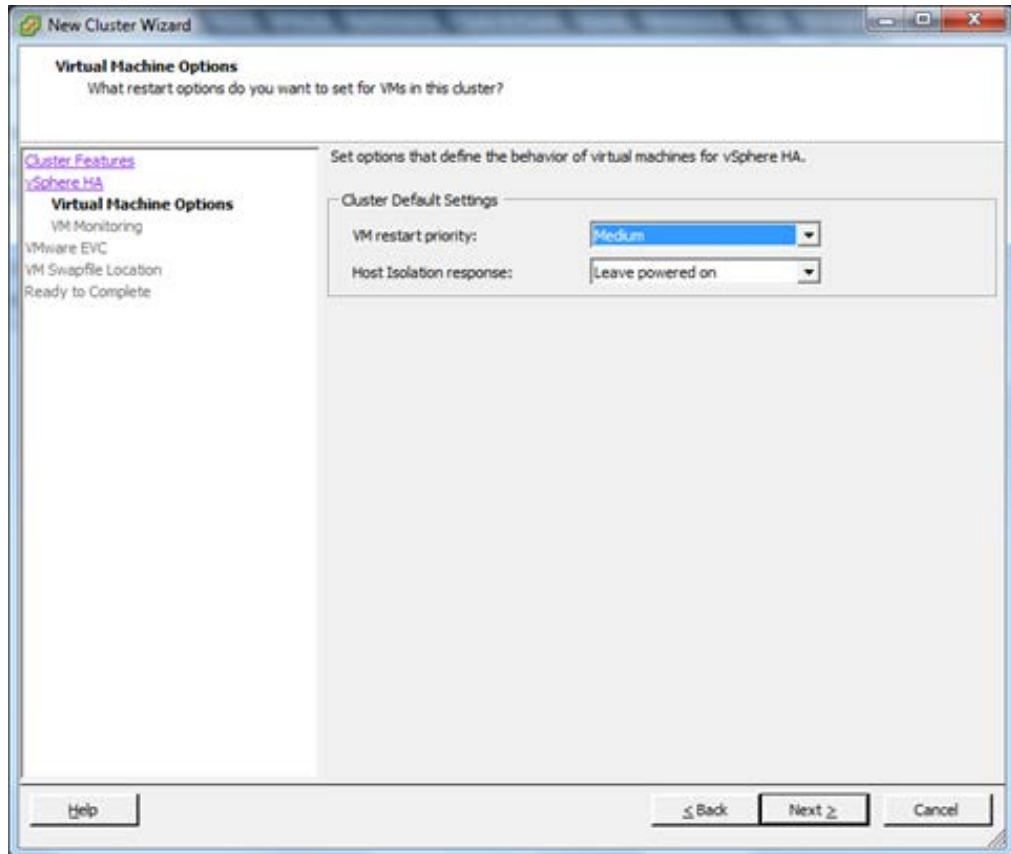
4. Click **Next**.



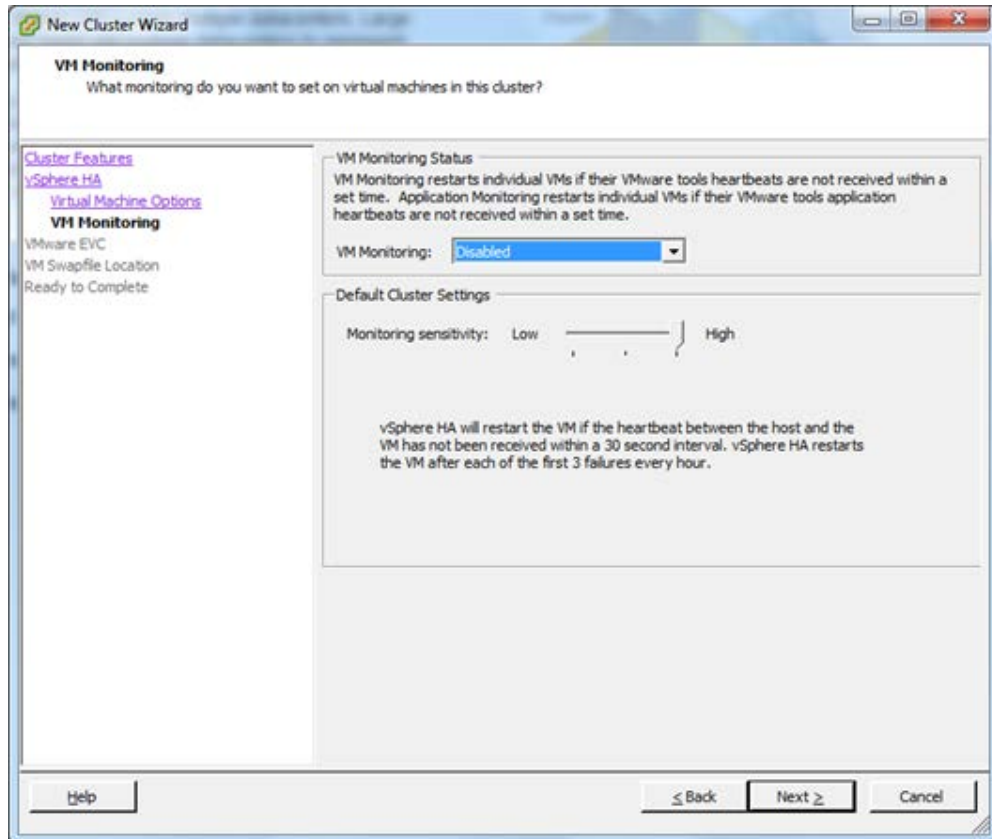
5. Select the **Enable Host Monitoring** checkbox.
6. Select **Enable: Disallow VM power on operations that violate availability constraints**.
7. Select **Host failures the cluster tolerates** and set the value to **1**.
8. Click **Next**.



9. Set the **VM restart priority** to **Medium**.
10. Set the **Host Isolation response** to **Leave powered on**.
11. Click **Next**.

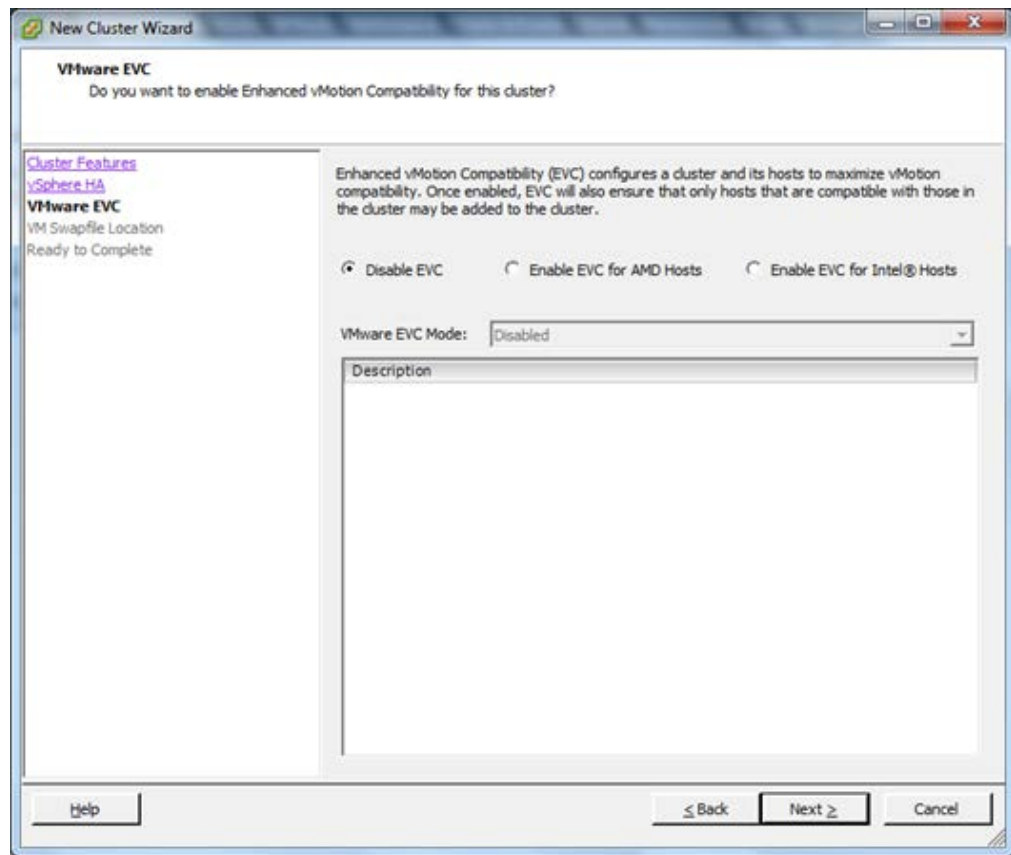


12. Set **VM Monitoring** to **Disabled**.
13. Drag the **Monitoring Sensitivity** pointer to **High**.
14. Click **Next**.



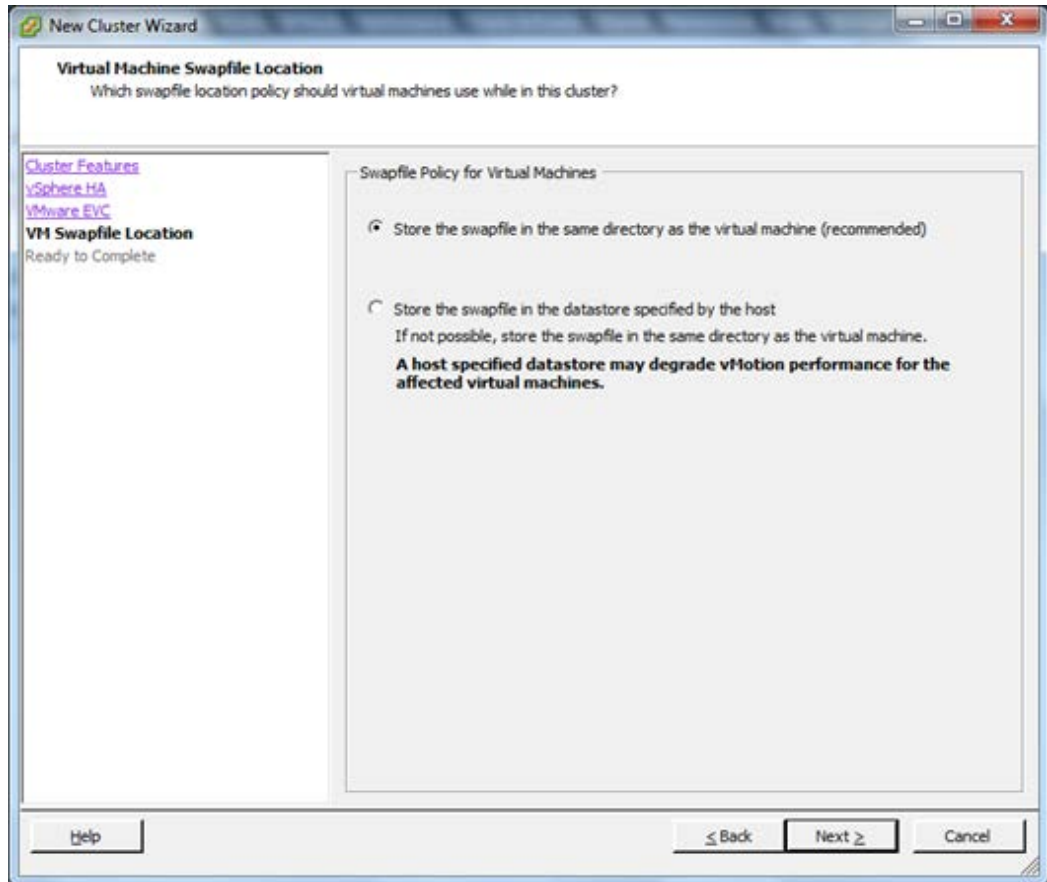
15. Select **Disable EVC**.

16. Click **Next**.

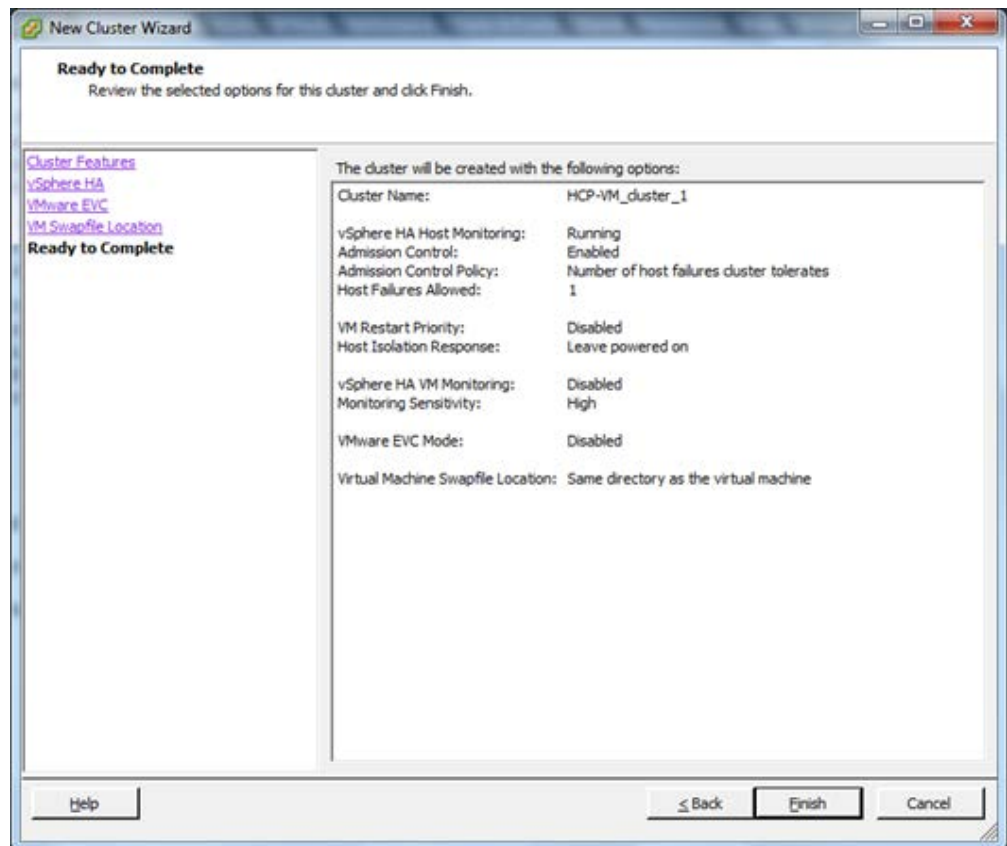


17. Select **Store the swap file in the same directory as the virtual machine (recommended)**.

18. Click **Next**.



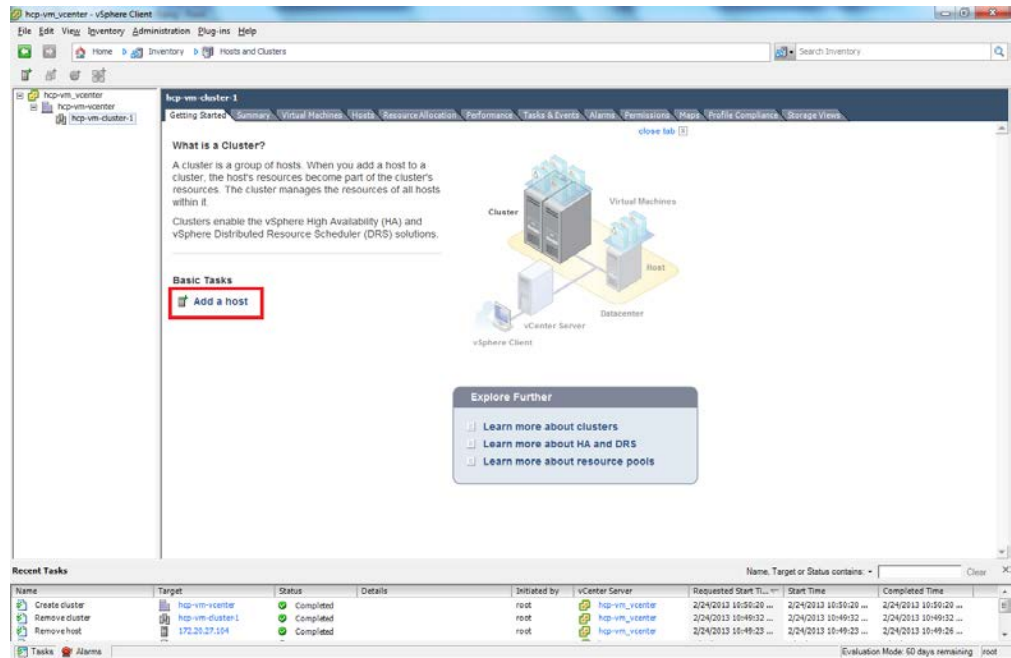
19. Review your preferences. Make sure they adhere to this manual.
20. Click **Finish** to create the new cluster.



Step 3: Add ESXi hosts to the HCP-VM cluster

To add ESXi hosts to the cluster:

1. On the vSphere Client home page, select the cluster you created on the left side navigation bar.
2. In the **Getting Started** tab, click on **Add a host**.

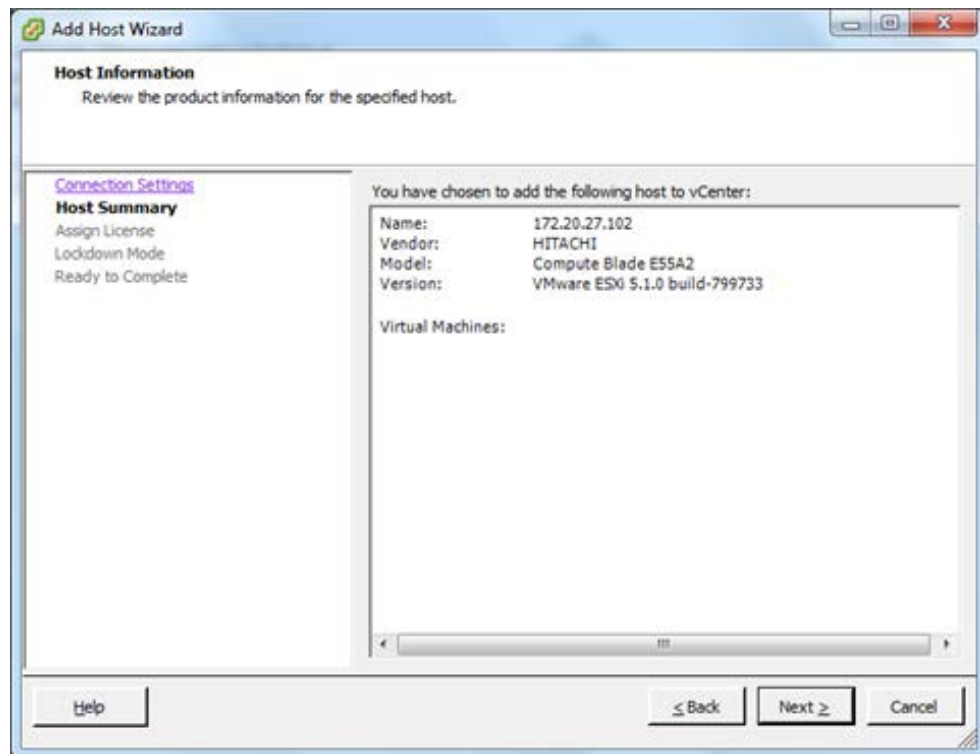


3. In the **Add Host Wizard**, enter the ESXi host connection information.
4. Enter the ESXi host Username and Password.
5. Click **Next**.

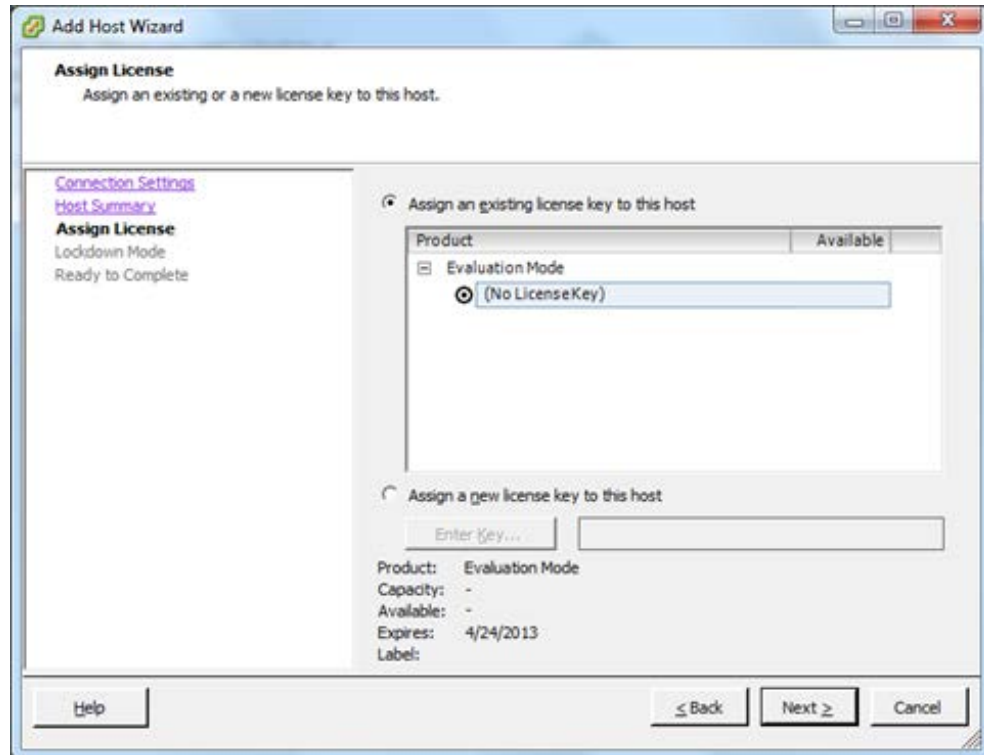
The screenshot shows the 'Add Host Wizard' window with the title bar 'Add Host Wizard'. The main content area is titled 'Specify Connection Settings' with the instruction 'Type in the information used to connect to this host.' On the left, under 'Connection Settings', there is a list of steps: 'Host Summary' (highlighted in blue), 'Assign License', 'Lockdown Mode', and 'Ready to Complete'. The main area is divided into two sections: 'Connection' and 'Authorization'. The 'Connection' section has the instruction 'Enter the name or IP address of the host to add to vCenter.' and a text box labeled 'Host:' containing the IP address '172.20.27.102'. The 'Authorization' section has the instruction 'Enter the administrative account information for the host. vSphere Client will use this information to connect to the host and establish a permanent account for its operations.' and two text boxes: 'Username:' containing 'root' and 'Password:' containing a masked password '*****'. At the bottom, there are four buttons: 'Help', '< Back', 'Next >', and 'Cancel'.

6. Review the Host Information.

7. Click **Next**.



8. Enter the license information for the ESXi host if it doesn't have any assigned.
9. Click **Next**.

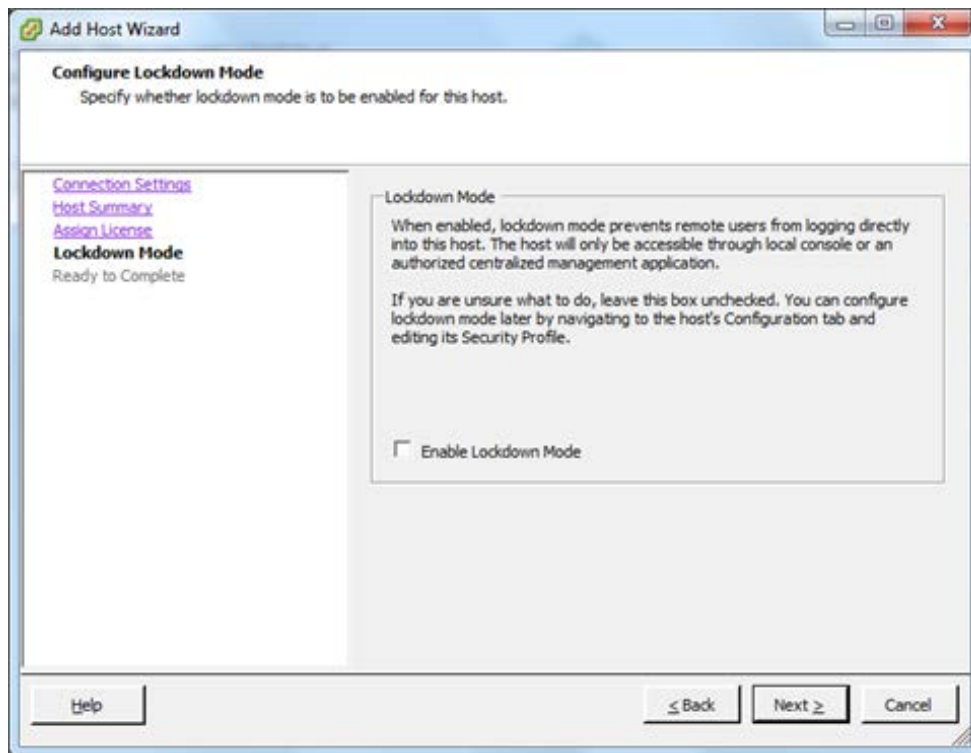


10. Select **Enable lock down mode** if you want to prevent remote users from logging in directly.

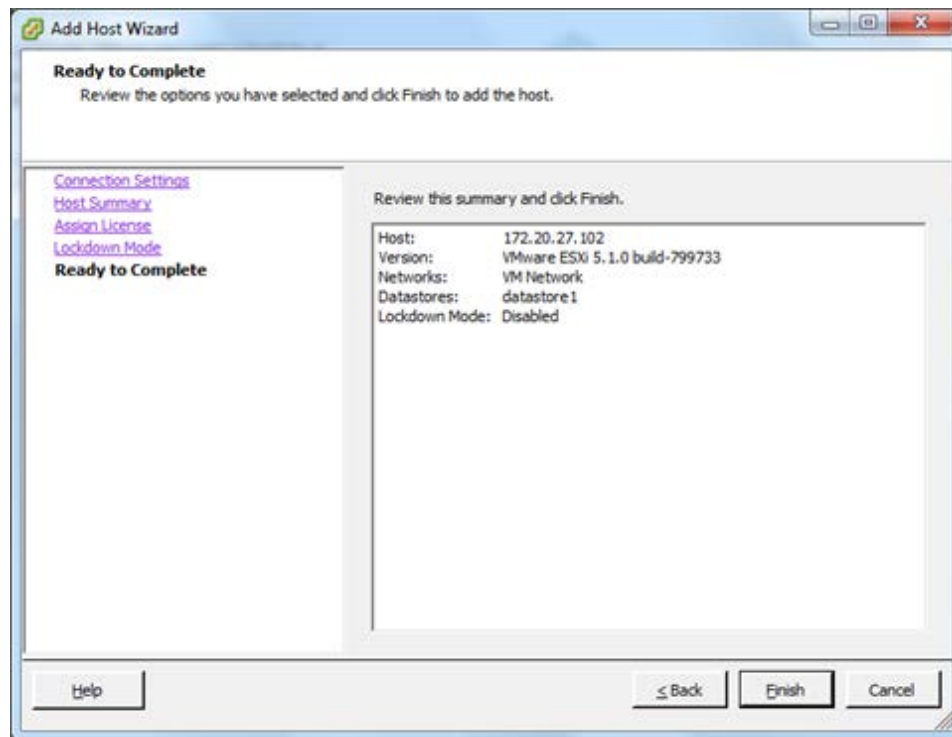


Note: The decision to implement lock down mode should be made by the customer.

11. Click **Next**.



12. Review your choices. Make sure they adhere to this guide.
13. Click **Finish** to add the ESXi host to the vSphere HA cluster.



14. Repeat Step 3 for all other ESXi hosts in the system.

After completing the configuration, it should be performed for all 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 (ex 6.0: vSphere Storage for ESXi 6.0 and vCenter Server 6.0) as well as the relevant storage vendor's VMware best practices guide.

It's 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:

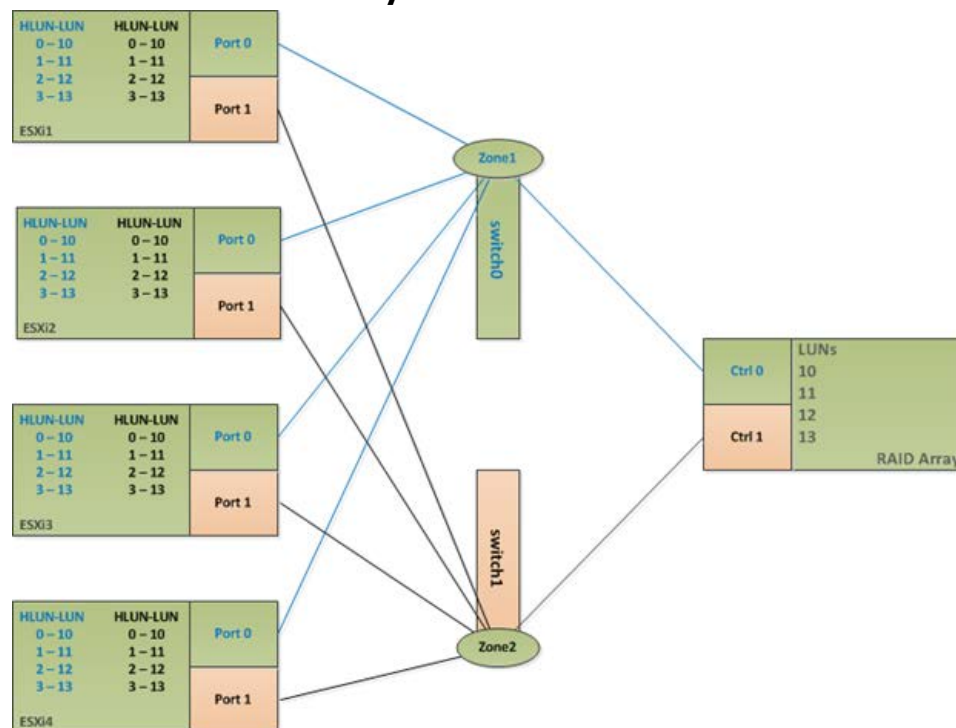
- Datastores used for HCP-VM nodes must be backed by shared RAID6 storage.
- Each datastore should only consist of one LUN.
- HCP-VM nodes cannot share datastores.
- All LUNs will be mapped to ESXi hosts in the vSphere HA cluster.
- All LUN IDs must be consistent across hosts. For example, LUN 1 should 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 HCP-VM VMDK OVF is configured with a 32GB OS LUN and two 500GB data LUNs.
 - Due to overhead (VMware, HCP system), you must configure 1.2 TB per LUN for each VMware datastore when using the default VMDK size in the VMDK OVF.
 - If the VMDK sizes included in the OVF need to be changed, refer to appendix C, [Appendix B: "Changing the VMDK target size"](#)

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



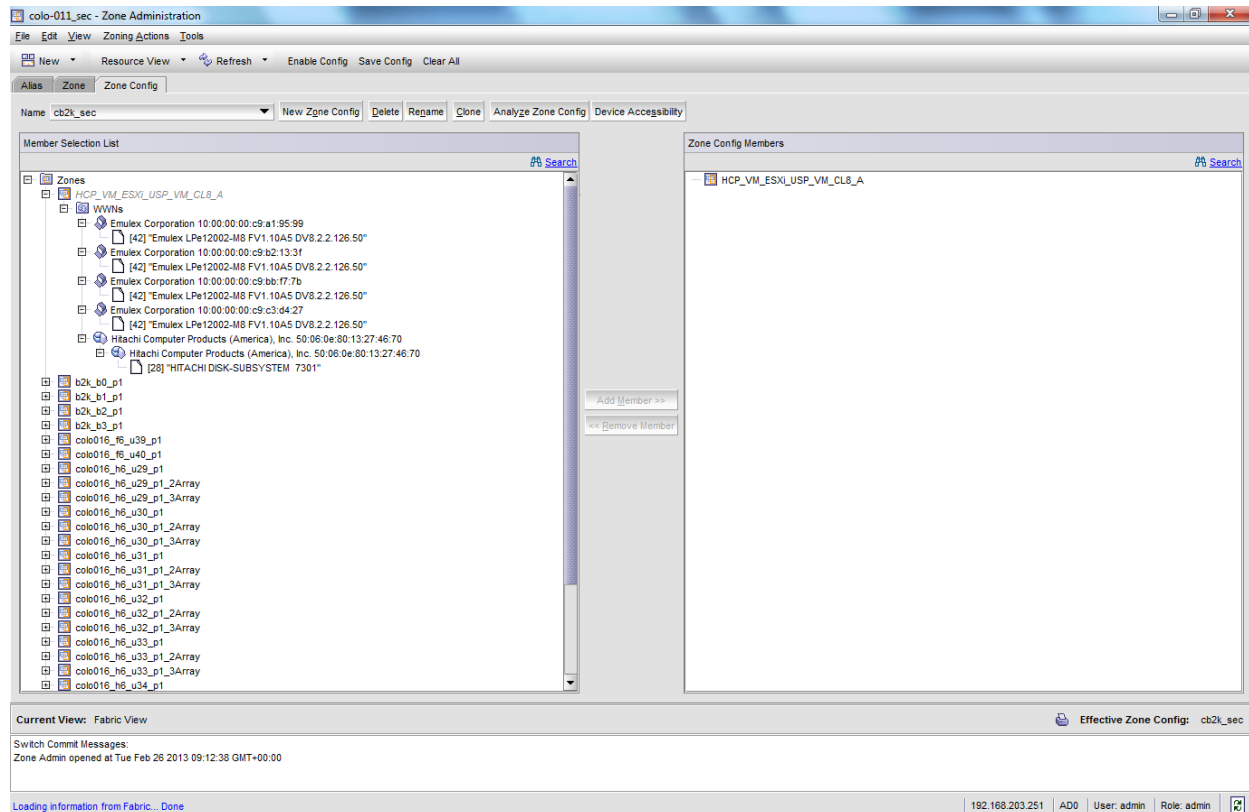
FC Switch 1, HCP-VM path 1

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

- Sample BrocadeZone containing four ESXi host WWNs, port 0 and a single array port on a HUS-VM.



- Sample HostGroup / LUN layout displaying the same LUNs mapped with the same HLUN to each ESXi host.
- This example assumes 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	1	10	hcp-vm_cluster-1_node_1_datastore_1
	ESXi-2	2	11	hcp-vm_cluster-1_node_2_datastore_1
	ESXi-3			
	ESXi-4	4	12	hcp-vm_cluster-1_node_3_datastore_1
		5	13	hcp-vm_cluster-1_node_4_datastore_1

Array path 2

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

This following image is an example of Storage Navigator view showing four datastores and LUN masking.



Note: Note that the same HLUN/LUN combination is assigned to all ESXi hosts.

Create LDEVs

1. Create LDEVs > 2. Select LDEVs > 3. Select Host Groups > **4. View/Change LUN Paths** > 5. Confirm

The LUN IDs are automatically set, but you can change a LUN by clicking Change LUN IDs. You must first select the check box for the host group (in the table subheading) you want to change, and select LDEVs you want to change and then click Change LUN IDs. Click Finish to confirm the LUN paths.

LUNs:

Added LUNs

Filter: ON OFF Select All Pages Options 1 / 1

LDEV ID	LDEV Name	Parity Group ID	Pool Name (ID)	Capacity	Provisioning Type	Attribute	LUN ID (8 Sets of Paths)							
							CL7- A/ESXi blade_0	CL7- A/ESXi blade_1	CL7- A/ESXi blade_2	CL7- A/ESXi blade_3	CL8- A/ESXi blade_0	CL8- A/ESXi blade_1	CL8- A/ESXi blade_2	CL8- A/ESXi blade_3
00:00:1C	HCP-VM_node_1_datastore_1	-	husVM_pool(0)	1228.80 GB	DP	-	1	1	1	1	1	1	1	1
00:00:1D	HCP-VM_node_2_datastore_1	-	husVM_pool(0)	1228.80 GB	DP	-	2	2	2	2	2	2	2	2
00:00:1F	HCP-VM_node_3_datastore_1	-	husVM_pool(0)	1228.80 GB	DP	-	4	4	4	4	4	4	4	4
00:00:30	HCP-VM_node_4_datastore_1	-	husVM_pool(0)	1228.80 GB	DP	-	5	5	5	5	5	5	5	5

Selected: 0 of 4

Change LDEV Settings Change LUN IDs

Back Next Finish Cancel Help

Add datastores to vSphere HA cluster

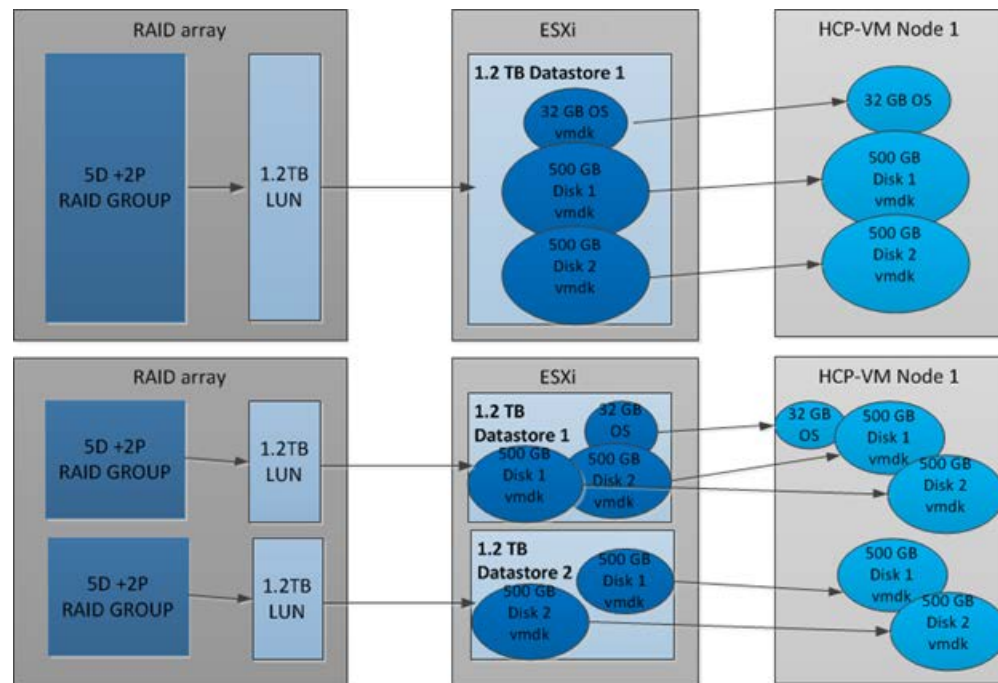
It is recommended to have only one LUN from a RAID Group in the HCP-VM system. Adding multiple LUNs from the same RAID Group increases the risk of data loss in the event of a failure.

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

During the initial OVF deploy, three VMDK's will be created from the initial datastore space. One 32GB OS LUN, and two 500GB data LUNs.

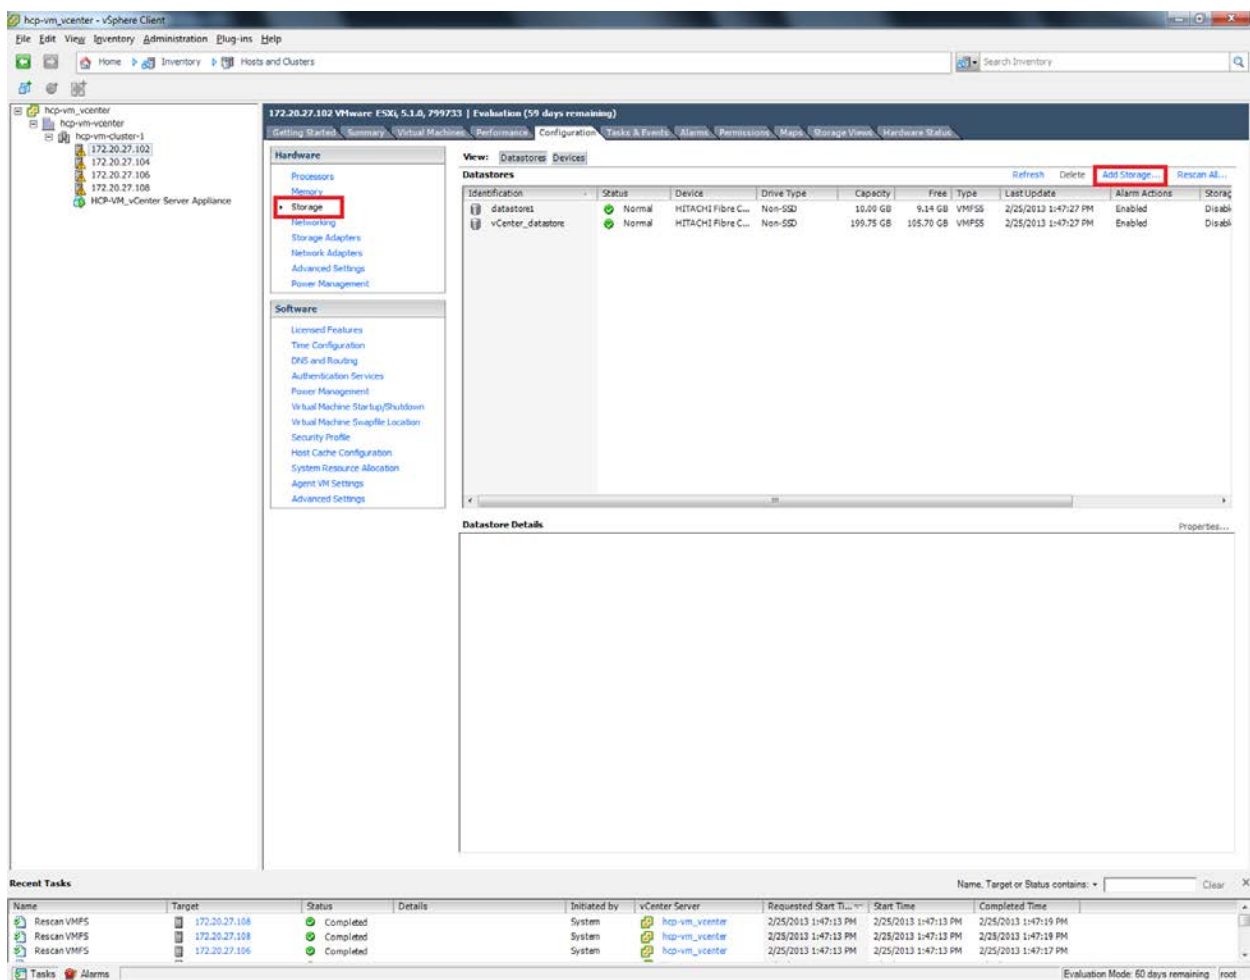
Currently, 2TB is the largest a disk can be in an HCP-VM system using VMware 5.0 and 5.1 VMDKs. The largest a disk can be in an HCP-VM system using VMware 5.5 and 6.0 is 16TB.

Here is a visual depiction of the cluster layout.

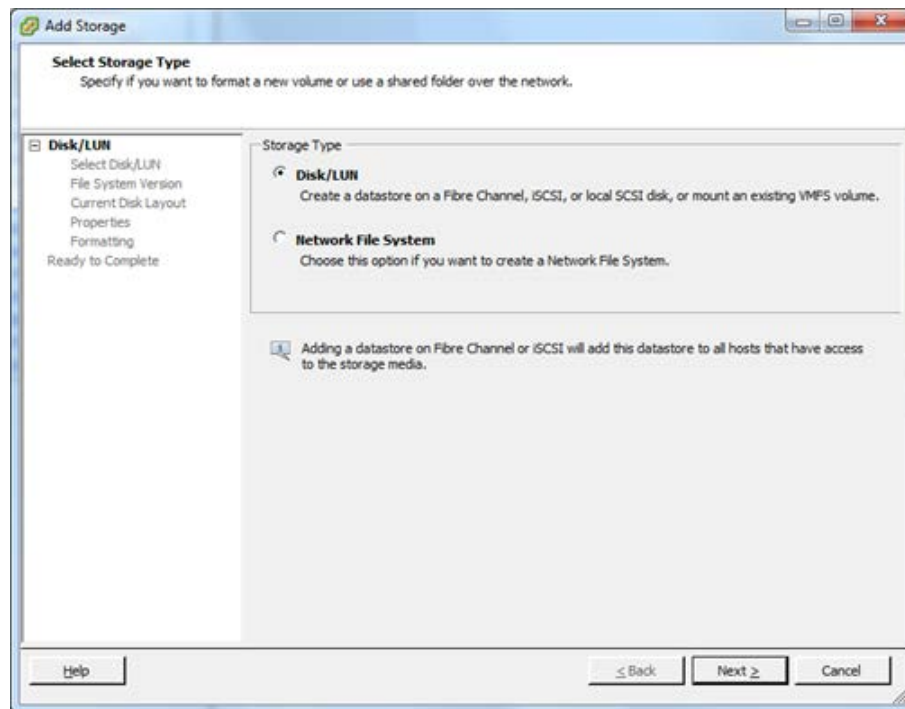


To add datastores to vSphere HA clusters:

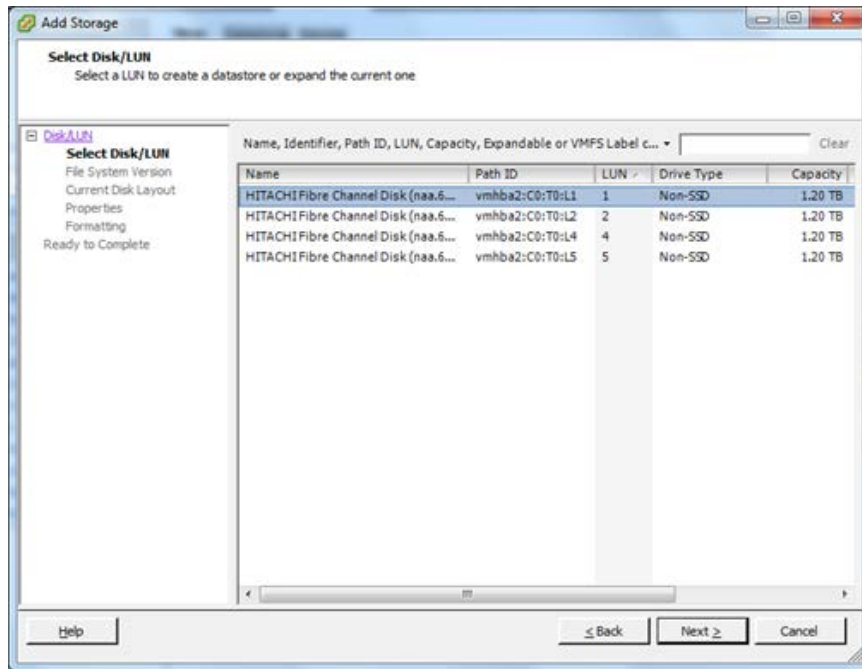
1. Access your vSphere Client.
2. In the left side navigation bar, click on the top ESXi host in your HCP-VM cluster.
3. In the right side window, click on the **Configuration** tab.
4. Click on **Storage** under the **Hardware** section.
5. In the **Datastores** section, click on **Add Storage**, located at the top right of the window.



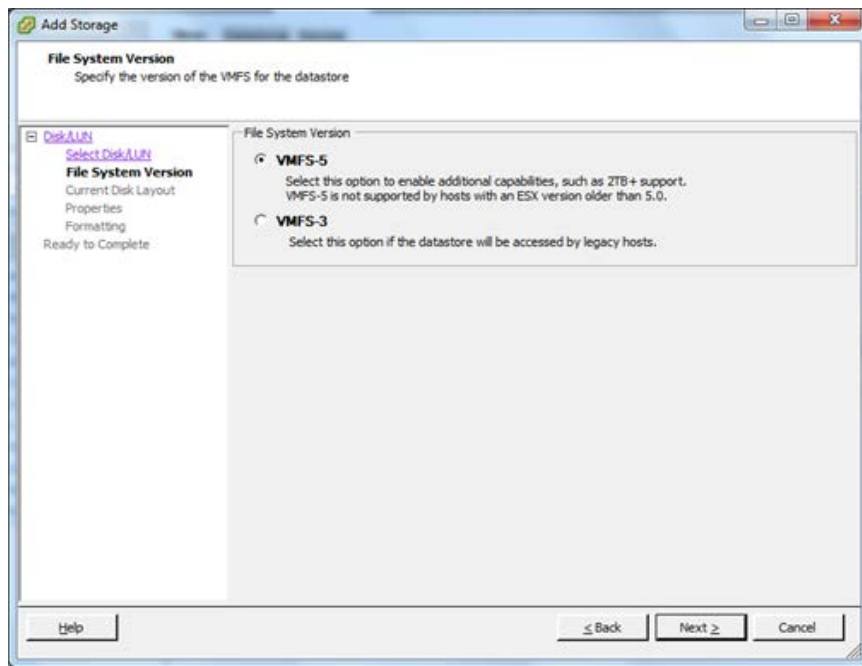
6. In the **Add Storage** window, select **Disk/LUN**.
7. Click **Next**.



8. Select the appropriate LUN in the list.
9. Click **Next**.

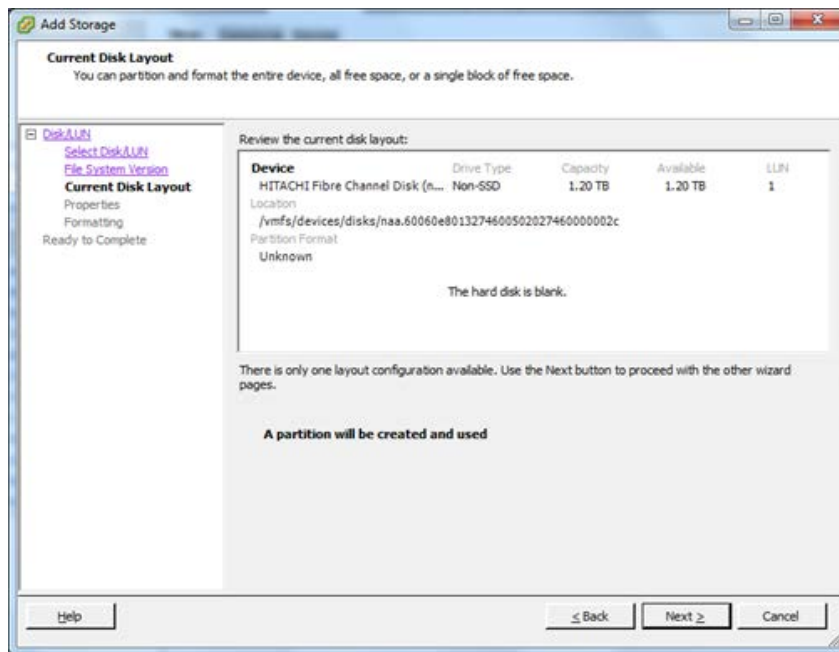


10. Select **VMFS-5**.
11. Click **Next**.



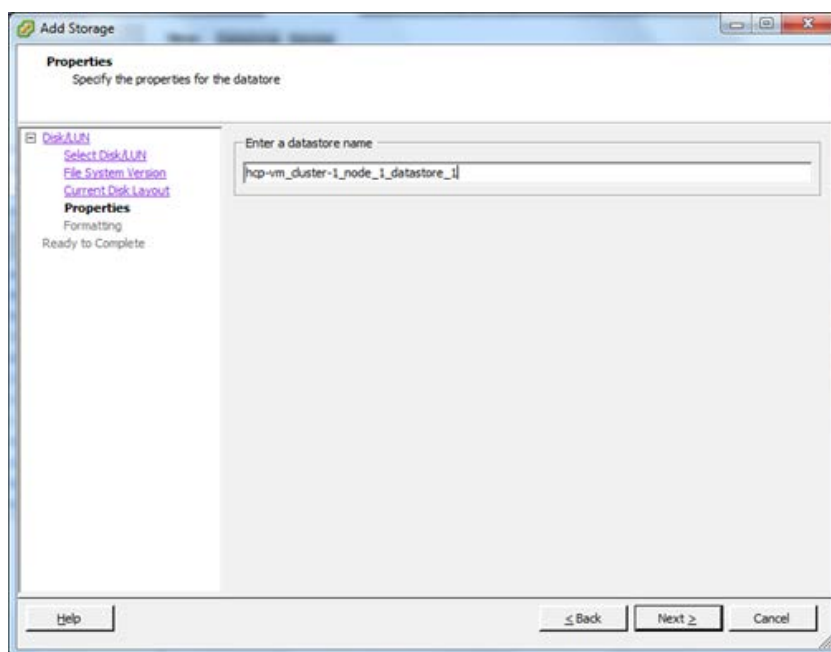
12. Review the the **Current Disk Layout** information.

13. Click **Next**.



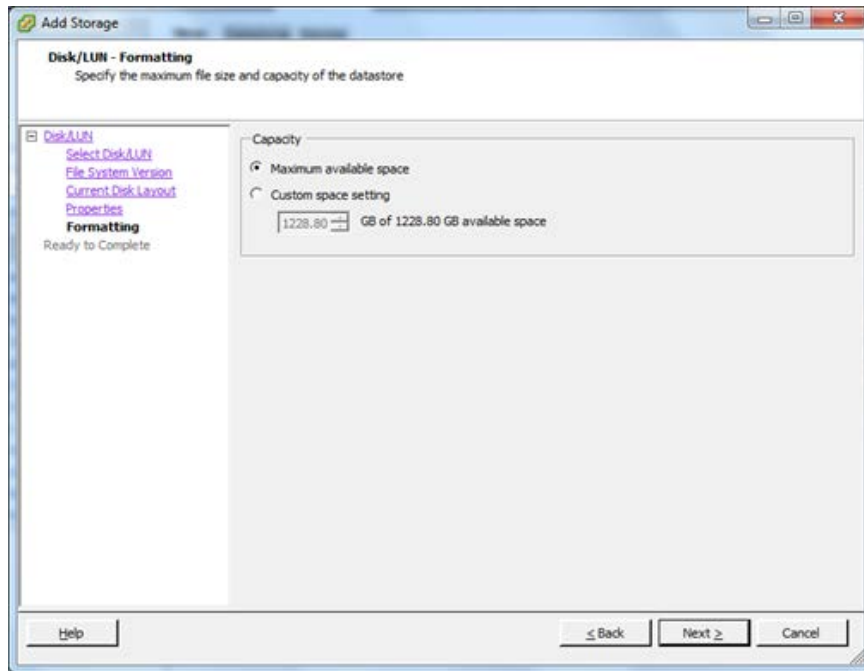
14. Enter a meaningful name for the datastore. A good example name is:
`hcp-vm_cluster_1_node_1_datastore_1`.

15. Click **Next**.



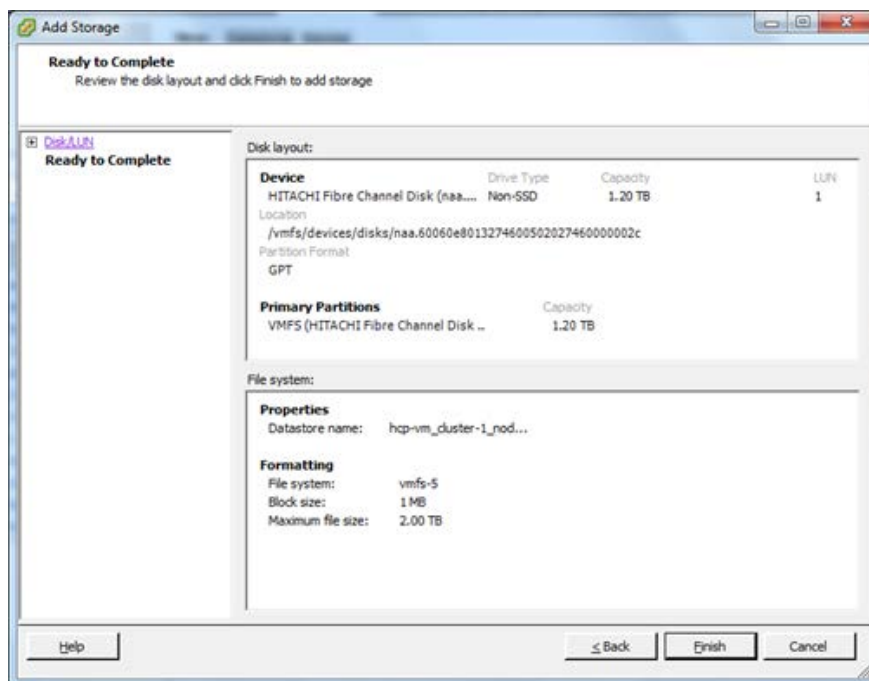
16. Select **Maximum available space**.

17. Click **Next**.



18. Review the **Disk layout** and **File System** information.

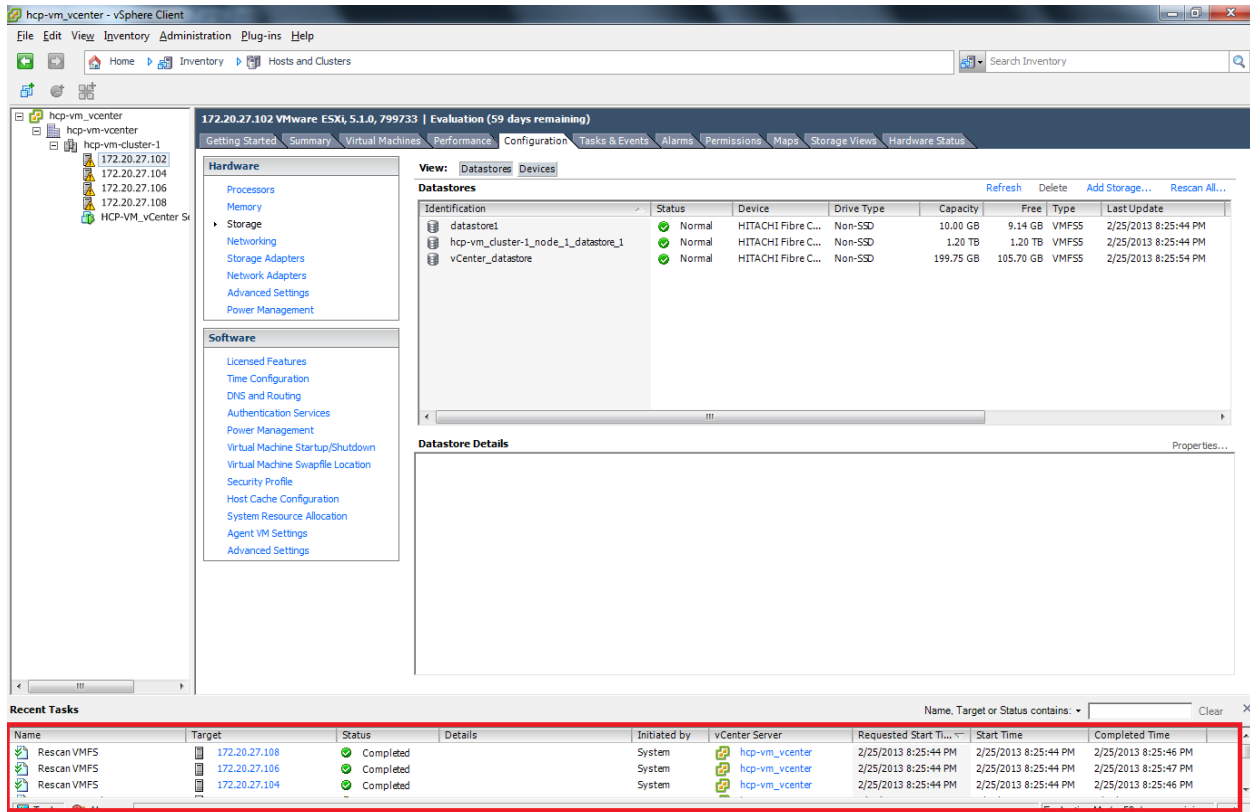
19. Click **Finish** to create the datastore.



Add datastores to vSphere HA cluster

The datastore should now be initialized and mounted. If it is, then in the **Recent Tasks** section, at the bottom of the vSphere Client, a **Rescan VMFS** alarm should be issued for all other ESXi hosts in the cluster.

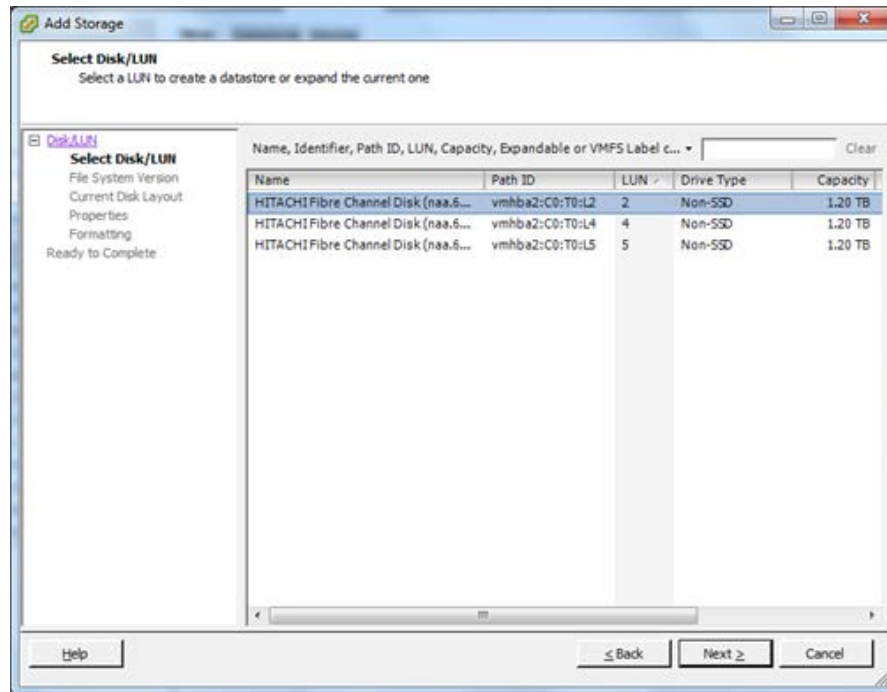
The new datastore should be automatically added to the inventory of all the other ESXi hosts.



Repeat the adding storage procedure for the 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
- LUN4 = hcp-vm-cluster_1_node_3_datastore_1
- LUN5 = hcp-vm-cluster_1_node_4_datastore_1



Once everything is completed, select the ESXi host and go to the **Configuration** tab. Click on the **Storage** under the **Hardware** section. Each ESXi host should appear with all datastores tagged with a normal status.

Add datastores to vSphere HA cluster

The screenshot shows the vSphere Client interface for a vSphere HA cluster. The 'Storage' tab is selected under the 'Hardware' section. The 'Datastores' table lists several datastores, including 'datastore1' and 'hcp-vm_cluster-1_node_3_datastore_1'. The 'Datastore Details' section for 'hcp-vm_cluster-1_node_3_datastore_1' shows its location, capacity, and storage capabilities. The 'Recent Tasks' section at the bottom shows a list of tasks, including 'Rename datastore' and 'Rescan VMFS'.

Identification	Status	Device	Drive Type	Capacity	Free	Type	Last Update
datastore1	Normal	HITACHI Fibre C...	Non-SSD	10.00 GB	9.14 GB	VMFSS	2/25/2013 8:34:52 PM
hcp-vm_cluster-1_node_1_datastore_1	Normal	HITACHI Fibre C...	Non-SSD	1.20 TB	1.20 TB	VMFSS	2/25/2013 8:34:56 PM
hcp-vm_cluster-1_node_2_datastore_1	Normal	HITACHI Fibre C...	Non-SSD	1.20 TB	1.20 TB	VMFSS	2/25/2013 8:34:56 PM
hcp-vm_cluster-1_node_3_datastore_1	Normal	HITACHI Fibre C...	Non-SSD	1.20 TB	1.20 TB	VMFSS	2/25/2013 8:34:56 PM
hcp-vm_cluster-1_node_4_datastore_1	Normal	HITACHI Fibre C...	Non-SSD	1.20 TB	1.20 TB	VMFSS	2/25/2013 8:34:56 PM
vCenter_datastore	Normal	HITACHI Fibre C...	Non-SSD	199.75 GB	105.70 GB	VMFSS	2/25/2013 8:34:56 PM

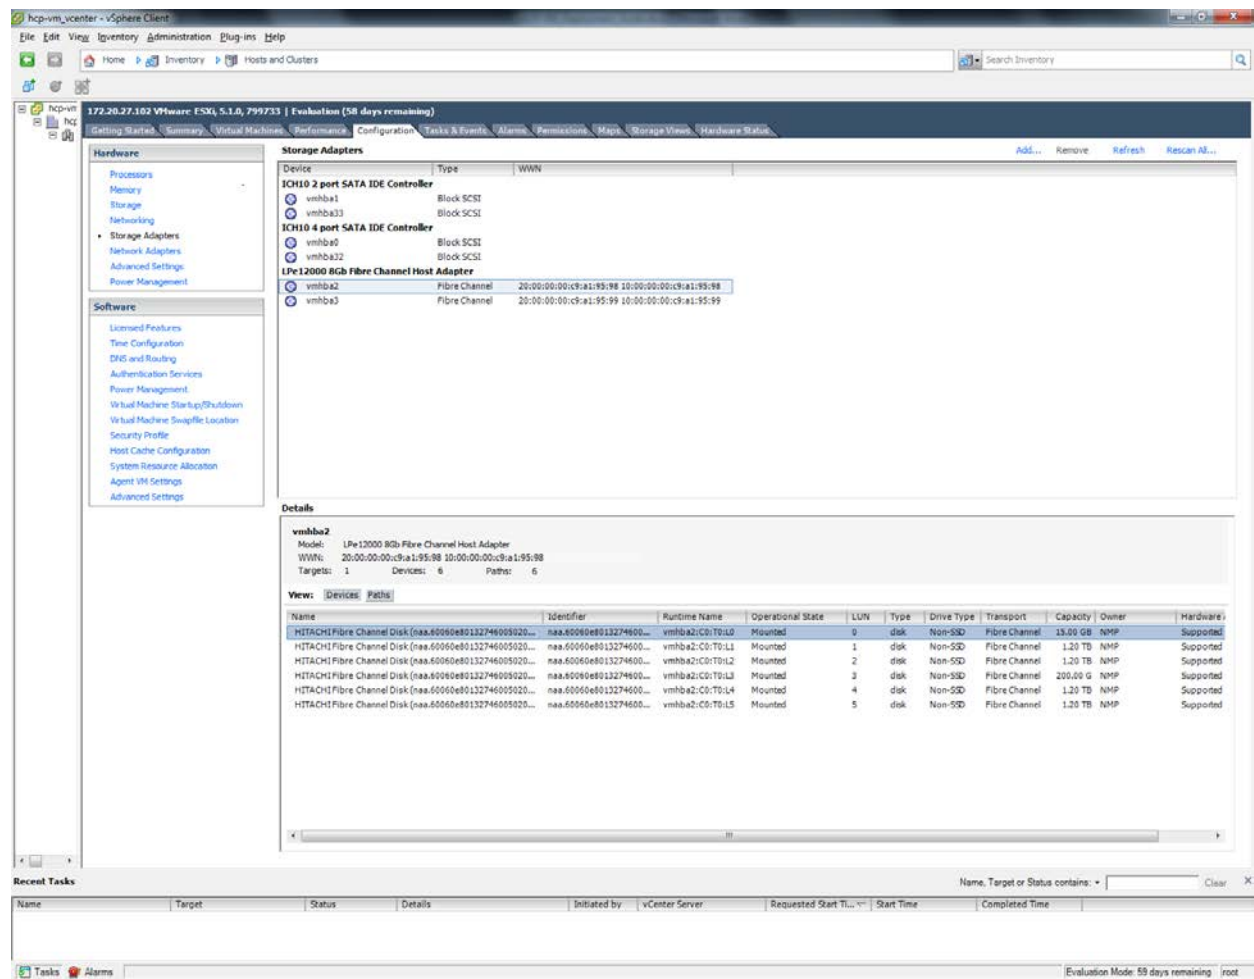
Path Selection	Properties	Extents	Storage I/O Control
Fixed (VMware)	Volume Label: hcp-vm_clus... Datastore Name: hcp-vm_clus...	HITACHI Fibre Channel Disk... Total Formatted Capacity: 1.20 TB	Disabled

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Rename datastore	hcp-vm_cluster-1_node...	Completed		root	hcp-vm_vcenter	2/25/2013 8:34:56 PM	2/25/2013 8:34:56 PM	2/25/2013 8:34:57 PM
Rescan VMFS	172.20.27.108	Completed		System	hcp-vm_vcenter	2/25/2013 8:34:22 PM	2/25/2013 8:34:22 PM	2/25/2013 8:34:24 PM
Rescan VMFS	172.20.27.106	Completed		System	hcp-vm_vcenter	2/25/2013 8:34:22 PM	2/25/2013 8:34:22 PM	2/25/2013 8:34:25 PM

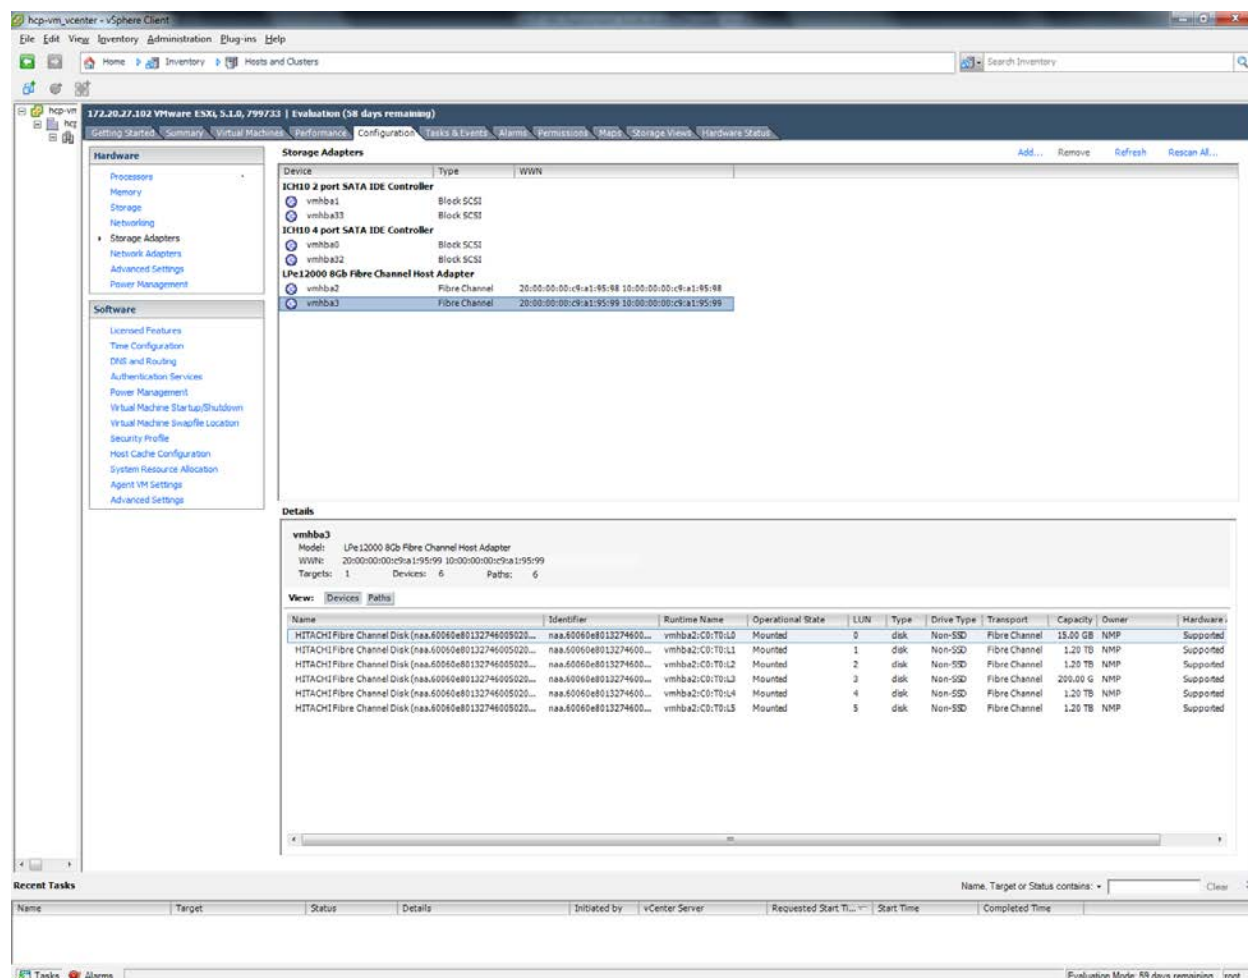
Alert should no longer appear for each ESXi node because there are now two datastores available for heartbeating.

Next, click on the **Configuration** tab, and click on **Storage Adapters** under the **Hardware** section. Make sure that the **Operational State** is **Mounted** for both paths.

Path 1



Path 2



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 VMware Virtual Infrastructure client.
2. In the left side navigation window, select an ESXi host.
3. In the right hand window, click on the **Configuration** tab.
4. Under the **Hardware** section in the right hand window, click **Storage**.
5. In the upper right hand corner of the right hand window, click on **Add Storage**(SCSI, SAN, and NFS).
6. In the **Storage Type window**, select the **Network File System** storage type.
7. Click **Next**.
8. In the **Locate Network File System** window, enter the NAS server name, the folder, and the datastore name,
9. Click **Next**.
10. Review your set up and click **Finish**.



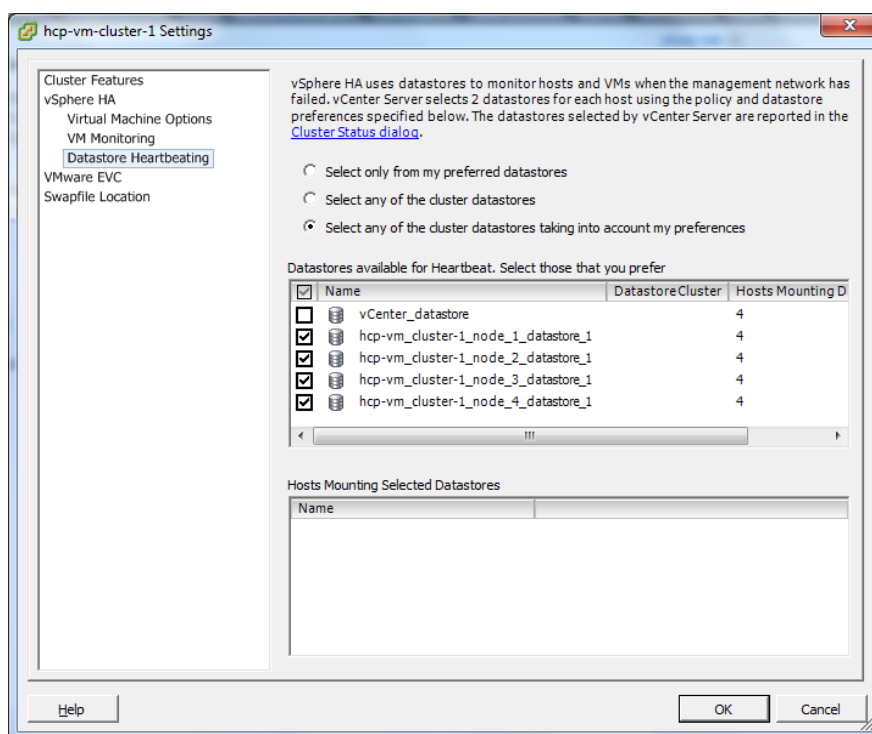
Important: Ensure that you mount datastores with the same volume label on all vSphere ESXi hosts within VMware high availability (HA) environments.

Heartbeat datastore selection

The Heartbeat Datastore function monitors hosts and Virtual Machines if the management network fails.

To activate Heartbeat datastore:

1. Access your vSphere Client.
2. On the left side navigation bar, right click on the cluster and in the sub menu click **Edit Settings**.
3. In the **Settings** window, select **Datastore Heartbeating** from the left side navigation bar.
4. Select four HCP-VM datastores.
5. Enable the option to **Select any of the cluster datastores** and mimic the preferences shown in the image below.
6. Click **OK** to commit the settings.



Preparing the ESXi network for the HCP-VM OVF deployment

For optimal performance, security, and high availability of an HCP-VM system, it is recommended to provide exclusive use of two physical NICs per node. These are used for private, Back-end communication within the

system. The Back-end network is responsible for such things as HCP **Heartbeating** and data traffic.

The Back-end NICs should be connected to dedicated, redundant Ethernet switches with spanning tree disabled and multi-cast enabled. Multi-cast should be configured for its vendor's specifications.

If the HCP-VM system is going to utilize dvSwitches, consult relevant VMware and vendor documentation for best practices.

Each HCP-VM node should have a least one physical NIC used for data access and system management. If utilizing 802.3ad in the customer environment, plan accordingly and follow VMware's best practices for configuration.

If the HCP-VM system is going to be used with the virtual network management feature, follow the guide in appendix B.

If the HCP-VM system will use NFS datastores, be sure to add the VM Kernel device for IP networking. Consult VMware documentation for more details on configuring ESXi with NFS datastores.

Configuring the Storage Network (HNAS Best Practice)

The IP protocol storage uses the **TCP/IP stack** as its foundation for communication. The stack includes Internet Small Computer System Interface **iSCSI** and Network Access Server **NAS** for ESXi hosts. A VMkernel uses the TCP/IP protocol stack to handle the data transport. Make sure the NFS server is enabled on all ESXi hosts.

To create a VMkernel:

1. Access the vSphere client.
2. On the left side navigation bar, select an ESXi host.
3. Click on the **Configuration** tab in the right side window.
4. In the **Hardware** section, click on **Networking**.
5. In the top right quadrant of the right side window, click on **Add Networking**.
6. In the **Add Network Wizard** window, select **VMkernel**.
7. Click **Next**.

8. Select one of the **Physical Network Cards**.
9. Click **Next**.
10. In the **Network Label** text box, enter VMkernel.
11. Click **Next**.
12. Enter the IP address and the subnet mask.
13. To provide the VMkernel default gateway, click **Edit** and enter the gateway address.
14. Click **OK**.
15. Back in the Wizard, click **Next**.
16. Click **Finish**.



Note:

- If using large 2TB NFS datastores, increase RPC timeout.
- Hitachi Vantara recommends that the VMkernel network be set up in a private network or with a unique VLAN ID that provides network isolation. For a full list of Hitachi Vantara recommendations for HNAS NFS datastores, review Hitachi NAS Platform Best Practices Guide for NFS with VMware vSphere.

Configuring networking for Front-end switching

The HCP Front-end network needs to be configured so that it can perform system management and provide client access. You are responsible for configuring the network.

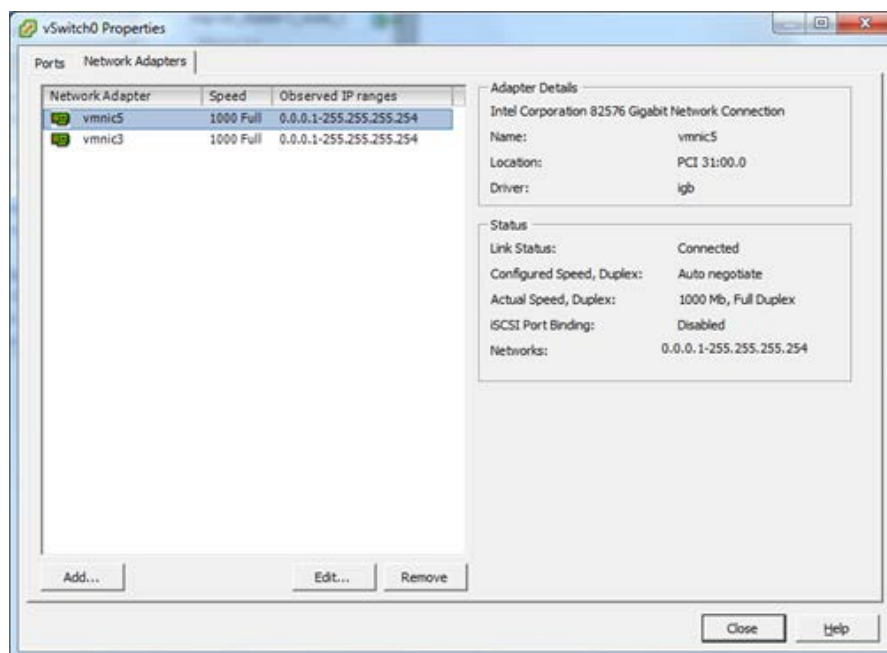
To configure the front end network:

1. Access the vSphere client.
2. In the left side navigation bar, select the first ESXi host.
3. In the right side window, click on the **Configuration** tab.
4. Click on **Networking** in the **Hardware** section.
5. Click **Properties** button located in the center of the right hand window.

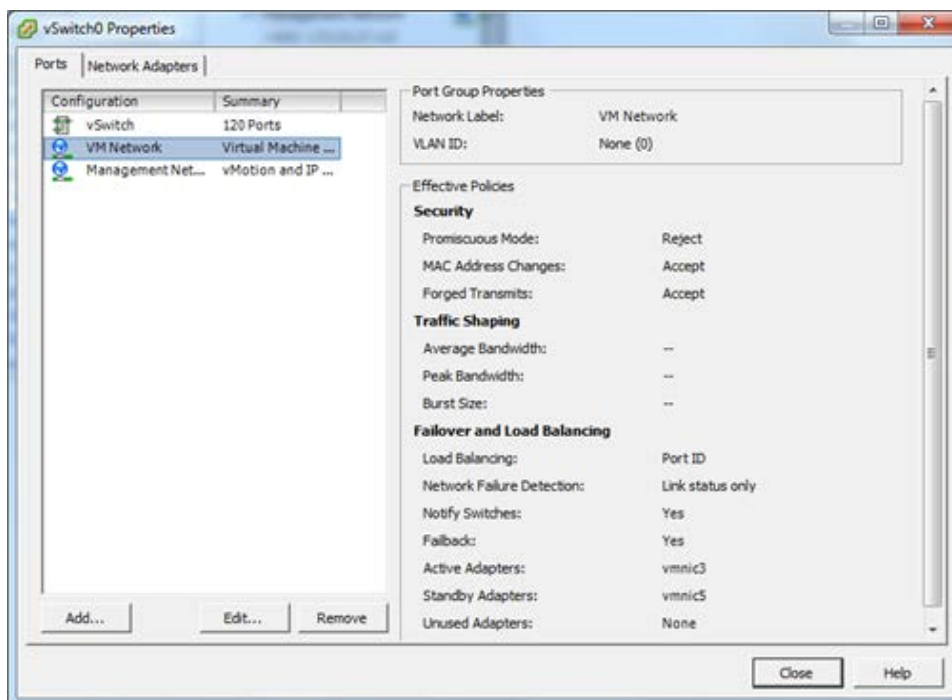


Note: There are multiple property buttons on the page. Make sure to click the right one or you will not open the appropriate window.

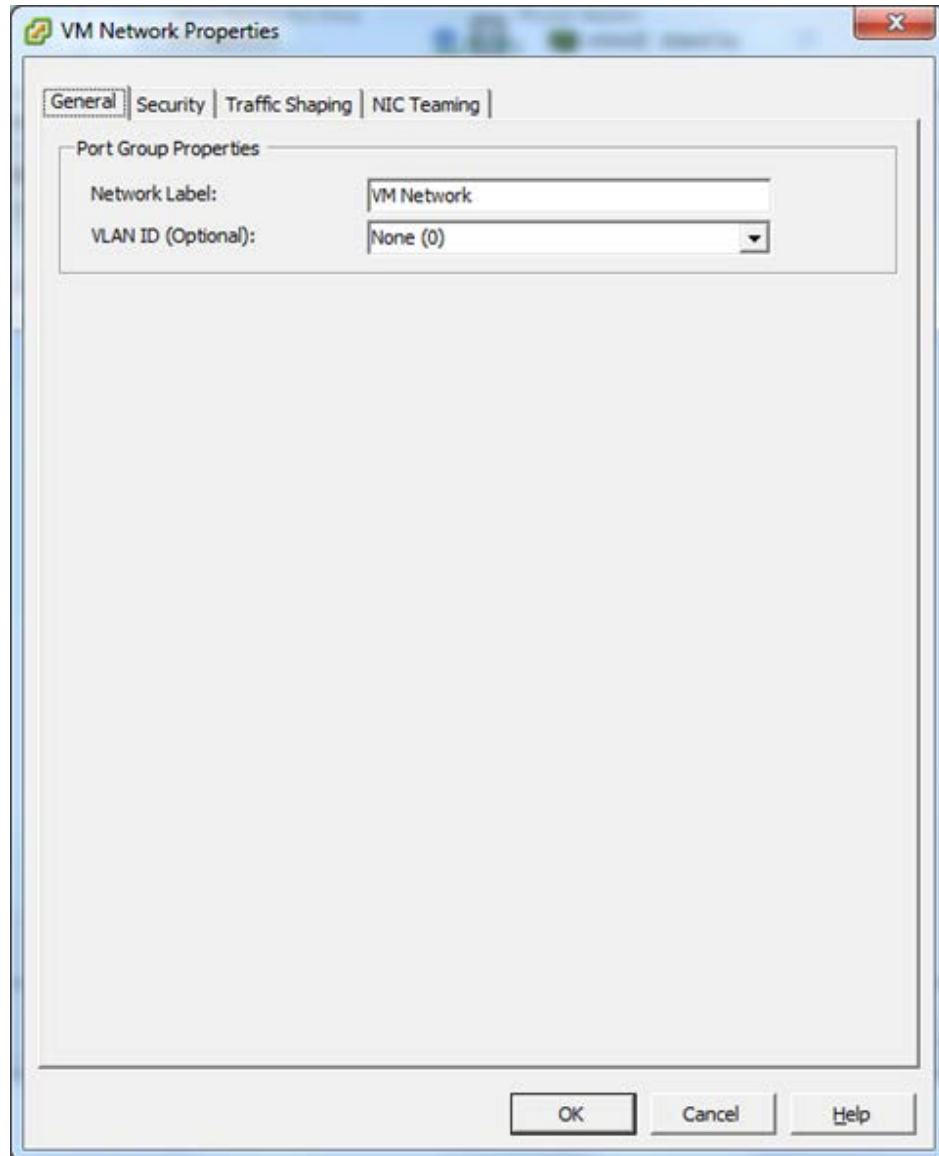
6. In the **vSwitch Properties** window, click on the **Network Adapters** tab.
7. Verify that the correct vmNICs are part of the Front-end Network. If they are incorrect:
 - a. Add the correct vmNICs
 - b. Remove the incorrect vmNICs.



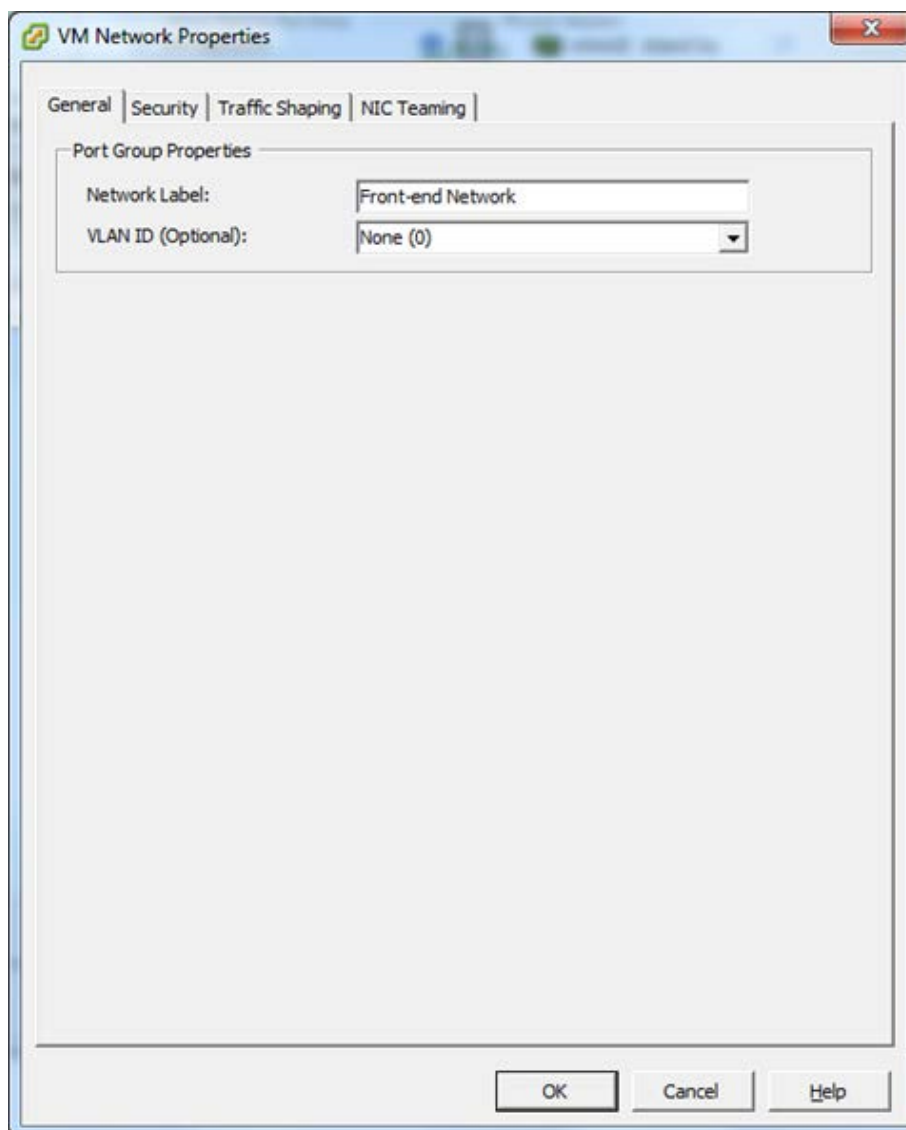
8. Click on the **Ports** tab.
9. In the left side window, select **VM Network** and click **Edit**.



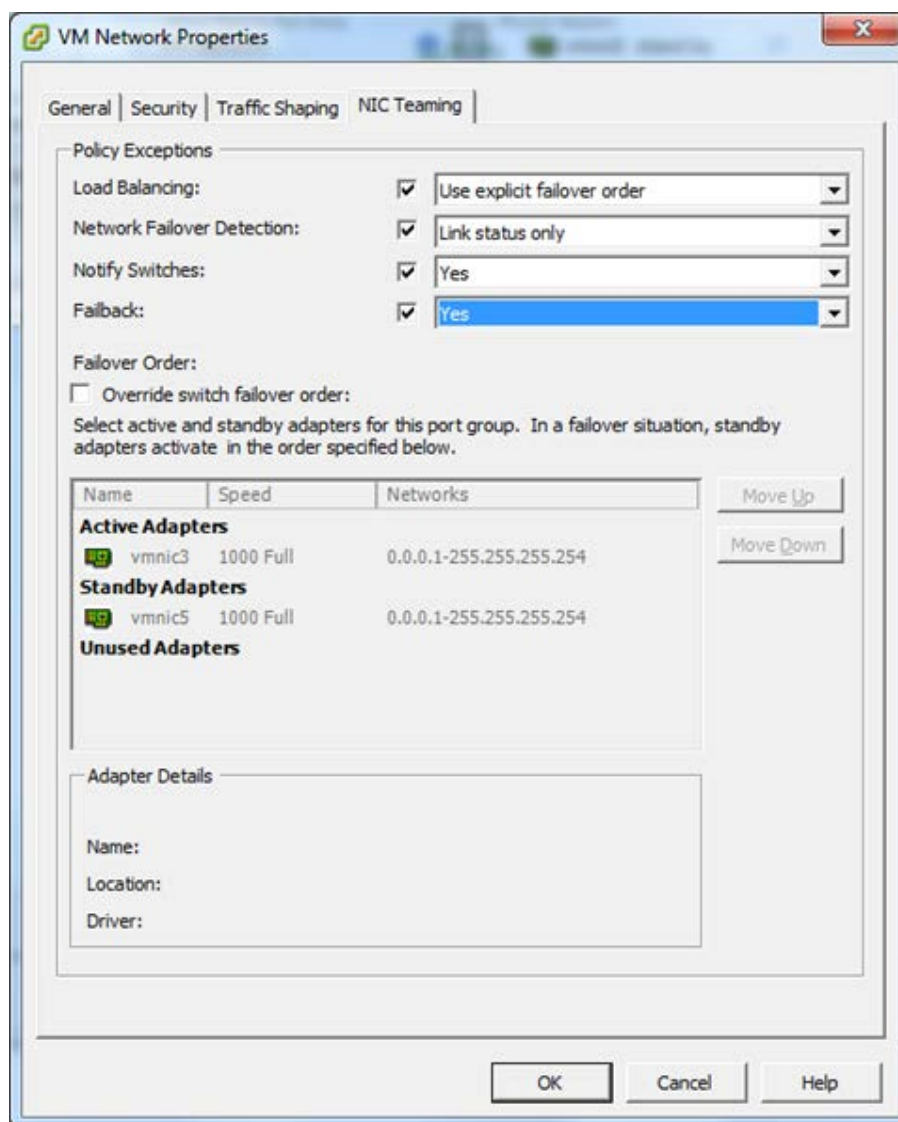
10. In the **VM Network Properties** window, change the Network Label to **Front-end Network**. Do **NOT** click **OK**.



11. Click on the **NIC Teaming** tab.



12. In the **NIC Teaming** tab, select the first four check boxes, and select the following for the drop down menus:
 - a. For **Load Balancing** select **Use explicit failover order**.
 - b. For **Network Failover Detection** select **Link status only**.
 - c. For **Notify Switches** select **Yes**.
 - d. For **Failback** select **Yes**.



13. Click **OK**.
14. In the **vSwitch Properties** window, click **Close**.

15. Repeat the steps to configure the Front-end Network for each ESXi host that will be part of the HCP-VM system.

Configure networking for Back-end switching


The HCP private Back-end network needs to be configured so that it can provide inter-node communication and data transfer. You are responsible for configuring the network.

To configure the back-end network for switching:

1. Access the vSphere Client.
2. In the left side navigation bar, select the first ESXi host.
3. In the right side window, click on the **Configuration** tab.
4. Click on **Networking** in the **Hardware** section.
5. Click **Add Networking** button located in the top right of the right hand window.
6. In the **Add Network Wizard** select **Virtual Machine**.
7. Click **Next**.

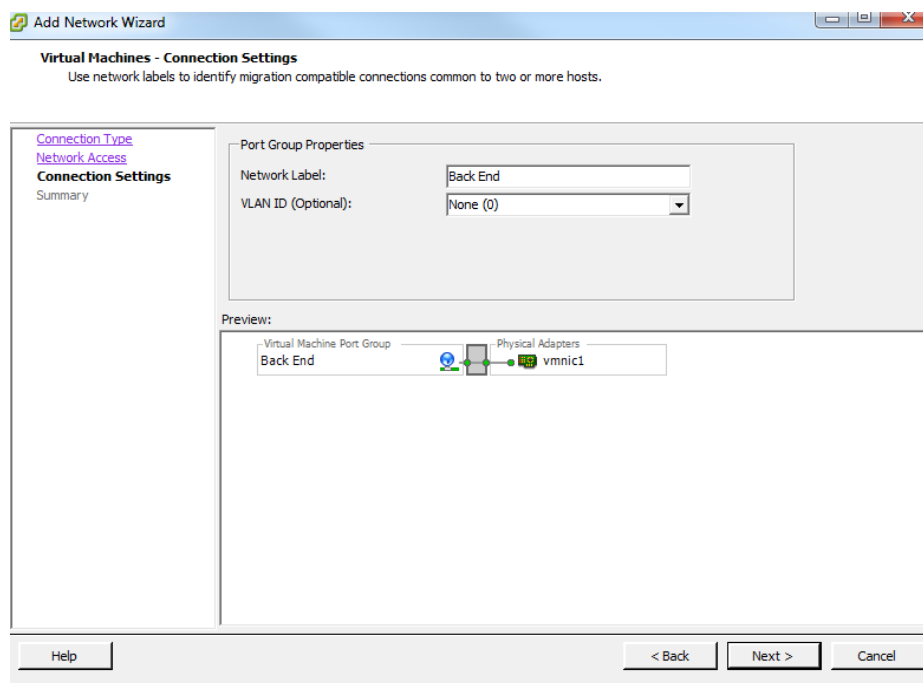
8. Select the Physical NIC to use for the Back-end network.

9. Click **Next**.



10. Name the Network label **Back-End**.

11. Click **Next**.



12. Review your changes and click **Finish**.
13. Repeat the steps to configure the Back-end Network for each ESXi host that will be part of the HCP-VM system.

Verifying ESXi configuration on all hosts

The Front-end and Back-end networks must be configured for each ESXi host added to the HCP-VM system. To make sure that all changes are correct, select a single ESXi host on the left side navigation bar, and click on the **Configuration** tab in the right side window. Beginning with **Processors**, click each components listed in the **Hardware** section and make that their specifications matches the images below.

No changes have been made to the **Advanced Settings** or **Power Management** sections.



Important: Repeat this verification on all ESXi hosts in the vSphere HA cluster.

Preparing the ESXi network for the HCP-VM OVF deployment

Processors

The screenshot displays the vSphere Client interface for configuring a virtual machine. The left sidebar shows the inventory tree with the path: hcp-vm_center > hcp-vm_vcenter > hcp-vm_vcenter-1 > 172.20.27.102 > HCP-VM_vCenter S. The main pane is titled '172.20.27.102 VMware ESXi, 5.1.0, 799733 | Evaluation (59 days remaining)'. The 'Processors' tab is selected under the 'Hardware' section. The 'Processors' configuration is shown in the 'General' tab, with the following details:

Property	Value
Model	Intel(R) Xeon(R) CPU X5675 @ 3.07GHz
Processor Speed	3.1 GHz
Processor Sockets	2
Processor Cores per Socket	6
Logical Processors	24
Hyperthreading	Enabled

The 'System' tab is also visible, showing the following details:

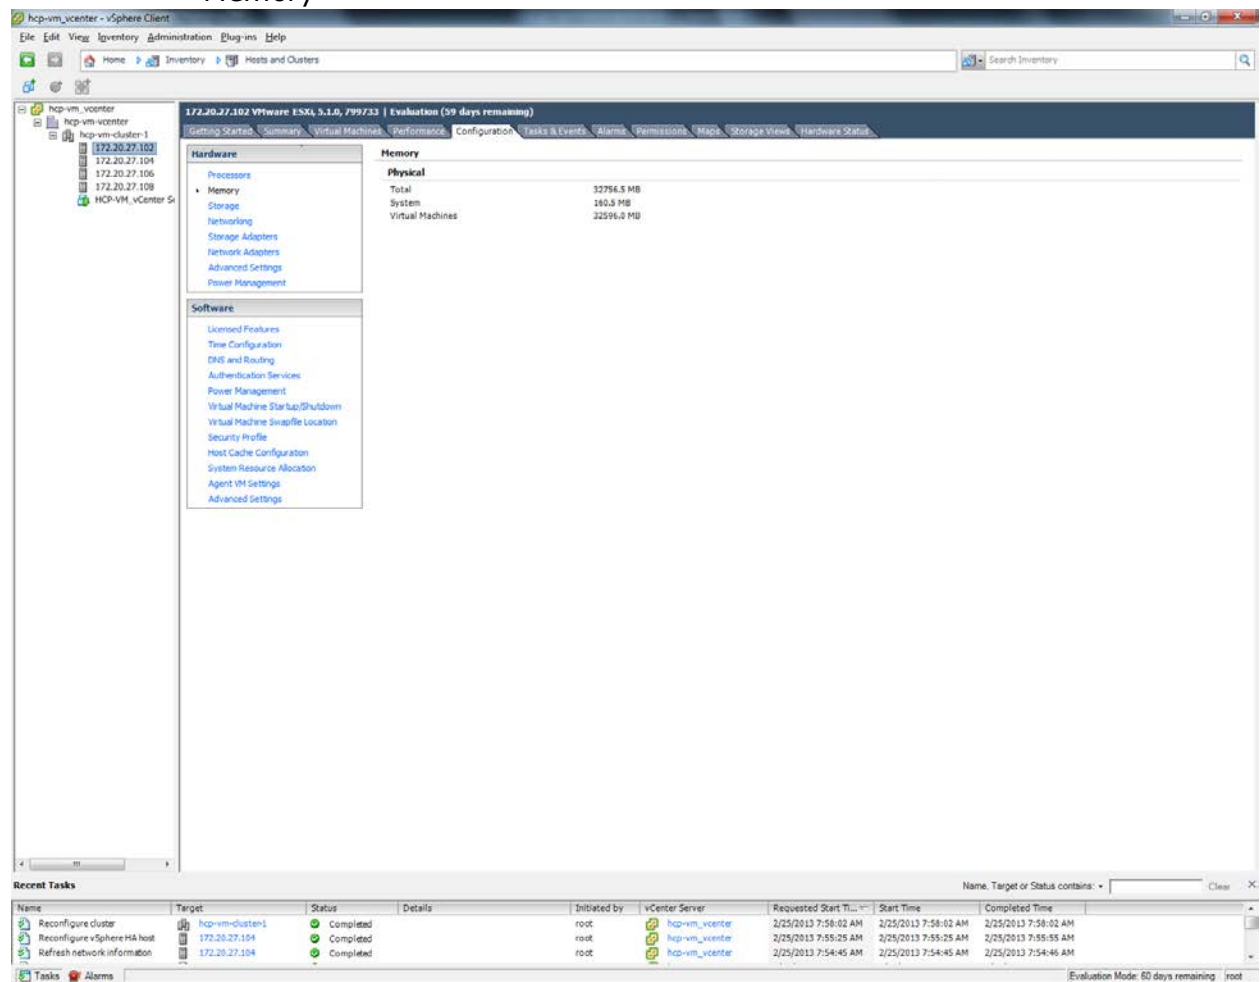
Property	Value
Manufacturer	HITACHI
Model	Compute Blade ES5A2
BIOS Version	4.6.3
Release Date	8/19/2011 12:00:00 AM
Service Tag	4600E70 T223000641
Asset Tag	None

The 'Recent Tasks' table at the bottom shows the following tasks:

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Reconfigure cluster	hcp-vm_vcenter-1	Completed		root	hcp-vm_vcenter	2/26/2013 7:54:02 AM	2/26/2013 7:54:02 AM	2/26/2013 7:54:02 AM
Reconfigure vSphere HA host	172.20.27.104	Completed		root	hcp-vm_vcenter	2/28/2013 7:55:28 AM	2/28/2013 7:55:28 AM	2/28/2013 7:55:28 AM
Refresh network information	172.20.27.104	Completed		root	hcp-vm_vcenter	2/25/2013 7:54:45 AM	2/25/2013 7:54:45 AM	2/25/2013 7:54:46 AM

The bottom status bar indicates 'Evaluation Mode: 60 days remaining' and 'root'.

Memory



Storage

Preparing the ESXi network for the HCP-VM OVF deployment

The screenshot shows the vSphere Client interface for the HCP-VM vCenter. The left sidebar displays the inventory tree with the path: hcp-vm_vcenter > hcp-vm-vcenter > hcp-vm-cluster-1 > 172.20.27.102 > HCP-VM_vCenter S. The main pane shows the configuration for the ESXi host 172.20.27.102. The 'Datastores' tab is selected, showing a table of datastores:

Identification	Status	Device	Drive Type	Capacity	Free	Type	Last Update	Actions
datastore1	Normal	HITACHI Fibre C...	Non-SSD	10.00 GB	9.14 GB	VMFS5	2/25/2013 8:35:43 PM	En
hcp-vm_cluster-1_node_1_datastore_1	Normal	HITACHI Fibre C...	Non-SSD	1.20 TB	1.20 TB	VMFS5	2/25/2013 8:35:44 PM	En
hcp-vm_cluster-1_node_2_datastore_1	Normal	HITACHI Fibre C...	Non-SSD	1.20 TB	1.20 TB	VMFS5	2/25/2013 8:35:44 PM	En
hcp-vm_cluster-1_node_3_datastore_1	Normal	HITACHI Fibre C...	Non-SSD	1.20 TB	1.20 TB	VMFS5	2/25/2013 8:35:44 PM	En
hcp-vm_cluster-1_node_4_datastore_1	Normal	HITACHI Fibre C...	Non-SSD	1.20 TB	1.20 TB	VMFS5	2/25/2013 8:35:44 PM	En
vCenter_datastore	Normal	HITACHI Fibre C...	Non-SSD	199.75 GB	105.70 GB	VMFS5	2/25/2013 8:35:44 PM	En

The 'Recent Tasks' pane at the bottom shows a list of completed tasks:

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Rename datastore	hcp-vm_cluster-1_node_1	Completed		root	hcp-vm_vcenter	2/25/2013 8:34:56 PM	2/25/2013 8:34:56 PM	2/25/2013 8:34:57 PM
Rescan VMFS	172.20.27.108	Completed		System	hcp-vm_vcenter	2/25/2013 8:34:22 PM	2/25/2013 8:34:22 PM	2/25/2013 8:34:24 PM
Rescan VMFS	172.20.27.106	Completed		System	hcp-vm_vcenter	2/25/2013 8:34:22 PM	2/25/2013 8:34:22 PM	2/25/2013 8:34:25 PM

Networking

The screenshot shows the vSphere Client interface for the HCP-VM vCenter, specifically the 'Networking' configuration page. The left sidebar shows the inventory tree with the path: hcp-vm_vcenter > hcp-vm-vcenter > hcp-vm-cluster-1 > 172.20.27.102 > HCP-VM_vCenter S. The main pane shows the configuration for the ESXi host 172.20.27.102. The 'Networking' tab is selected, showing a diagram of the network configuration:

Standard Switch: vSwitch0

- Virtual Machine Port Group: Front-end Network
- Physical Adapters: vmnic5 (stand by), vmnic3 (1000 Full)

Standard Switch: vSwitch1

- Virtual Machine Port Group: Back-end Network
- Physical Adapters: vmnic4 (1000 Full), vmnic6 (1000 Full)

The 'Recent Tasks' pane at the bottom shows a list of completed tasks:

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Rename datastore	hcp-vm_cluster-1_node_1	Completed		root	hcp-vm_vcenter	2/25/2013 8:34:56 PM	2/25/2013 8:34:56 PM	2/25/2013 8:34:57 PM
Rescan VMFS	172.20.27.108	Completed		System	hcp-vm_vcenter	2/25/2013 8:34:22 PM	2/25/2013 8:34:22 PM	2/25/2013 8:34:24 PM
Rescan VMFS	172.20.27.106	Completed		System	hcp-vm_vcenter	2/25/2013 8:34:22 PM	2/25/2013 8:34:22 PM	2/25/2013 8:34:25 PM

Storage Adapters

hcp-vm_vcenter - vSphere Client

File Edit View Inventory Administration Plug-ins Help

Home Inventory Hosts and Clusters

Search Inventory

hcp-vm_vcenter

hcp-vm_vcenter

hcp-vm-cluster-1

172.20.27.104

172.20.27.106

172.20.27.108

HCP-VM_vCenter S...

172.20.27.102 VMware ESXi 5.1.0, 799733 | Evaluation (59 days remaining)

Getting Started Summary Virtual Machines Performance Configuration Tasks & Events Alarms Permissions Maps Storage Views Hardware Status

Hardware

Processors

Memory

Storage

Networking

Storage Adapters

Network Adapters

Advanced Settings

Power Management

Software

Licensed Features

Time Configuration

DNS and Routing

Authentication Services

Power Management

Virtual Machine Startup/Shutdown

Virtual Machine Swapfile Location

Security Profile

Host Cache Configuration

System Resource Allocation

Agent VM Settings

Advanced Settings

Storage Adapters

Add... Remove Refresh Rescan All...

Device	Type	WWN
ICH10 2 port SATA IDE Controller		
vmhba1	Block SCSI	
vmhba33	Block SCSI	
ICH10 4 port SATA IDE Controller		
vmhba0	Block SCSI	
vmhba32	Block SCSI	
LPe12000 8Gb Fibre Channel Host Adapter		
vmhba2	Fibre Channel	20:00:00:00:c9:a1:95:98 10:00:00:00:c9:a1:95:98
vmhba3	Fibre Channel	20:00:00:00:c9:a1:95:99 10:00:00:00:c9:a1:95:99

Details

vmhba2

Model: LPe12000 8Gb Fibre Channel Host Adapter

WWN: 20:00:00:00:c9:a1:95:98 10:00:00:00:c9:a1:95:98

Targets: 1 Devices: 6 Paths: 6

View: Devices Paths

Name	Identifier	Runtime Name	Operational St...	LUN	Type	Drive Type	Transport	Capacity	OW
HITACHI Fibre Channel Disk (naa....	naa.60060e8013274600...	vmhba2:C0:T0:L0	Mounted	0	disk	Non-SSD	Fibre Channel	15.00 GB	NM
HITACHI Fibre Channel Disk (naa....	naa.60060e8013274600...	vmhba2:C0:T0:L1	Mounted	1	disk	Non-SSD	Fibre Channel	1.20 TB	NM
HITACHI Fibre Channel Disk (naa....	naa.60060e8013274600...	vmhba2:C0:T0:L2	Mounted	2	disk	Non-SSD	Fibre Channel	1.20 TB	NM
HITACHI Fibre Channel Disk (naa....	naa.60060e8013274600...	vmhba2:C0:T0:L3	Mounted	3	disk	Non-SSD	Fibre Channel	200.00 G	NM
HITACHI Fibre Channel Disk (naa....	naa.60060e8013274600...	vmhba2:C0:T0:L4	Mounted	4	disk	Non-SSD	Fibre Channel	1.20 TB	NM
HITACHI Fibre Channel Disk (naa....	naa.60060e8013274600...	vmhba2:C0:T0:L5	Mounted	5	disk	Non-SSD	Fibre Channel	1.20 TB	NM

Recent Tasks

Name, Target or Status contains: Clear X

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Start Time	Completed Time
Rename datastore	hcp-vm_cluster-1_node...	Completed		root	hcp-vm_vcenter	2/25/2013 8:34:56 PM	2/25/2013 8:34:56 PM	2/25/2013 8:34:57 PM
Rescan VMFS	172.20.27.108	Completed		System	hcp-vm_vcenter	2/25/2013 8:34:22 PM	2/25/2013 8:34:22 PM	2/25/2013 8:34:24 PM
Rescan VMFS	172.20.27.106	Completed		System	hcp-vm_vcenter	2/25/2013 8:34:22 PM	2/25/2013 8:34:22 PM	2/25/2013 8:34:25 PM

Tasks Alarms

Evaluation Mode: 59 days remaining root

Verify Network Adapters

The screenshot shows the vSphere Client interface with the 'Configuration' tab selected for a virtual machine. The left sidebar shows the hierarchy: hcp-vm_center > hcp-vm-vcenter > hcp-vm-cluster-1 > 172.20.27.102 > HCP-VM_vCenter Sx. The main pane displays the 'Network Adapters' section under the 'Hardware' category. A table lists the network adapters and their configurations.

Device	Speed	Configured	Switch	MAC Address	Observed IP ranges	Wake on LAN Supported
Intel Corporation 82567LF-2 Gigabit Network Connection						
vmnic0	100 Half	Negotiate	None	00:25:90:58:98:e2	172.20.27.1-172.20.27.127	Yes
Intel Corporation 82576 Gigabit Network Connection						
vmnic6	1000 Full	Negotiate	None	00:1b:21:af:7f:ff	None	No
vmnic5	1000 Full	Negotiate	vSwitch0	00:1b:21:af:7f:fe	172.20.27.1-172.20.27.127	No
vmnic4	1000 Full	Negotiate	None	00:1b:21:af:7f:fd	None	No
vmnic3	1000 Full	Negotiate	vSwitch0	00:1b:21:af:7f:fc	172.20.27.1-172.20.27.127	Yes
vmnic2	1000 Full	Negotiate	None	00:25:90:58:98:e1	None	No
vmnic1	1000 Full	Negotiate	None	00:25:90:58:98:e0	None	Yes

At the bottom, the 'Recent Tasks' table shows the following entries:

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Reconfigure cluster	hcp-vm-cluster-1	Completed		root	hcp-vm_vcenter	2/25/2013 7:54:02 AM	2/25/2013 7:54:02 AM	2/25/2013 7:54:03 AM
Reconfigure vSphere HA host	172.20.27.104	Completed		root	hcp-vm_vcenter	2/25/2013 7:55:25 AM	2/25/2013 7:55:25 AM	2/25/2013 7:55:55 AM
Refresh network information	172.20.27.104	Completed		root	hcp-vm_vcenter	2/25/2013 7:54:45 AM	2/25/2013 7:54:45 AM	2/25/2013 7:54:46 AM

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 the OVF Zip file

On your computer, access the DVD that contains the virtual machine image file and unpack the zip vmdkIso: HS421_x.x.x.x.iso.zip or the zip rdmIso: HS433_x.x.x.x.iso.zip file into a directory of your choice.

To unpack the file that contains the virtual machine image:

1. On your computer, unpack the zip vmdkIso: HS421_x.x.x.x.iso.zip or the zip rdmIso: HS433_x.x.x.x.iso.zip file into a directory of your choice.
2. Navigate into the folder you unpacked the zip.
3. Unpack the ISO files vmdkIso: HS421_x.x.xx.iso or the rdmIso: HS433_x.x.x.x.iso.

Deploying the HCP-VM OVF VDMK

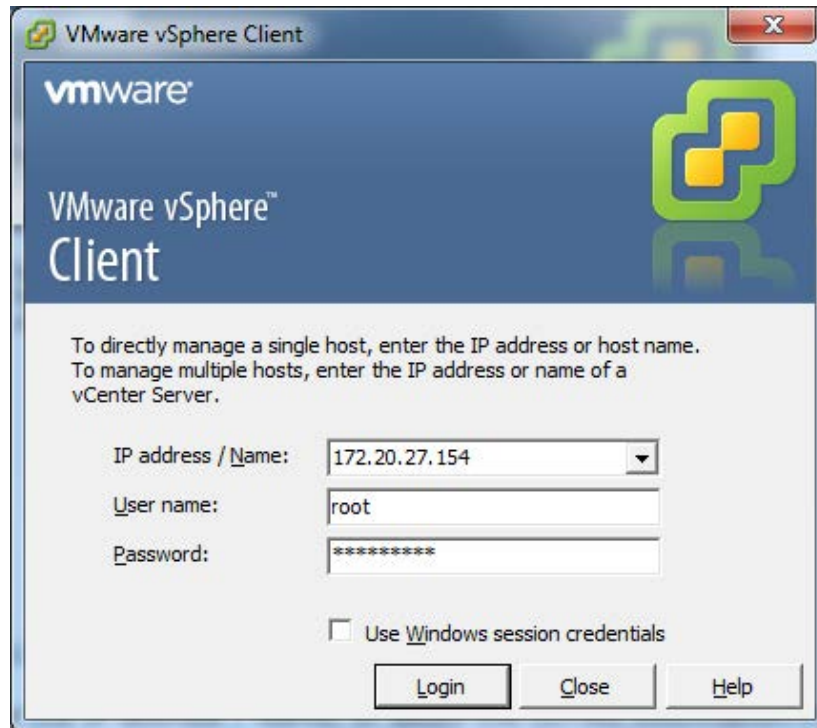
There are two different OVFs that can be deployed. These steps are for the VMDK deploy. The RDM procedure is identical to this one except for some minor differences. You only need to install one of them.

Step 1: Log into the ESXi server

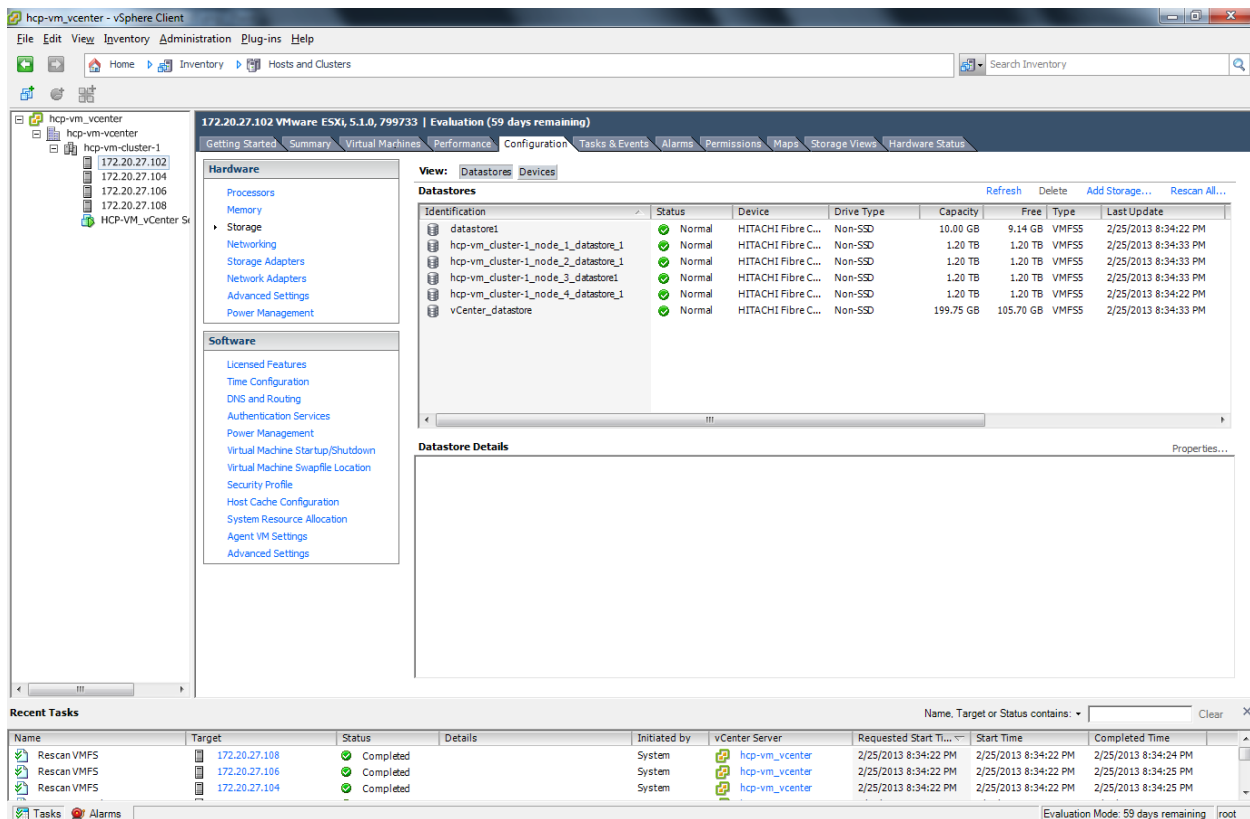
To deploy the HCP-VM OVF:

1. Launch the vSphere client.

2. Enter the IP address / Name, or select the correct information from the drop down menu to connect to the vCenter server where the vSphere HA cluster was configured for the HCP-VM system.
3. Enter the User name and Password.
4. Click **Login**.

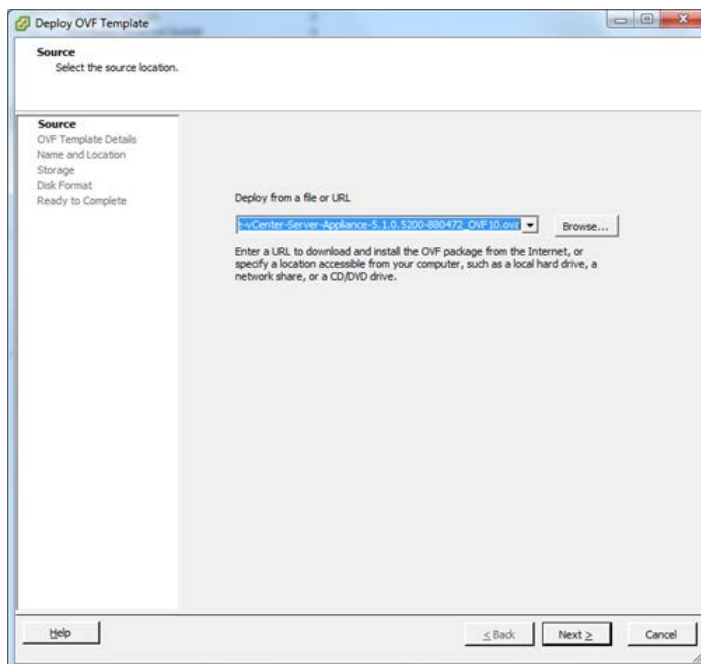


5. Once logged in to the vSphere Client, you should see the datacenters, clusters and ESXi nodes on the left side navigation bar that were previously added to vCenter.
6. In the navigation bar on the left hand side, select the ESXi host to target for the deploy and click **File** in the toolbar at the top of the screen and in the submenu click **Deploy OVF Template**.

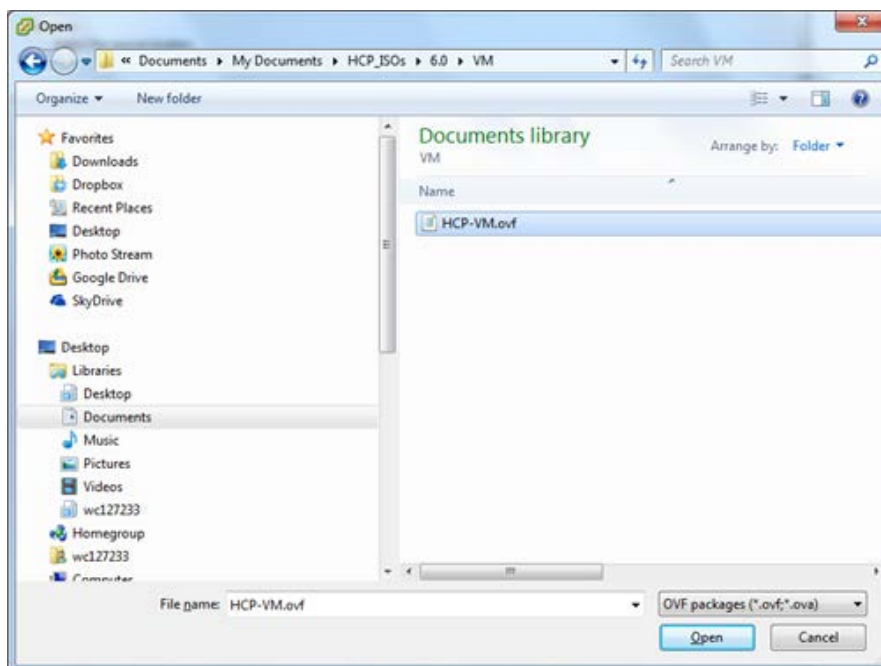


Step 2: Deploy VMDK OVF Template

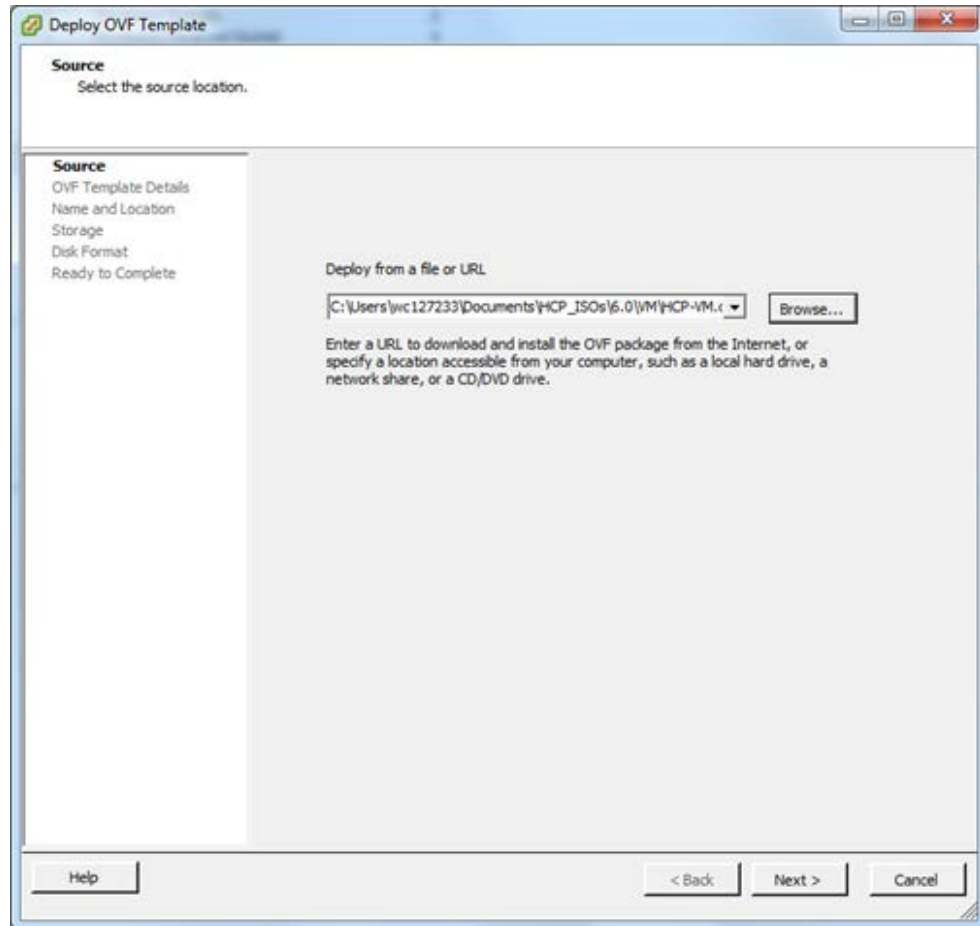
1. In the **Deploy OVF Template** window, click on the **Browse** button and navigate to the local file system to the location that HS421_7.0.XX.zip you extracted.




2. Select the HCP-VM-VMDK.ovf file and click **Open**.



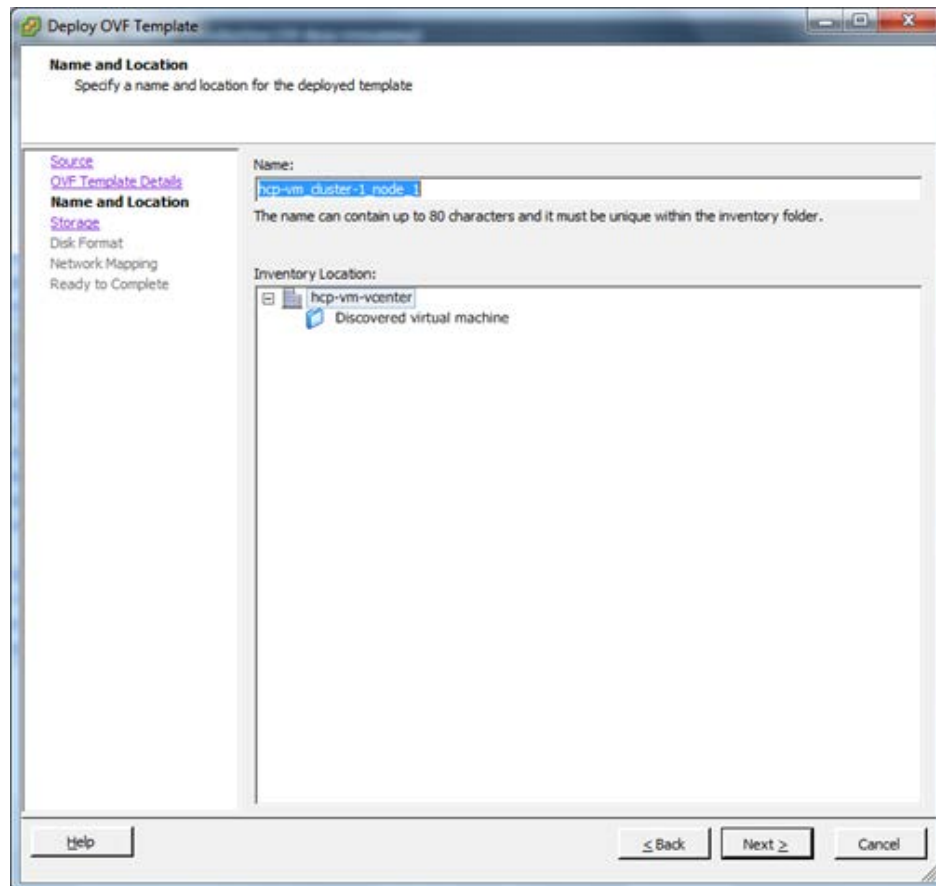
3. Once the path to the OVF file has been selected, click **Next**.



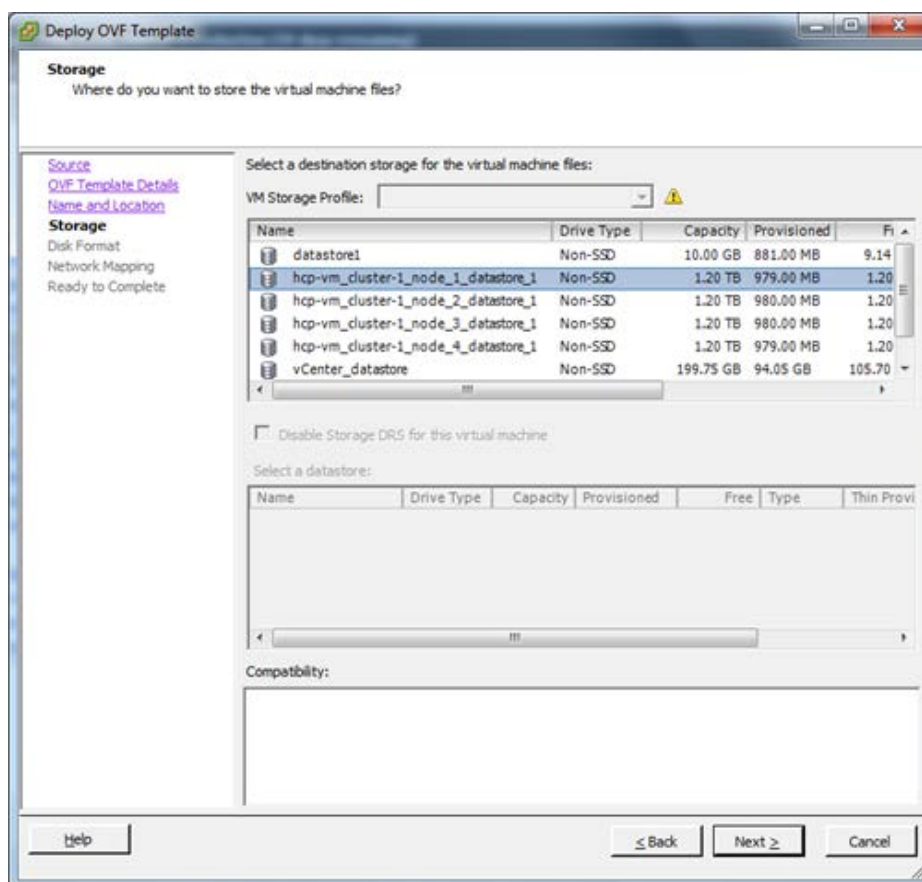
4. Verify that the OVF template details show the product is HCP-VM and that the **Size on disk** is 1.0TB (thick provisioned).
5. Click **Next**.



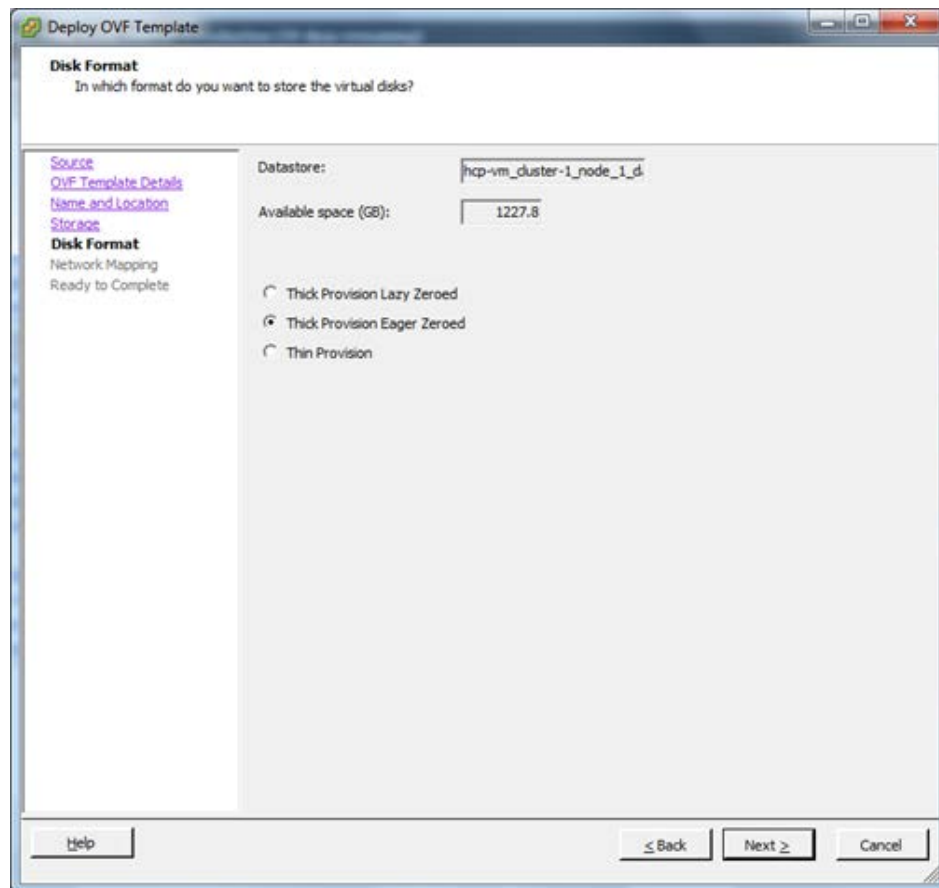
6. Enter a name for the node that is being deployed. It should be named something meaningful for the installation. For example: `hcp-vm_cluster-1_node-1_node_1`.
7. Once the name has been entered, click **Next**.



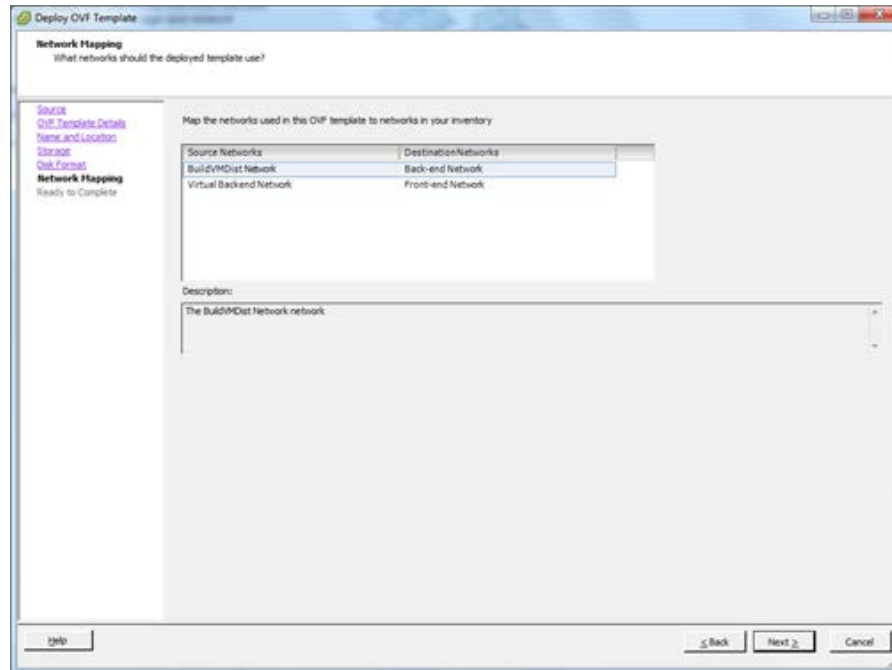
8. Select 1 datastore from the list you previously added to the ESXi hosts. If you're doing a consecutive load, make sure to select the next datastore down (from the previous load) on the list. The selected datastore should have a capacity of at least 1.2TB.
9. Click **Next**.



10. Verify that the datastore you selected matches the **Available space** expected for the datastore.
11. Select **Thick Provision Eager Zeroed**.
12. Click **Next**.

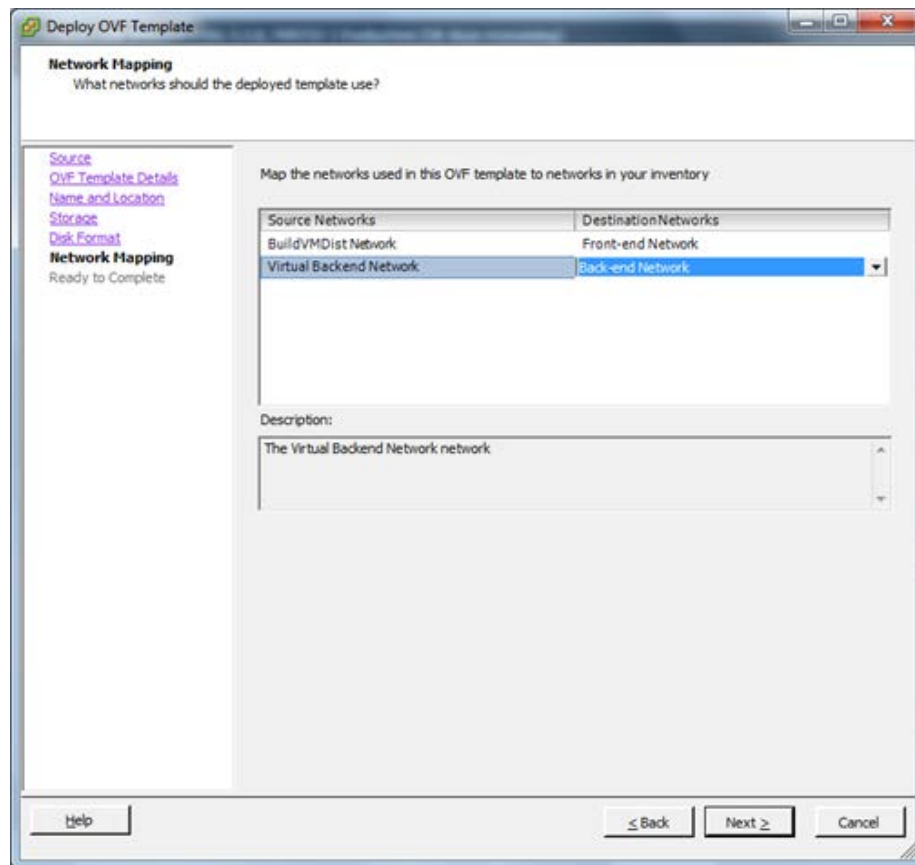


- 13.** Hover your cursor over a Destination network to make a drop down menu button appear. Click on the drop down menu for Destination Networks.
- 14.** Change the Destination Networks so that the Front-end Network aligns with the BuildVMDisk Network.
- 15.** Change the Destination Networks so that the Back-end Network aligns with the Virtual Back-end Network.



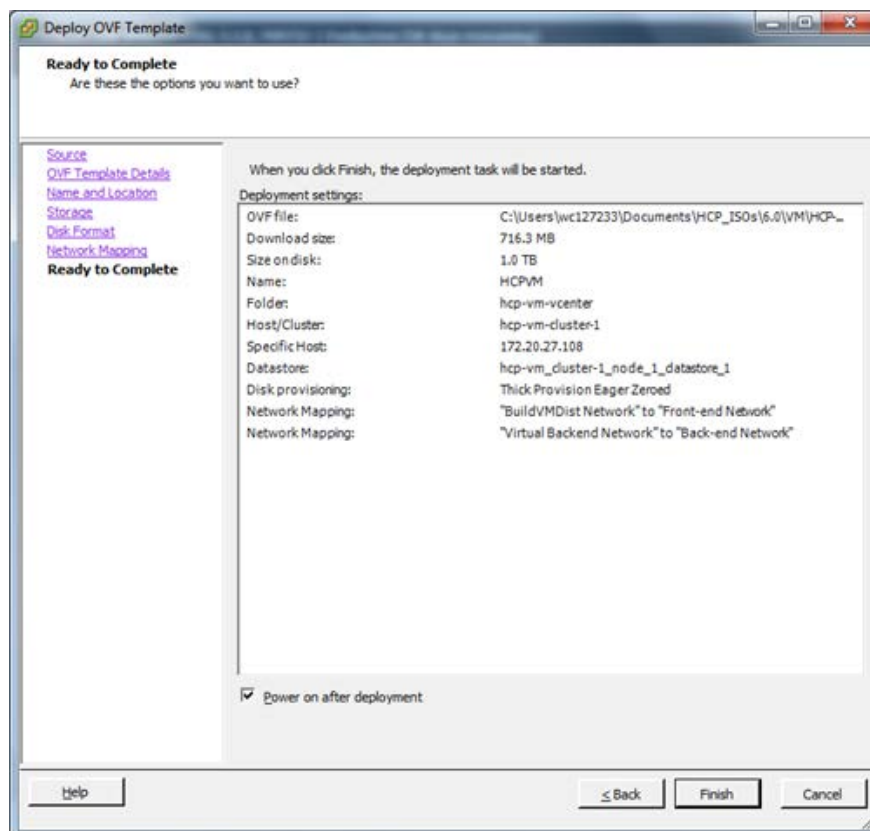
16. Verify the Destination Networks mimic the following image.

17. Click **Next**.



Important: Do **NOT** select the Power on checkbox.

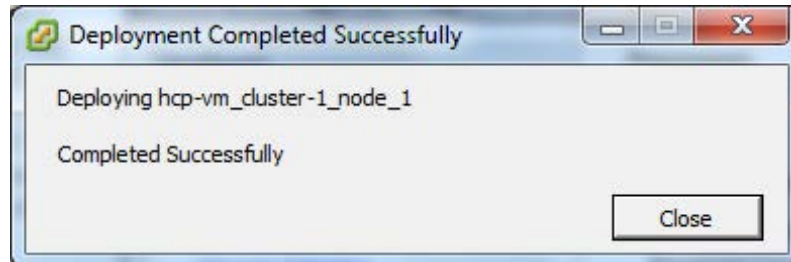
18. Verify the information in Deployment settings matches what was previously entered:
 - a. If so click on **Finish** to begin the OVF deploy.
 - b. If not, go back and correct any information that needs to be changed.



Important:

- The VMDK OVF deploy can take up to an hour or more. This is due to the fact that VMware is preparing the vmdk's for use by the HCP-VM node. There will not be any indication of progress in the OVF deploy window (just a spinning cursor) or in the deploy task at the bottom of vSphere client (just "in progress"). The only indication will come when checking the available capacity of the datastore. This will show a decrease in available capacity when the first vmdk has been prepared.
- You must repeat the OVF deployment for each of the nodes that are going to be part of the HCP-VM system.
- Make sure that you have highlighted the desired ESXi host that you want the HCP-VM node to run on initially before importing the OVF.

Once the OVF Deploy is completed, you will see the following message.



Deploy the HCP-VM OVF RDM


There are two different OVFs that can be deployed. These steps are for the RDM deploy. The VDMK procedure is identical to this one except for some minor differences. You only need to install one of them.

Step 1: Log into the ESXi server


To deploy the HCP-VM OVF:

1. Launch the vSphere client.

2. Enter the IP address / Name, or select the correct information from the drop down menu to connect to the vCenter server where the vSphere HA cluster was configured for the HCP-VM system.
3. Enter the User name and Password.
4. Click **Login**.

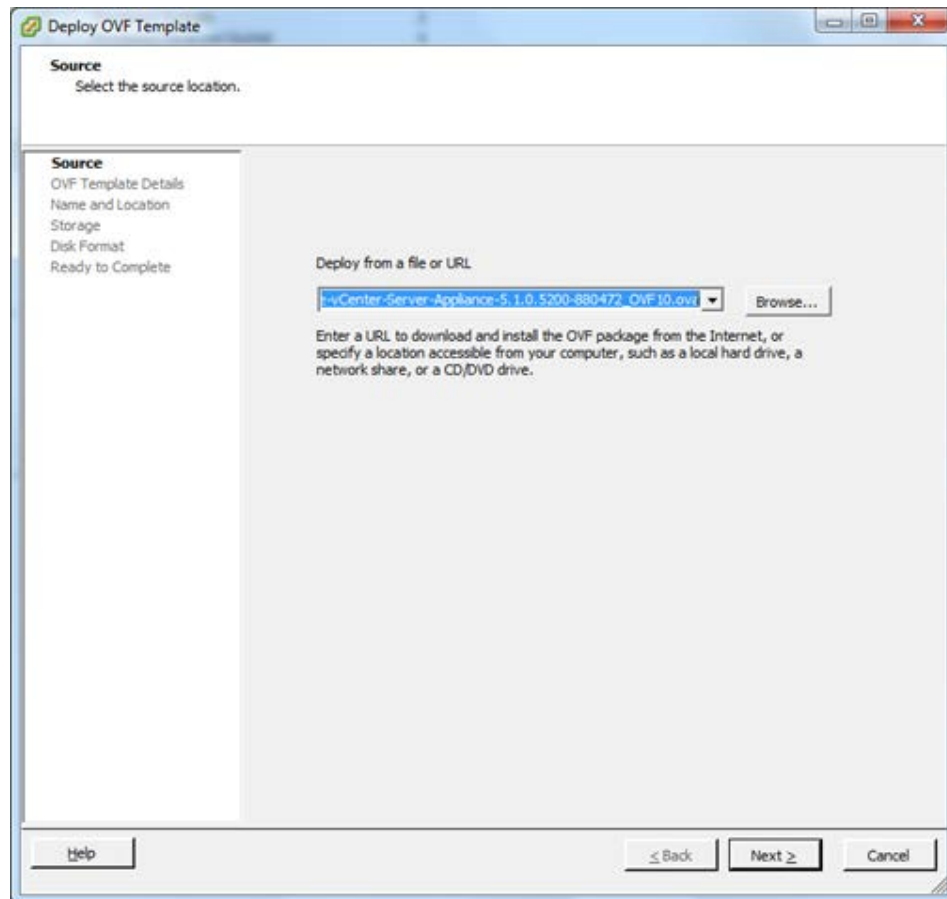


5. Once logged in to the vSphere Client, you should see the datacenters, clusters and ESXi nodes that were previously added to vCenter in the left side navigation bar.
6. In the navigation bar on the left hand side, select the ESXi host to target for the deploy and click **File > Deploy OVF Template** from the toolbar at the top of the screen.



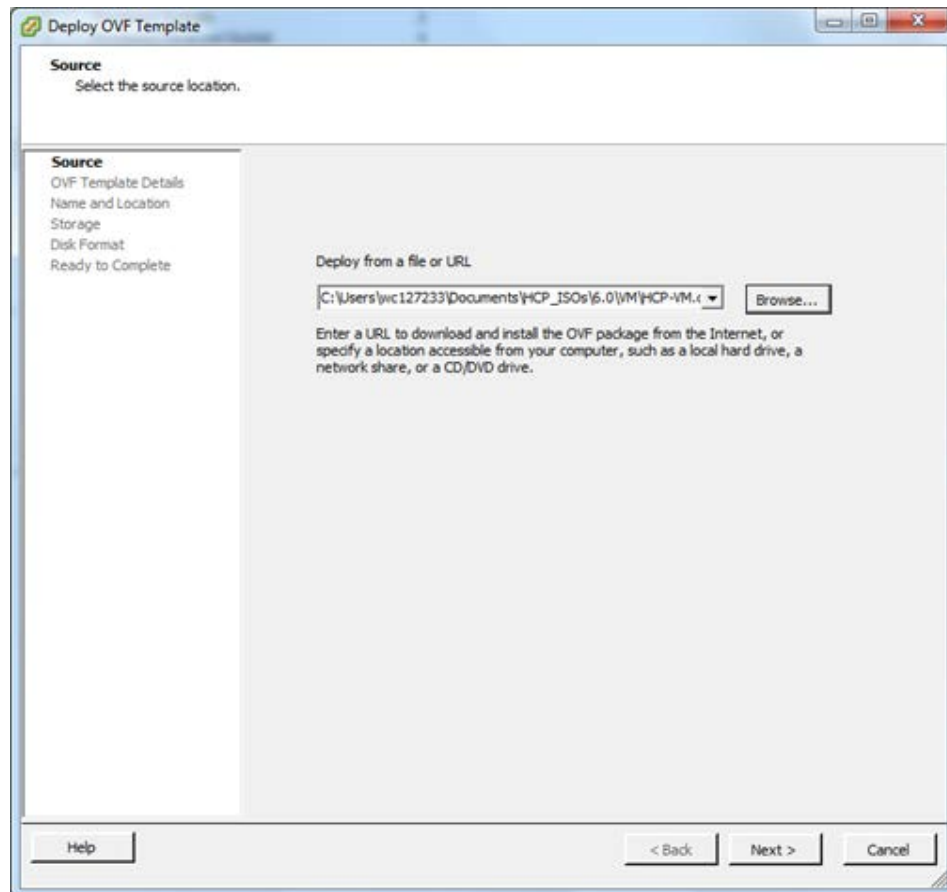
Step 2: Deploy RDM OVF Template

1. In the Deploy OVF Template window, click on the **Browse** button and navigate to the local file system to the location that HS433_7.0.XX.zip you extracted.

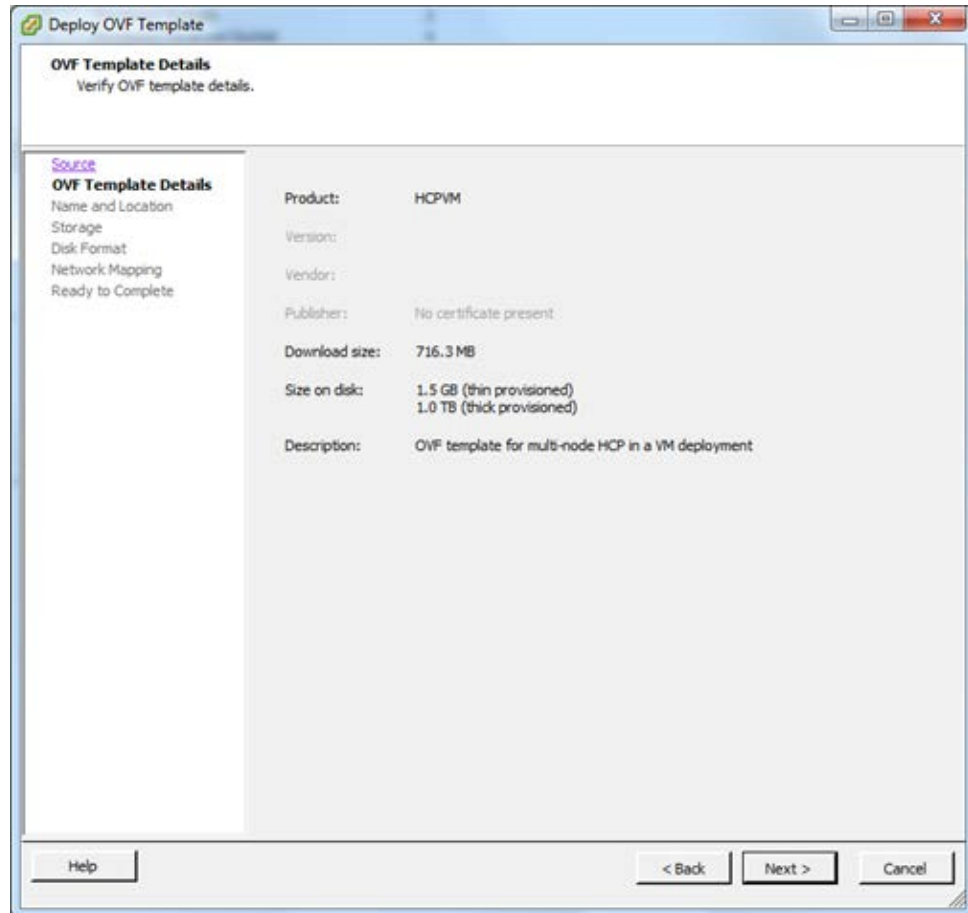


2. Select the HCP-VM-RDM.ovf file and click **Open**.

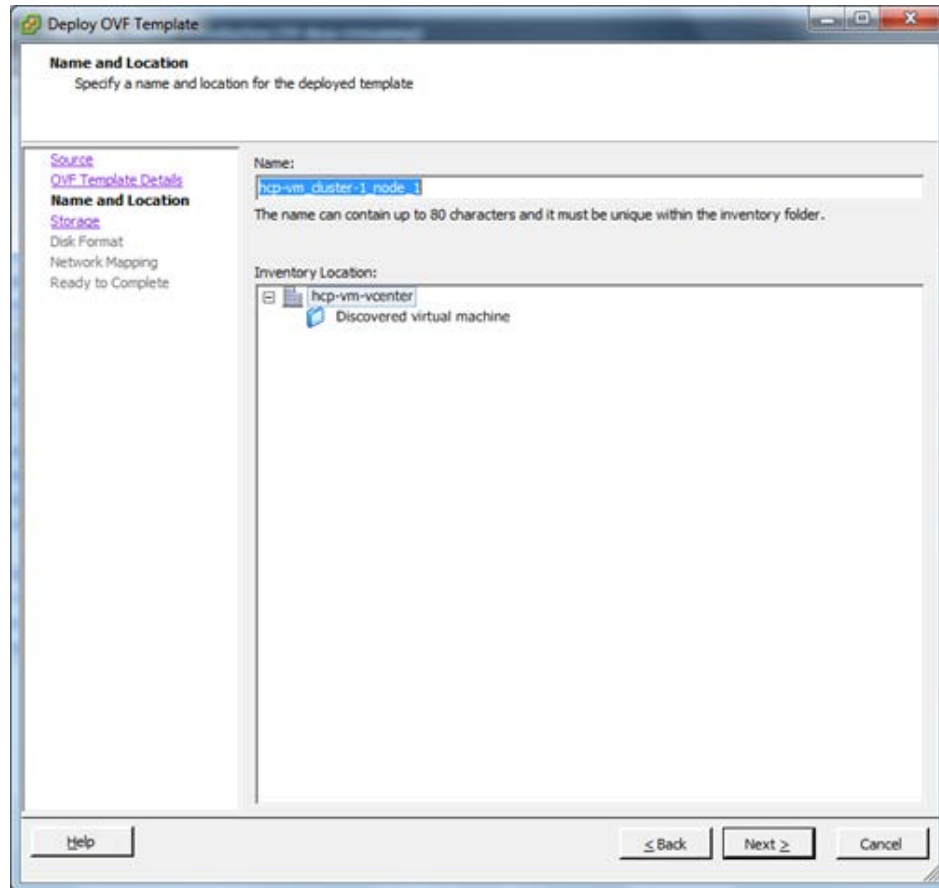
3. Once the path to the OVF file has been selected, click on **Next** to proceed with the **Deploy OVF Template** wizard.



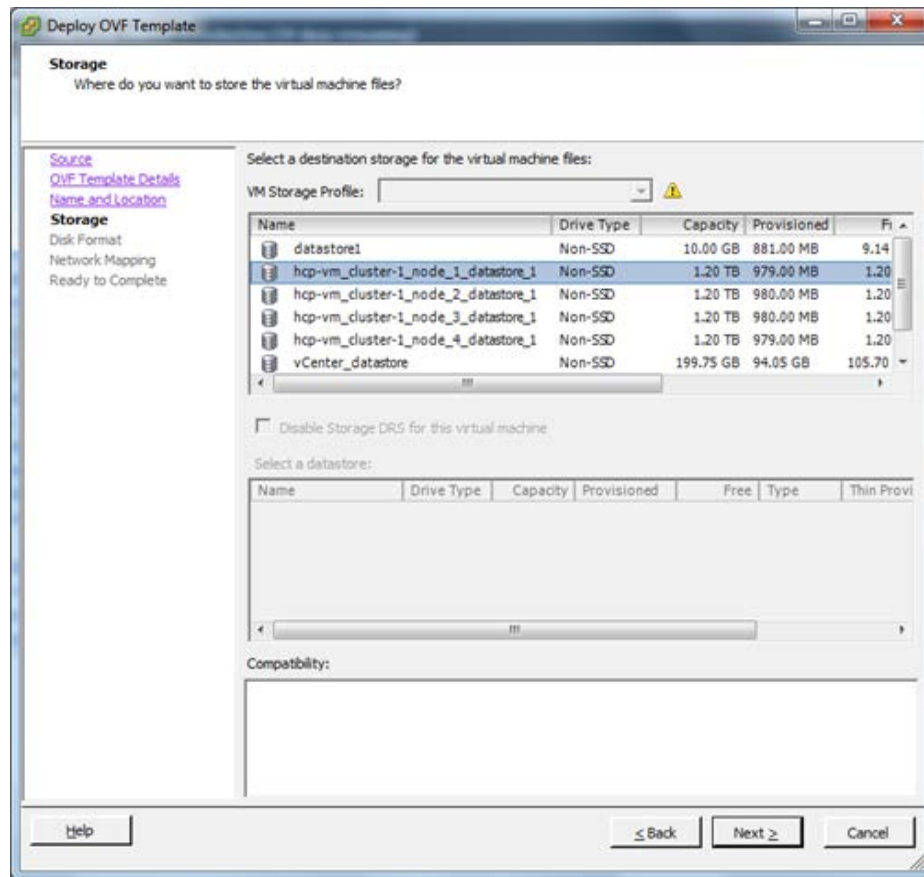
4. Verify that the OVF template details show that the product is HCP-VM and the **Size on disk** is 32.0GB (thick provisioned).
5. Click **Next**.



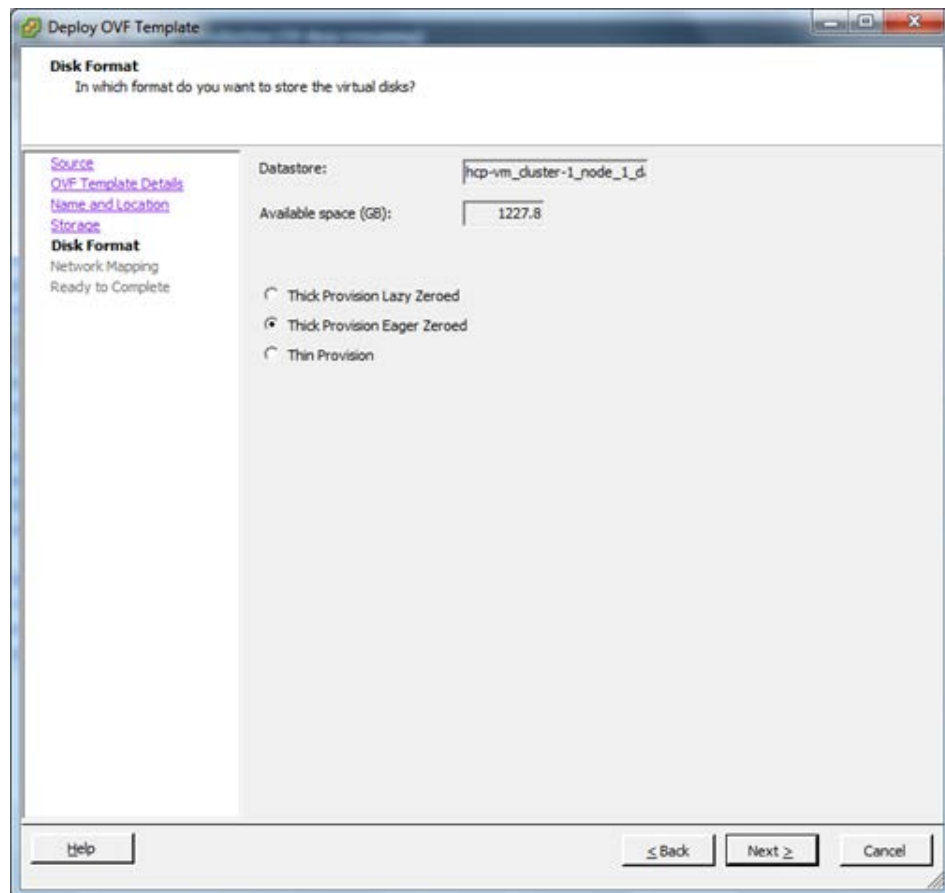
6. Enter a name for the node that is being deployed. It should be named something meaningful for the installation. For example: `HCPVM-node-1`.
7. Once the name has been entered, click **Next**.



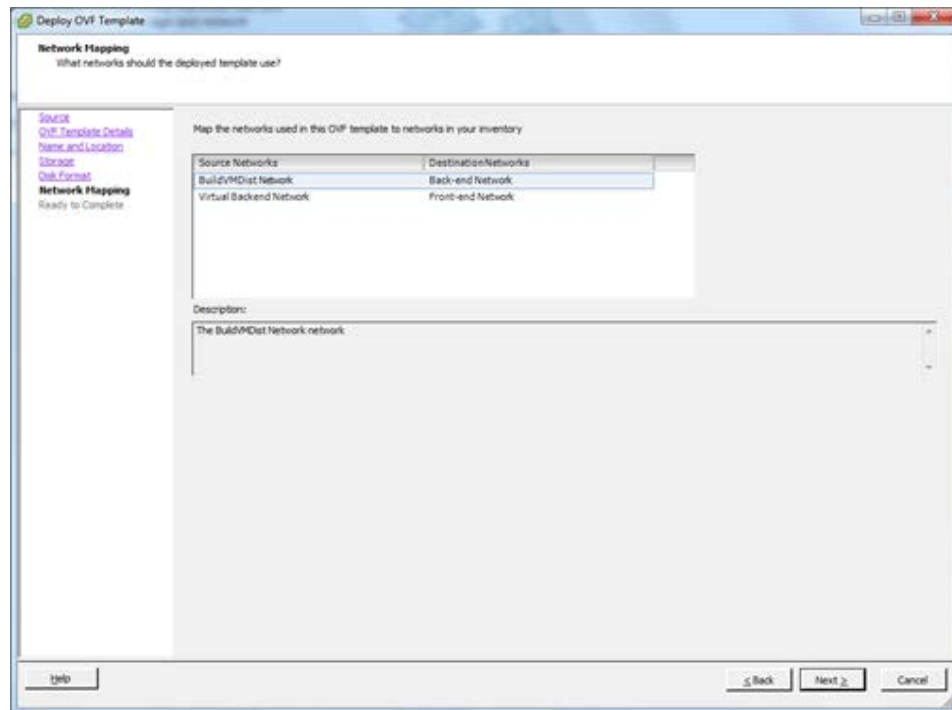
8. Select 1 datastore from the list you previously added to the ESXi hosts. If you're doing a consecutive load, make sure to select the next datastore down (from the previous load) on the list. The selected datastore should have a capacity of 50GB.
9. Click **Next**.



10. Verify that the datastore you selected matches the Available space size expected for the datastore.
11. Select **Thick Provision Eager Zeroed**.
12. Click **Next**.

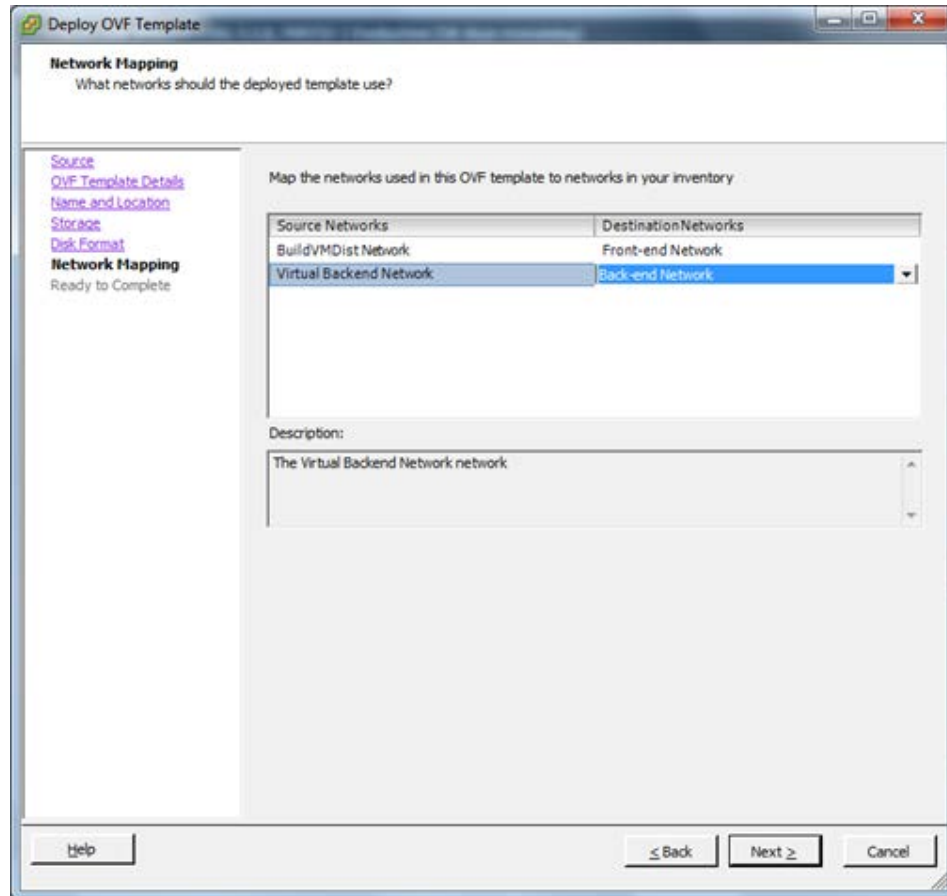


- 13.** Hover your cursor over a Destination network to make a drop down menu button appear. Click on the drop down menu for Destination Networks.
- 14.** Change the Destination Networks so that the Front-end Network lines up with the BuildVMDisk Network.
- 15.** Change the Destination Networks so that the Back-end Network lines up with the Virtual Back-end Network.



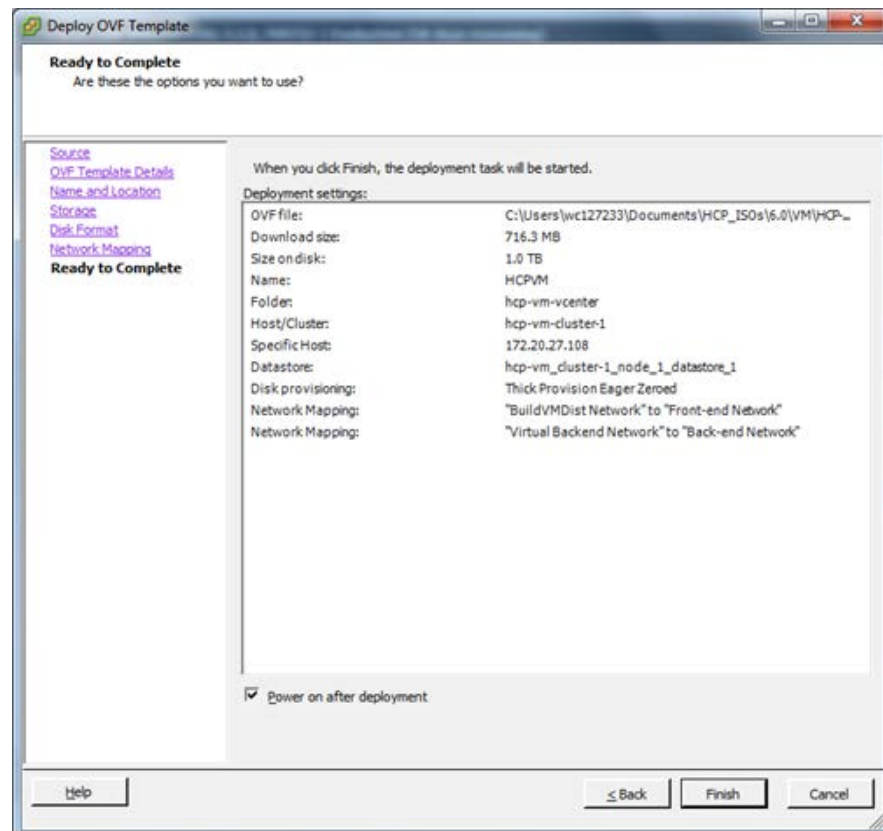
16. Verify your Destination Networks are set the same way as in the image below.

17. Click **Next**.



Important: Do **NOT** select the Power on checkbox.

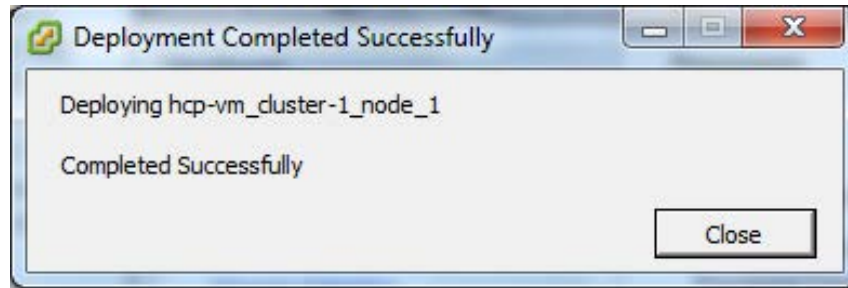
- 18.** Verify the information in Deployment settings matches what was previously entered:
 - a.** If so click on **Finish** to begin the OVF deploy.
 - b.** If not, go back and correct any information that needs to be changed.



Important:

- You must repeat the OVF deployment for each of the nodes that are going to be part of the HCP-VM system.
- Make sure that you have highlighted the desired ESXi host that you want the HCP-VM node to run on initially before importing the OVF.

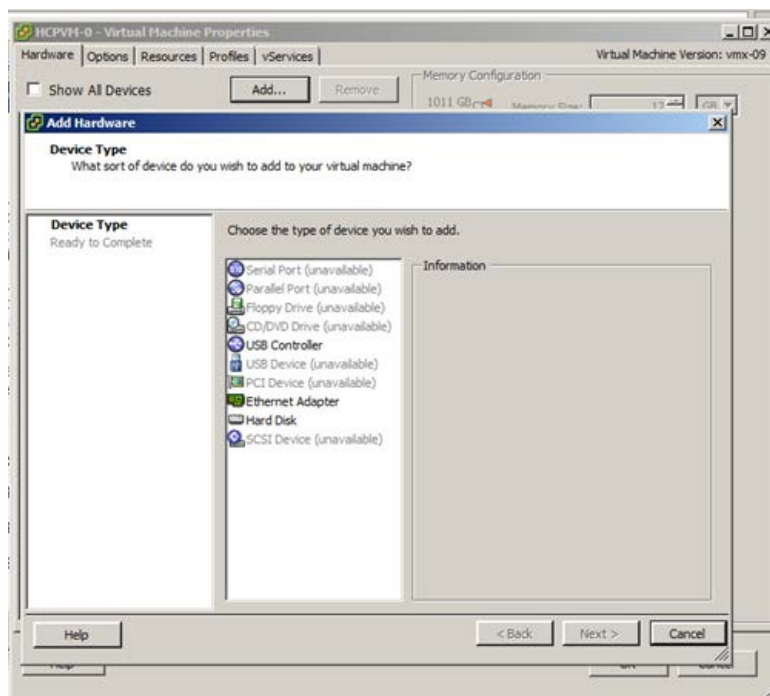
Once the OVF Deploy is completed, you will see the following message.



Step 3: Complete the Deployment

To complete the deployment:

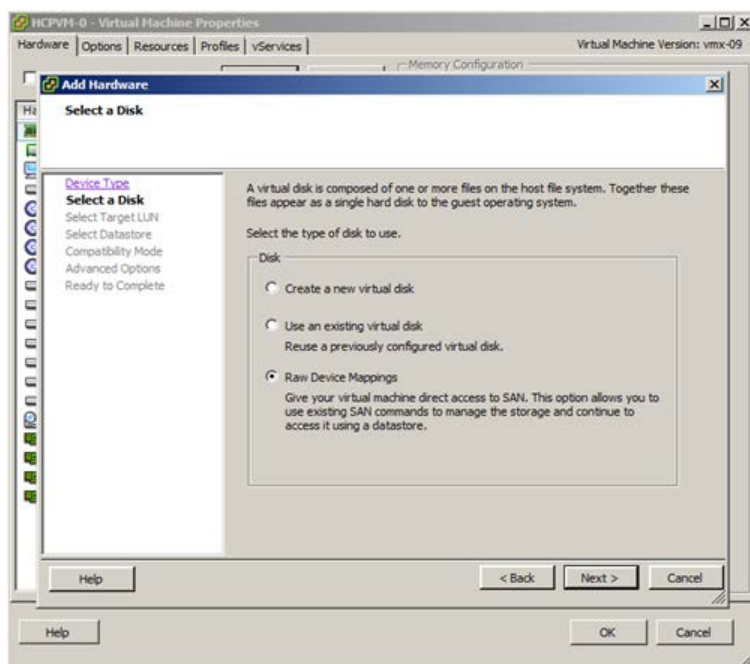
1. After the OVFs have deployed successfully, in the left side navigation bar right click on the HCP-VM and select **Edit Settings**.
2. In the **Settings** window, click **Add**.
3. In the **Add Hardware** window, select **Hard Disk**.
4. Click **Next**.



Deploy the HCP-VM OVF RDM

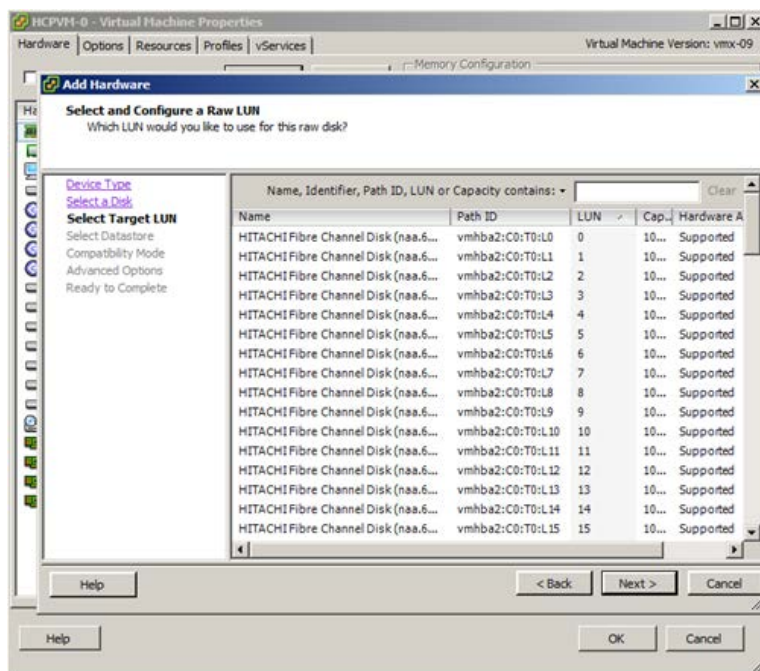
5. Select **Raw Device Mapping**.

6. Click **Next**.



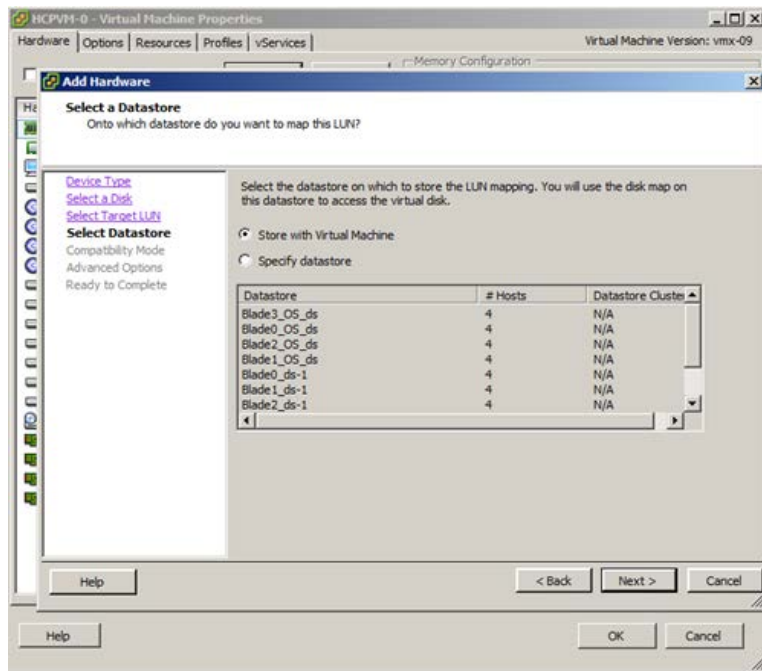
7. Select the desired LUN.

8. Click **Next**.



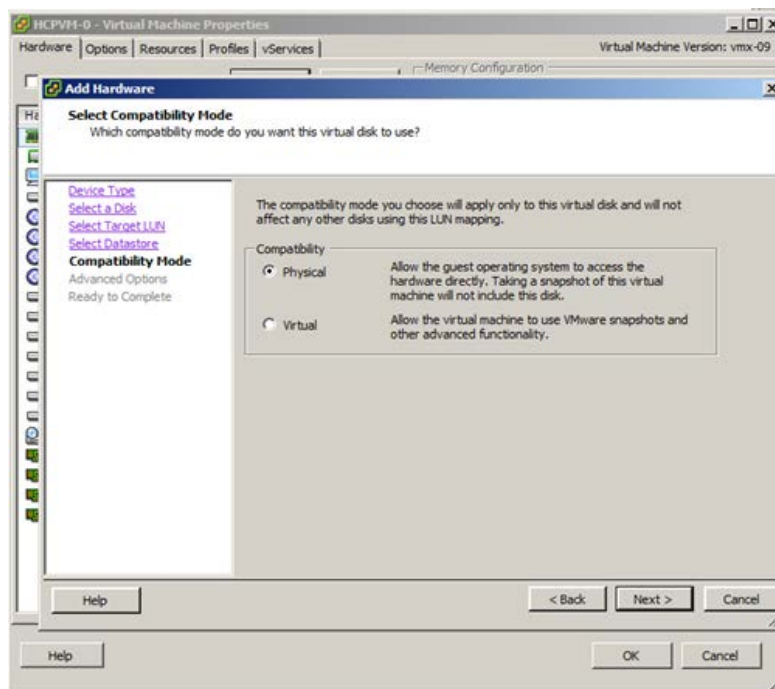
9. Select **Store with Virtual Machine**.

10. Click **Next**.



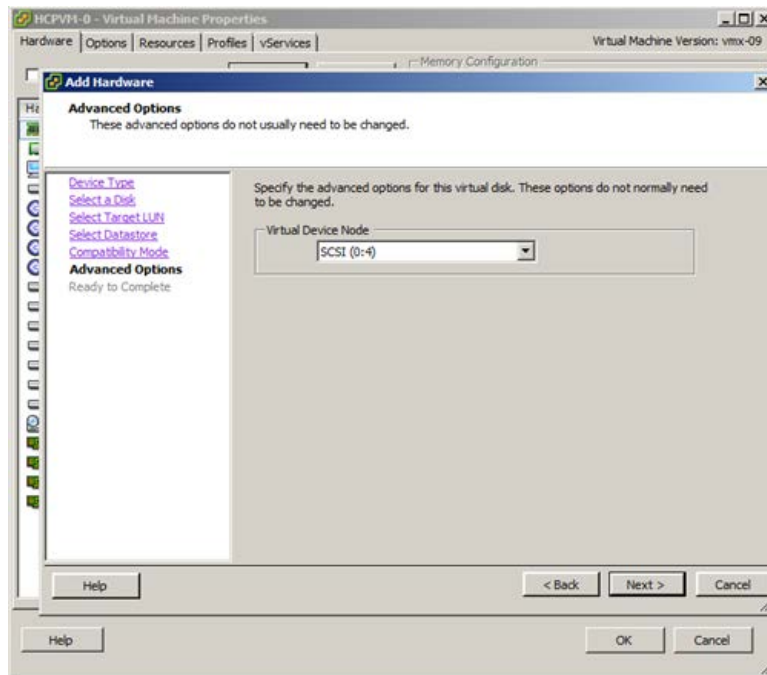
11. Select **Physical**.

12. Click **Next**.



13. Select the next SCSI device.

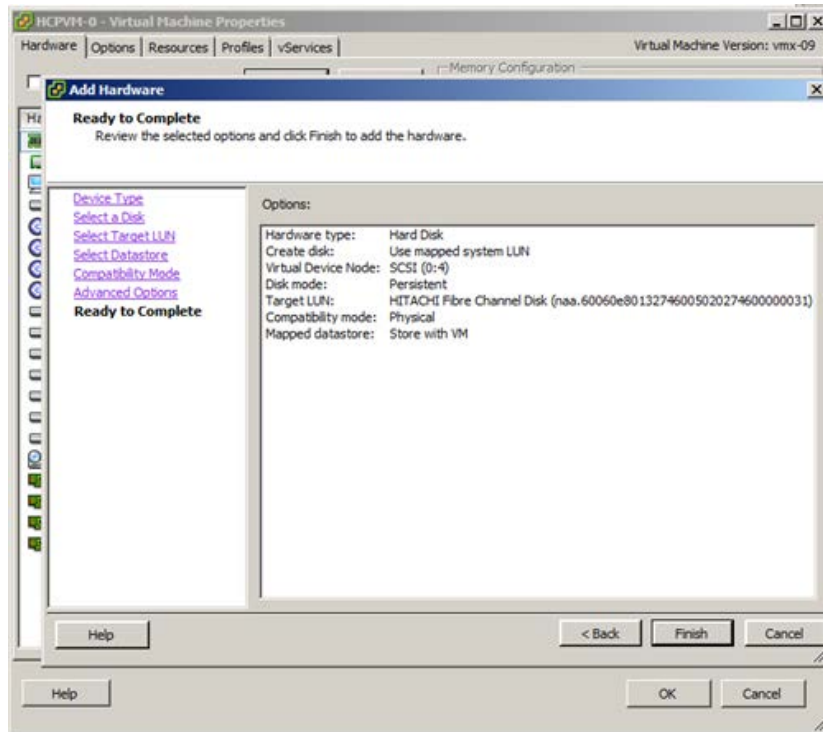
14. Click **Next**.



15. Review your actions and click **Finish**.
16. Repeat the steps to add a second data LUN.



Important: Do **NOT** power on the HCP-VM yet.



Configuring the HCP-VM small instance

If you are deploying this 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 **OK** to save your changes and close the **Virtual Machine Properties** window.

Configuring the HCP-VM network

After deploying the OVF, 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.

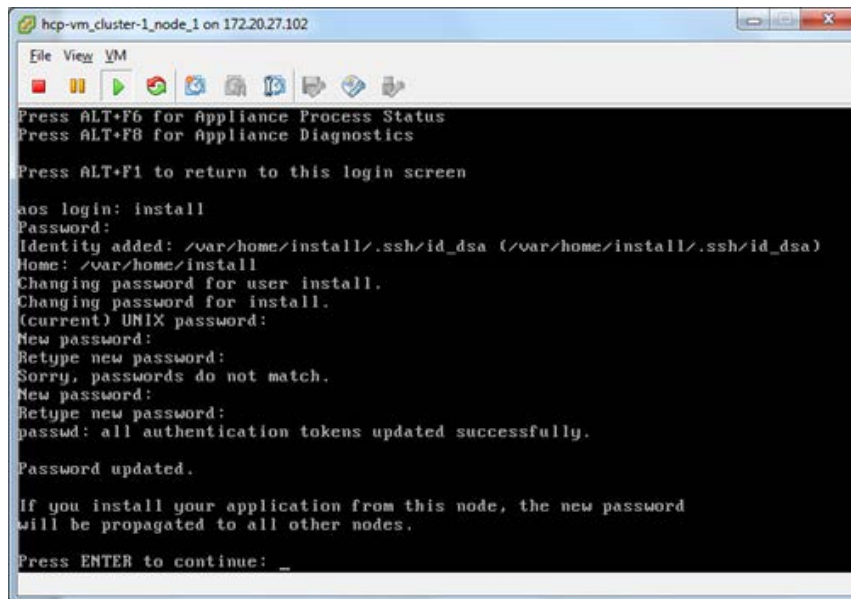
To configure the HCP-VM network:

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



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

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



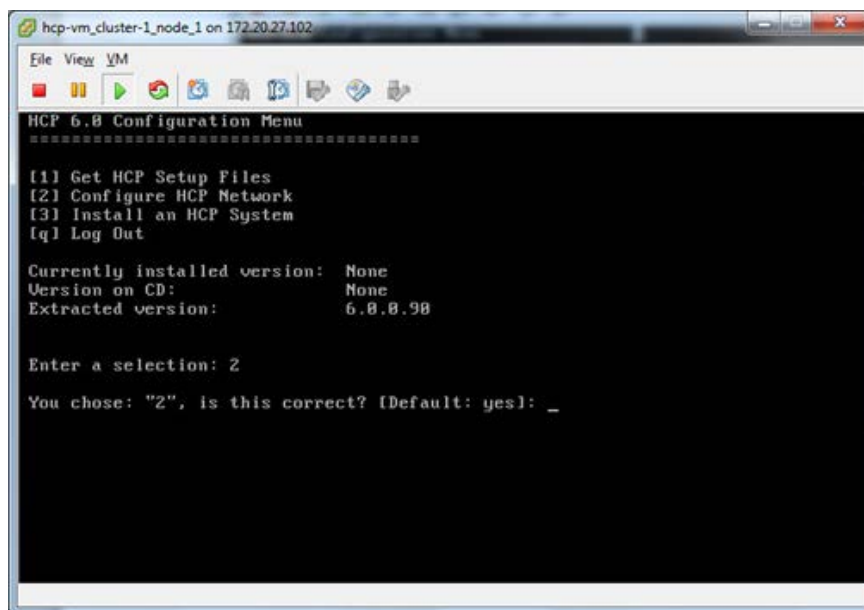
```
hcp-vm_cluster-1_node_1 on 172.20.27.102
File View VM
Press ALT+F6 for Appliance Process Status
Press ALT+F8 for Appliance Diagnostics
Press ALT+F1 to return to this login screen
aos login: install
Password:
Identity added: /var/home/install/.ssh/id_dsa (/var/home/install/.ssh/id_dsa)
Home: /var/home/install
Changing password for user install.
Changing password for install.
(current) UNIX password:
New password:
Retype new password:
Sorry, passwords do not match.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.

Password updated.

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

Press ENTER to continue: _
```

5. Enter 2 to access the **Configure HCP Network** menu.



```
hcp-vm_cluster-1_node_1 on 172.20.27.102
File View VM
HCP 6.8 Configuration Menu
=====
[1] Get HCP Setup Files
[2] Configure HCP Network
[3] Install an HCP System
[q] Log Out

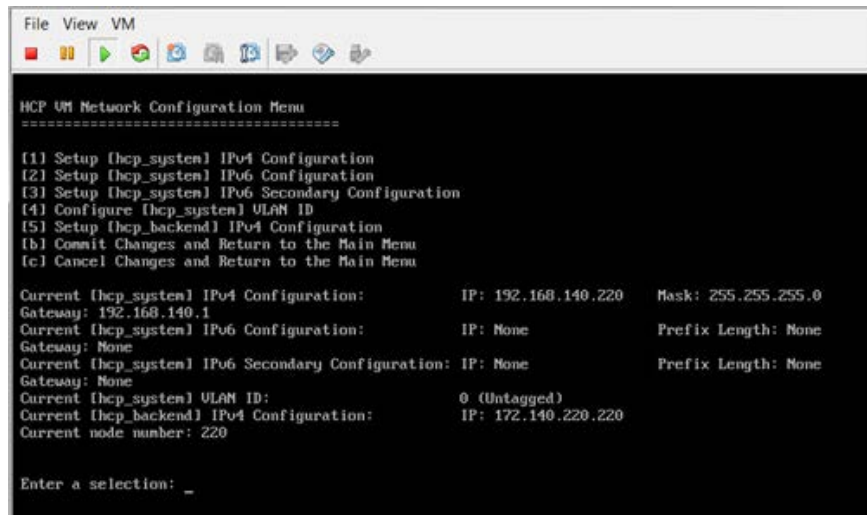
Currently installed version: None
Version on CD: None
Extracted version: 6.8.8.98

Enter a selection: 2

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

6. Update options 1 and 5 with information provided by the customer.

7. Ignore option 4 unless the customer wants to deploy with VLAN support turned on. See appendix A for configuring the ESXi Networking to support this.



```

File View VM
HCP VM Network Configuration Menu
=====
[1] Setup [hcp_system] IPv4 Configuration
[2] Setup [hcp_system] IPv6 Configuration
[3] Setup [hcp_system] IPv6 Secondary Configuration
[4] Configure [hcp_system] VLAN ID
[5] Setup [hcp_backend] IPv4 Configuration
[b] Commit Changes and Return to the Main Menu
[c] Cancel Changes and Return to the Main Menu

Current [hcp_system] IPv4 Configuration:      IP: 192.168.140.220      Mask: 255.255.255.0
Gateway: 192.168.140.1
Current [hcp_system] IPv6 Configuration:      IP: None                Prefix Length: None
Gateway: None
Current [hcp_system] IPv6 Secondary Configuration: IP: None                Prefix Length: None
Gateway: None
Current [hcp_system] VLAN ID:                 0 (Untagged)
Current [hcp_backend] IPv4 Configuration:      IP: 172.140.220.220
Current node number: 220

Enter a selection: _

```

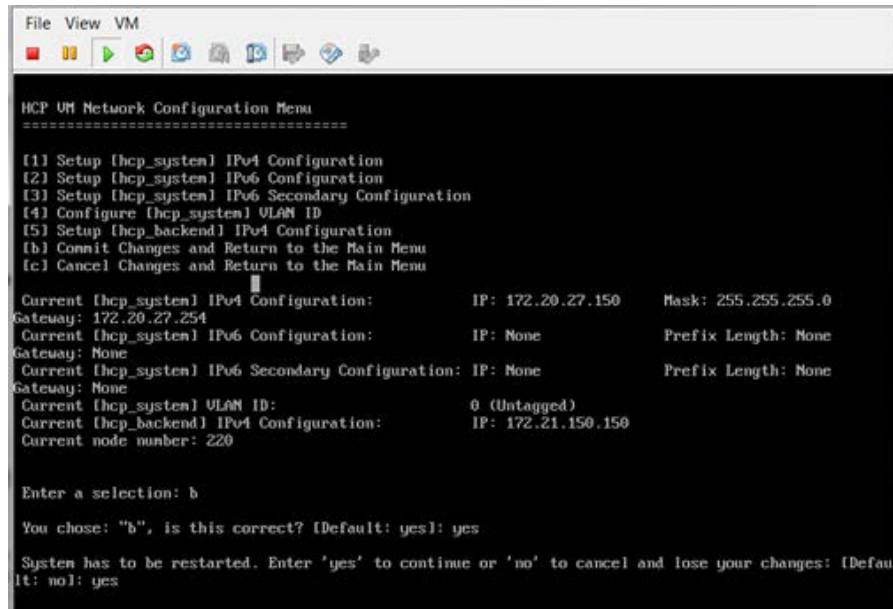
8. For the example system, the following was changed and is reflected in the next image:

- Front-end IP: *172.20.27.150*
- Gateway address: *172.20.27.254*
- Back-end IP: *172.21.150.150*



Note: For configuring separate clusters, if you use similar Back-end IPs the third octet has to be unique, otherwise the nodes will communicate across clusters.

9. Confirm the information and enter *B* to commit the changes.



```

File View VM
=====
[1] Setup [hcp_system] IPv4 Configuration
[2] Setup [hcp_system] IPv6 Configuration
[3] Setup [hcp_system] IPv6 Secondary Configuration
[4] Configure [hcp_system] VLAN ID
[5] Setup [hcp_backend] IPv4 Configuration
[b] Commit Changes and Return to the Main Menu
[c] Cancel Changes and Return to the Main Menu

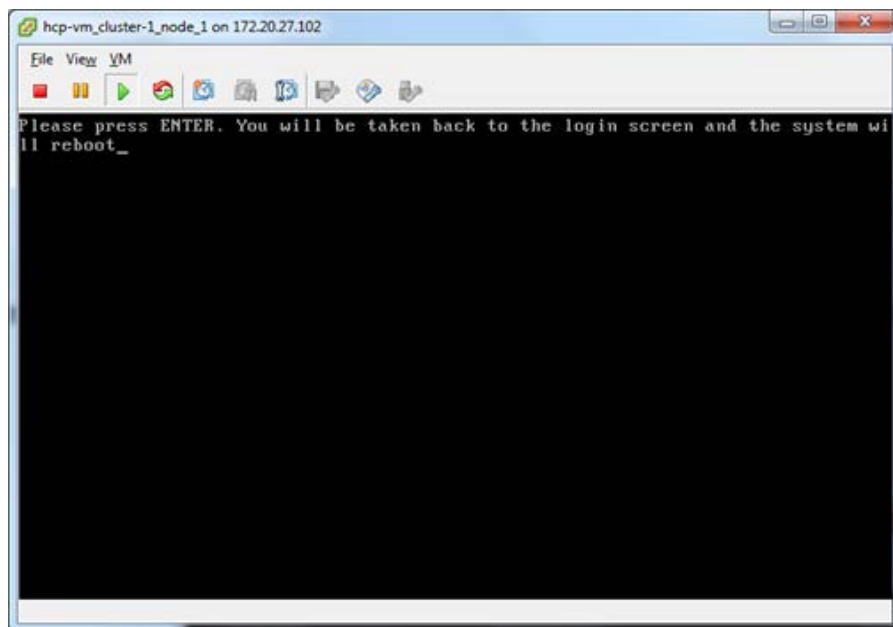
Current [hcp_system] IPv4 Configuration:      IP: 172.20.27.150      Mask: 255.255.255.0
Gateway: 172.20.27.254
Current [hcp_system] IPv6 Configuration:      IP: None              Prefix Length: None
Gateway: None
Current [hcp_system] IPv6 Secondary Configuration: IP: None      Prefix Length: None
Gateway: None
Current [hcp_system] VLAN ID:                  0 (Untagged)
Current [hcp_backend] IPv4 Configuration:      IP: 172.21.150.150
Current node number: 220

Enter a selection: b

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

System has to be restarted. Enter 'yes' to continue or 'no' to cancel and lose your changes: [Default: no]: yes
  
```

10. Press enter to reboot the HCP-VM node.



```

hcp-vm_cluster-1_node_1 on 172.20.27.102
File View VM
=====
Please press ENTER. You will be taken back to the login screen and the system will reboot_
  
```


11. The HCP-VM node will begin to reboot. Do not touch it until the reboot is complete.



Note: The previous steps must be completed for each VM you set up.

12. Once the HCP-VM node finishes rebooting, login with the username and password:

- Username: *install*
- Password: *hcpinsta11*



Install HCP software

The HCP install is performed from the node with the highest last octet in its Back-end IP address. For example, the four Back-end IP addresses for the example system are:

- 172.21.150.150
- 172.21.150.151
- 172.21.150.152
- 172.21.150.153


So 172.21.150.153 is the proper node for the HCP software installation.



Note: While HCP software is customer installable, it is not possible to enable data at rest encryption (DARE) on customer installed systems. 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 a console under 172.20.27.153 (hcp-vm_cluster-1_node_4).
3. Right click on the console and click on **Open Console**.
4. Login with the username and password:
 - Username: *install*
 - Password: *hcpinsta11*
5. Enter 3.
6. Hit **Enter**.



Step 1: Identify the nodes in the HCP system

To identify the nodes in the HCP system:

1. From the **HCP 7.3.3 Configuration** menu, enter **3** to run the HCP Setup wizard.
2. In response to the confirming prompt, enter *y* or *yes* to confirm your entry or *n* or *no* to try again.

When you enter *y* or *yes*, the HCP Setup wizard **New Install** menu appears.

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

[1] HCP Nodes

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

[x] Install a New HCP System with This Configuration

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

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

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

The **HCP Nodes** menu appears.

```
HCP Nodes Menu
=====

[1] Storage Node Back-end IP Addresses

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

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

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



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

5. From the **HCP Nodes** menu, enter **b** to return to the **New Install** menu.

The **New Install** menu now includes additional options for configuring the HCP system.

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

[1] HCP Nodes
[2] Distributor/OEM Key Access (Arizona)
[3] Networking Settings
[4] DNS Settings
[5] Time Settings
[6] Internal Configuration Settings
[7] Encryption Settings

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

[x] Install a New HCP System with This Configuration

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

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

Configure the HCP system

From the **New Install** menu, 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.

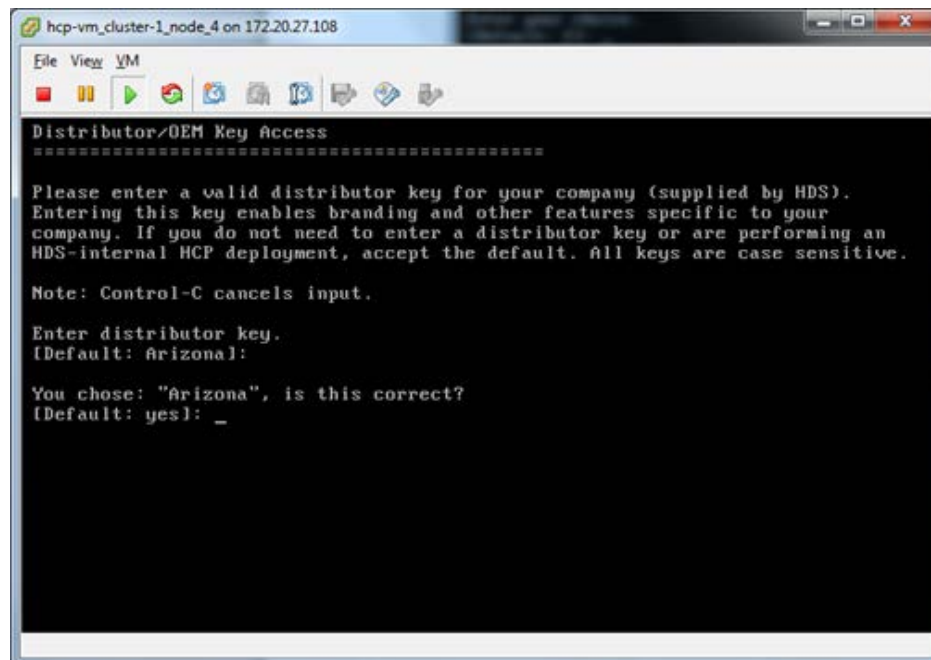
To configure the HCP system:

1. Enter 2 in the **New Install** menu to open the **Key Access** menu.
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.



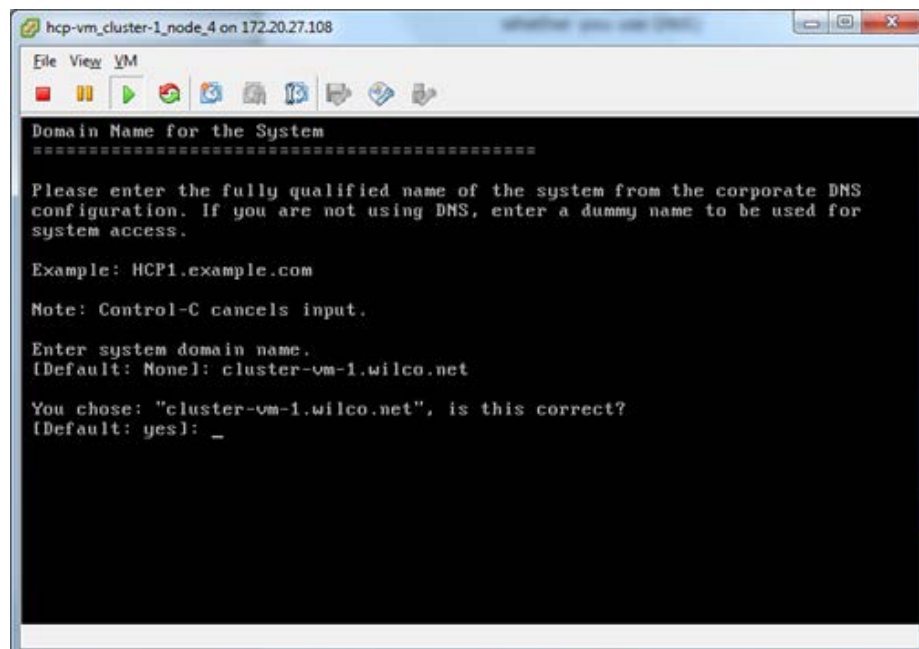
4. Enter 3 to open the **HCP Networking** menu.
5. Enter 1 and change the Gateway router IP address.
6. Enter 2 and change the Multicast Network.
7. Enter *b* to return to the **New Install** menu.



8. Enter 4 to open the **HCP DNS Options** menu.
9. Enter 2 to input the domain name for the system.




10. Enter the system domain name.

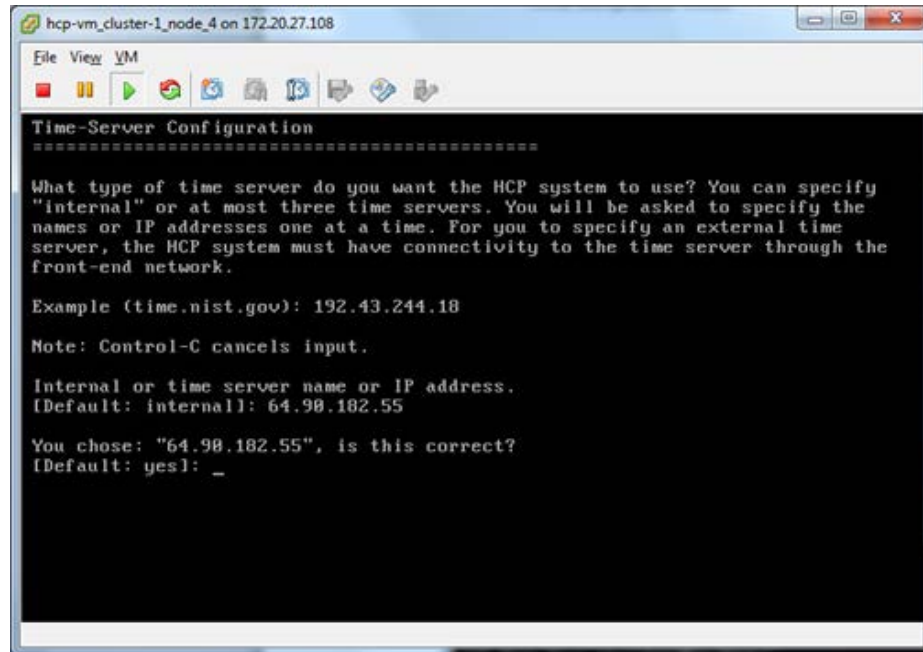


11. If **Option 1: Enable DNS** is not set to yes, change it to yes.
12. If **Option 3: DNS Servers** is not set to the proper corporate DNS server, change it accordingly.
13. Enter *b* to return to the **New Install** menu.
14. Enter *5* open the **HCP Time Options** menu.
15. Enter *1* and set the time configuration to an external 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 the ["Enabling NTP for the ESXi hosts"](#).




16. Enter an external time server.



17. Enter **6** to change the internal settings.

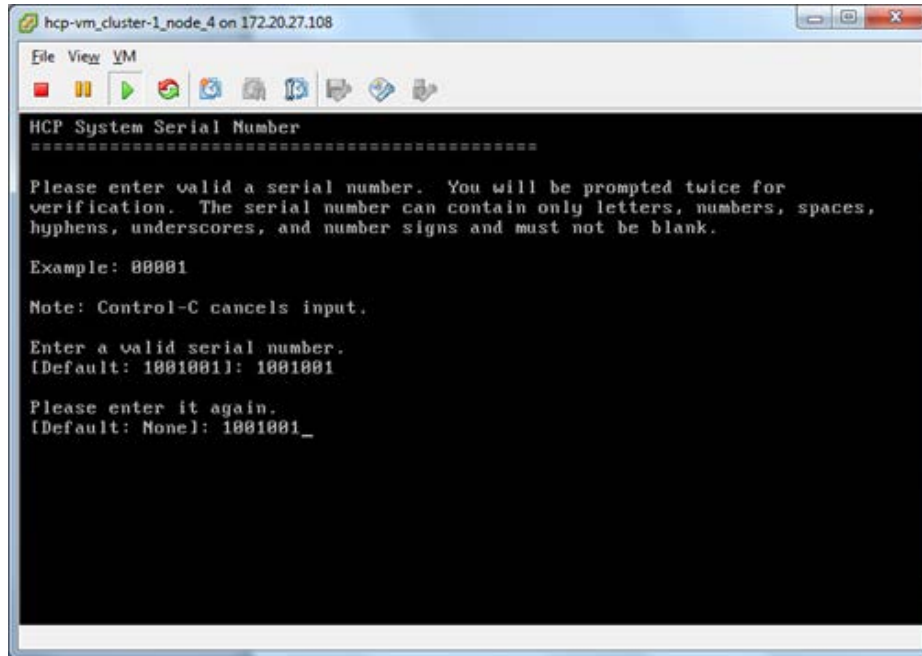
18. Select **Internal**.



19. Select the option to set the serial number for the system and enter the unique serial number for this HCP system.



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.



20. Decide if 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.

- Contact information for HCP customer support.

21. To specify no contact information, hit **space**.

22. If you want to enable encryption, contact your authorized service provider.

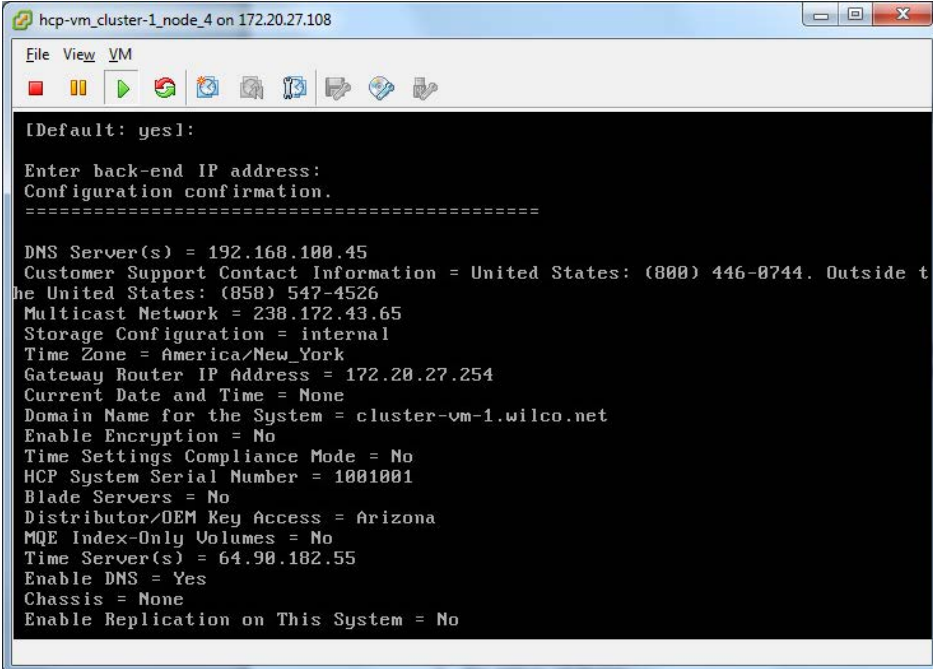
Execute the installation

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**.

The wizard should display the following information.



```

[Default: yes]:

Enter back-end IP address:
Configuration confirmation.
=====

DNS Server(s) = 192.168.100.45
Customer Support Contact Information = United States: (800) 446-0744. Outside t
he United States: (858) 547-4526
Multicast Network = 238.172.43.65
Storage Configuration = internal
Time Zone = America/New_York
Gateway Router IP Address = 172.20.27.254
Current Date and Time = None
Domain Name for the System = cluster-vm-1.wilco.net
Enable Encryption = No
Time Settings Compliance Mode = No
HCP System Serial Number = 1001001
Blade Servers = No
Distributor/DEM Key Access = Arizona
MQE Index-Only Volumes = No
Time Server(s) = 64.90.102.55
Enable DNS = Yes
Chassis = None
Enable Replication on This System = No
  
```

```

hcp-vm_cluster-1_node_4 on 172.20.27.108
File View VM
Gateway Router IP Address = 172.20.27.254
Current Date and Time = None
Domain Name for the System = cluster-vm-1.wilco.net
Enable Encryption = No
Time Settings Compliance Mode = No
HCP System Serial Number = 1001001
Blade Servers = No
Distributor/OEM Key Access = Arizona
MQE Index-Only Volumes = No
Time Server(s) = 64.90.182.55
Enable DNS = Yes
Chassis = None
Enable Replication on This System = No
Spindown Volumes = No
HCP Storage Nodes: 4
  172.21.150.150
  172.21.150.151
  172.21.150.152
  172.21.150.153
HCP Search Nodes: 0

Use SHIFT+PGUP to review the Configuration.

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

2. Review the configuration.
3. Perform one of the following steps:
 - a. If the configuration is not correct:
 1. Enter *n* or *no*.
 2. In response to the confirmation prompt, enter *y* or *yes*.
 3. Correct any mistakes you made in the previous sections.
 - b. If the configuration is correct:
 1. Enter *y* or *yes*.
 2. In response to the confirmation prompt, enter *y* or *yes*.

When you enter *y* or *yes*, HCP Setup performs a set of installation prechecks and, if they are successful, installs the HCP software on all nodes in the system. This can take from several minutes to several hours, depending on the size of the logical volumes.



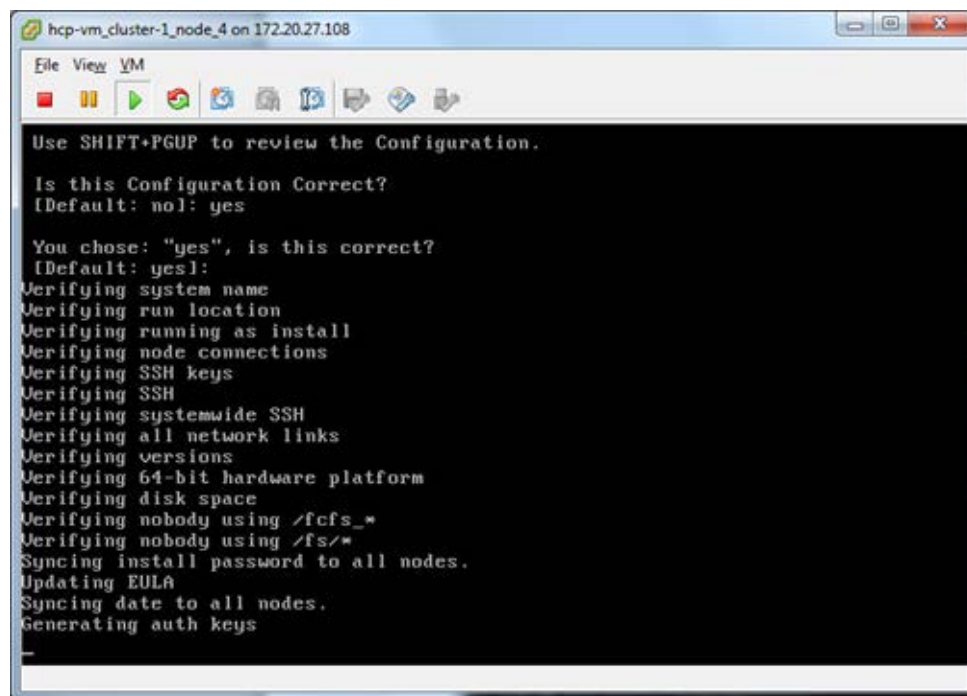
Important: If you enabled encryption in the system configuration, HCP Setup displays the encryption key after doing some initial setup. It then prompts you to enter the key. Before entering the encryption key, write it down on paper.

After you enter the key, HCP Setup proceeds with 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.

Beginning of HCP install. Pre-checks for system health.



The screenshot shows a terminal window titled "hcp-vm_cluster-1_node_4 on 172.20.27.108". The terminal displays the following text:

```

Use SHIFT+PGUP to review the Configuration.

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

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 all network links
Verifying versions
Verifying 64-bit hardware platform
Verifying disk space
Verifying nobody using /fcfs_*
Verifying nobody using /fs/*
Syncing install password to all nodes.
Updating EULA
Syncing date to all nodes.
Generating auth keys
_

```

HCP installation formatting the data LUNs.

```

node 150: 45% Complete (23/51): Formatting 36% complete (8 / 2 volumes)
node 151: 45% Complete (23/51): Formatting 36% complete (8 / 2 volumes)
node 152: 45% Complete (23/51): Formatting 37% complete (8 / 2 volumes)
node 153: 45% Complete (23/51): Formatting 36% complete (8 / 2 volumes)
Tue Feb 26 13:49:56 2013 Current status:
node 150: 45% Complete (23/51): Formatting 36% complete (8 / 2 volumes)
node 151: 45% Complete (23/51): Formatting 36% complete (8 / 2 volumes)
node 152: 45% Complete (23/51): Formatting 37% complete (8 / 2 volumes)
node 153: 45% Complete (23/51): Formatting 41% complete (8 / 2 volumes)
Tue Feb 26 13:50:01 2013 Current status:
node 150: 45% Complete (23/51): Formatting 41% complete (8 / 2 volumes)
node 151: 45% Complete (23/51): Formatting 41% complete (8 / 2 volumes)
node 152: 45% Complete (23/51): Formatting 41% complete (8 / 2 volumes)
node 153: 45% Complete (23/51): Formatting 41% complete (8 / 2 volumes)
Tue Feb 26 13:51:00 2013 Current status:
node 150: 45% Complete (23/51): Formatting 46% complete (8 / 2 volumes)
node 151: 45% Complete (23/51): Formatting 41% complete (8 / 2 volumes)
node 152: 45% Complete (23/51): Formatting 41% complete (8 / 2 volumes)
node 153: 45% Complete (23/51): Formatting 46% complete (8 / 2 volumes)
Tue Feb 26 13:51:05 2013 Current status:
node 150: 45% Complete (23/51): Formatting 46% complete (8 / 2 volumes)
node 151: 45% Complete (23/51): Formatting 46% complete (8 / 2 volumes)
node 152: 45% Complete (23/51): Formatting 46% complete (8 / 2 volumes)
node 153: 45% Complete (23/51): Formatting 46% complete (8 / 2 volumes)

```

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

After the reboot, you can also check the runlevel of the node by hitting 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-889-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

```

Verifying the HCP installation

Access the HCP System Management Console to verify that the HCP system 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 inadvertently use 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 this 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 processing, 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 the Update Password button.

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, mouse 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. Perform additional system configuration, as described in ["Setting additional configuration options"](#). Do this only if the installation was successful.
- a. Log out of the System Management Console, and close the browser window to ensure that no one can return to the Console without a logging in.

Setting additional configuration options

After verifying that the HCP system was correctly installed, you can perform additional system configurations. For example, you can 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're not already logged in).
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 to ensure that no one can return to the Console without logging in.

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

Monitoring and alerting

The HCP hardware based appliance has built in redundant hardware, monitoring, alerting and failover behavior that cannot be leveraged in a virtualized VMware environment. To maintain performance and data integrity, all underlying hardware associated with the HCP-VM system must be treated as mission critical and monitored for failures. Whenever Hitachi servers, storage, and networking are part of the HCP-VM system, they must be connected to HiTrack. Any non-Hitachi equipment should be closely monitored using the vendor or customer equivalent to HiTrack. 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 lengthy RAID rebuild times.

Monitoring and alerts

In general, HCP-VM can be managed like other HCP platforms except when monitoring the physical environment and monitoring the VMware environment. Here are some of the differences between HCP-VM monitoring and other platforms:

- HCP-VM System hardware monitoring is the responsibility of the customer and should be treated with the utmost priority and importance.
- HCP IPMI monitoring is not available in the HCP-VM environment.
- Storage is not restricted to Hitachi arrays. Array health monitoring and maintenance is the responsibility of the customer.

Software monitoring

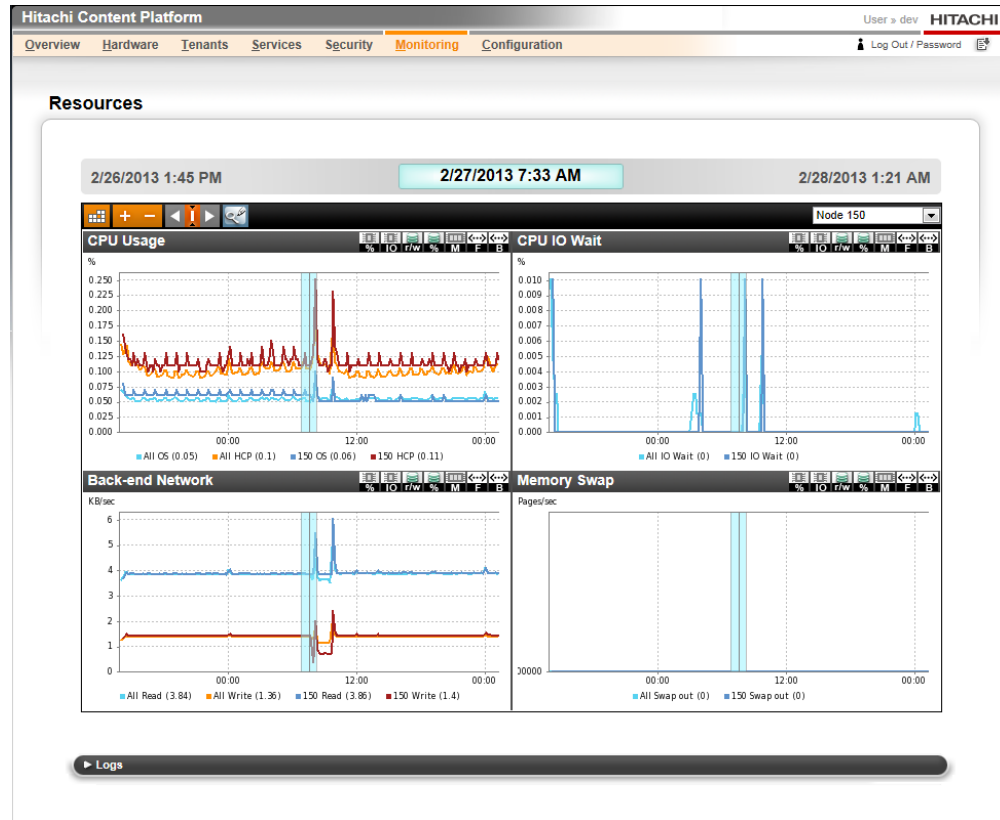
HCP maintains a system log which logs all events that happen within the system. You can view this log in the HCP **System Management Console**. You can send system log messages to syslog servers, System Network Management Protocol (**SNMP**) managers, and/or email addresses. Additionally, you can use 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 via HiTrack.

HCP 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 representing 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 diagnostic menu

For any HCP 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.

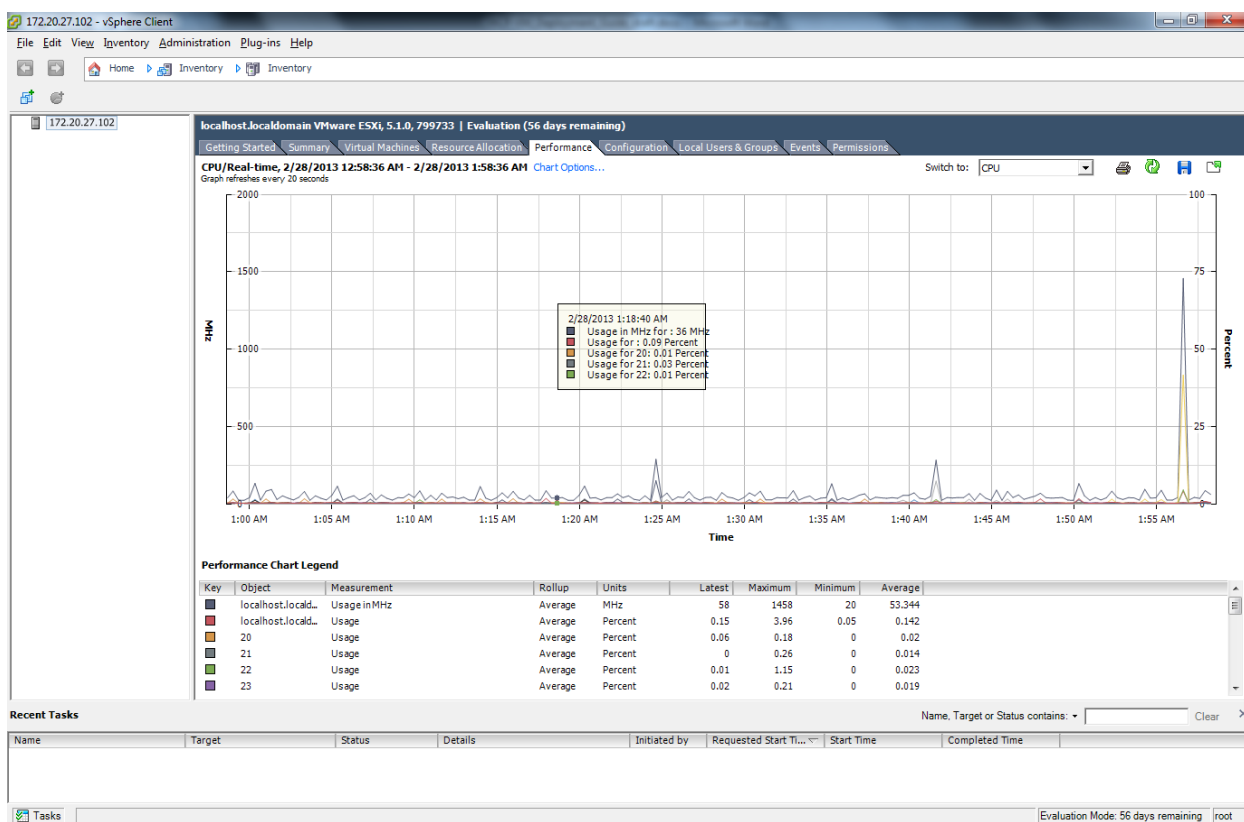
The diagnostics let you:

- **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.

More details about HCP system monitoring facilities can be found in the *Administering HCP* manual.

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.



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>

Setting the HCP-VM network adapters

This chapter covers changing between the supported e1000 and VMXNET3 network adapters on your HCP-VMs.

About network adapters

HCP supports two types of network adapters: e1000 and VMXNET3. 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.

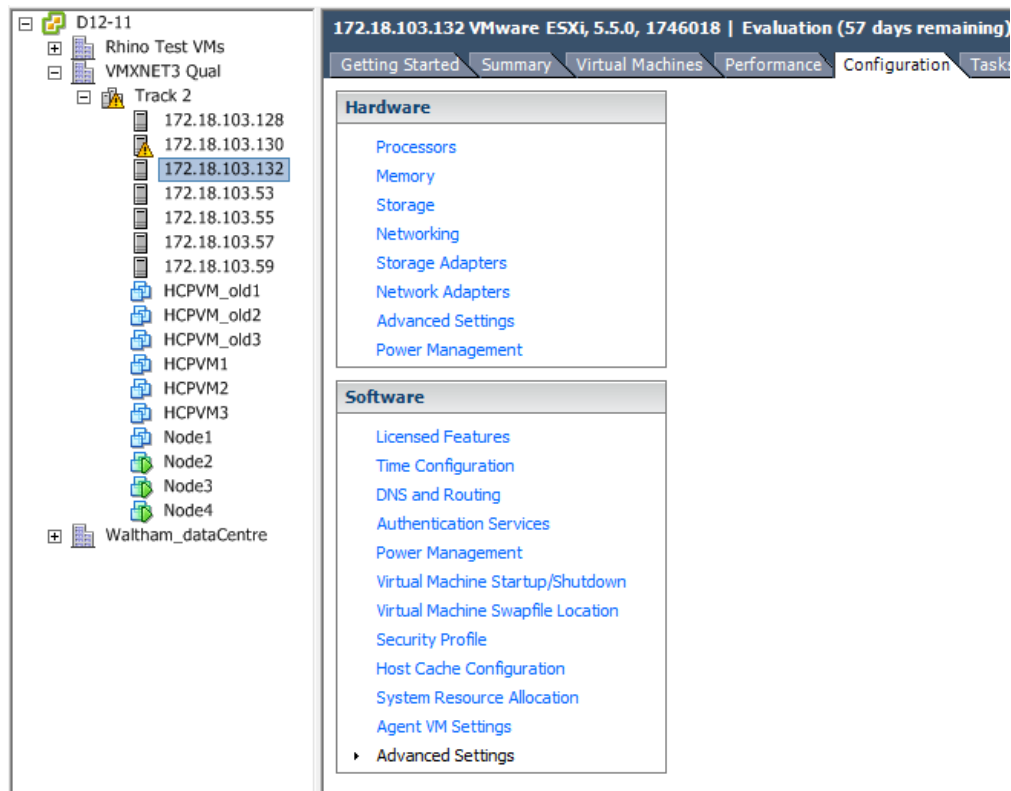
e1000 network adapters only support one gigabyte network configurations. VMXNET3 support both one and 10 gigabyte network configurations. If you have a 10 gigabyte network configuration it's recommended that you use VMXNET3 network adapters.

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:

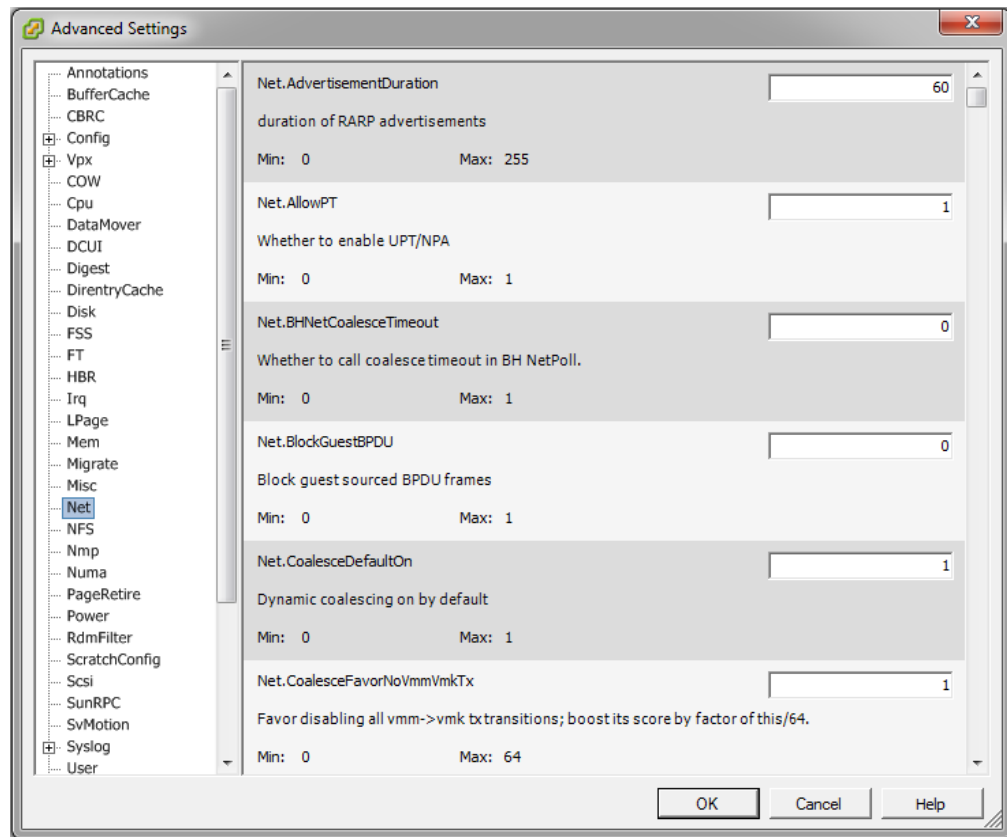
1. Log in to the vSphere Client.
2. Click on a server that hosts ESXi for your HCP-VMs.
3. Click on the **Configuration** tab.



4. In the **Software** panel, click on **Advanced Settings**.

The **Advanced Settings** window appears.


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



- 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
```

- Click **Ok**.
- From the vSphere client, reboot the server by right clicking on it, and in the dropdown menu clicking **Reboot**.



Setting 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, or you can remove the existing VMXNET3 network adapters and replace them with e1000 adapters. 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**.



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:

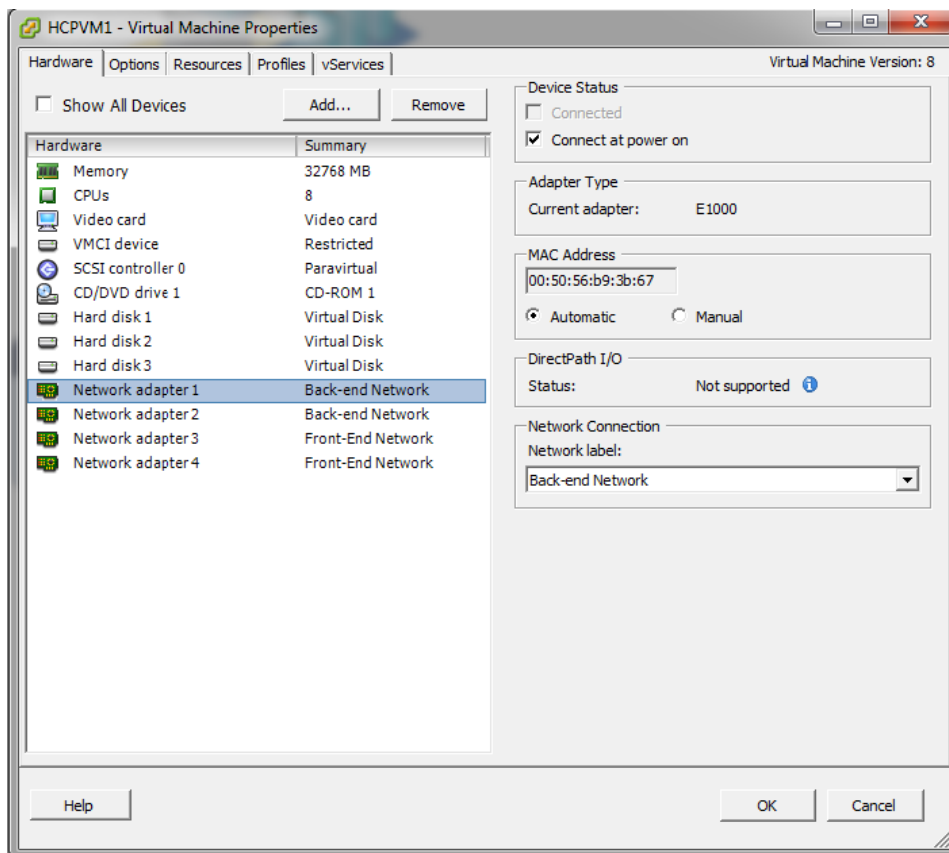
1. From the vSphere client, right click on one of the powered off HCP-VMs.

A dropdown menu appears.

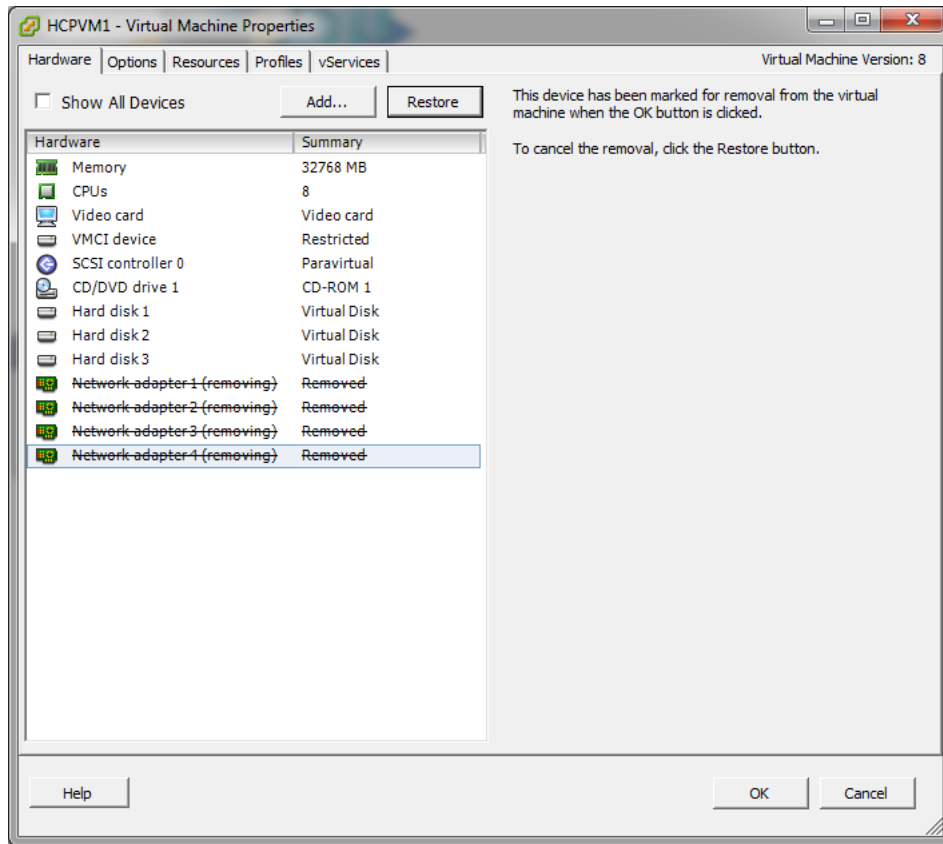
2. In the drop down menu that appears, click on **Edit Settings....**

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



4. Click on **OK**.

Step 3: Change the guest OS

Once the old network adapters are removed, you need to change the guest OS. To change the guest OS:

1. From the vSphere client, right click on your 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. From the **Edit Settings** window, click on the **Options** tab.

4. In the **Options** tab **Guest Operating System** panel, click on the **Version:** field and from the drop down menu select **Enterprise Linux 6 (64 bit)**. If **Enterprise Linux 6 (64 bit)** is already selected, leave it unchanged.



5. Click on **OK**.

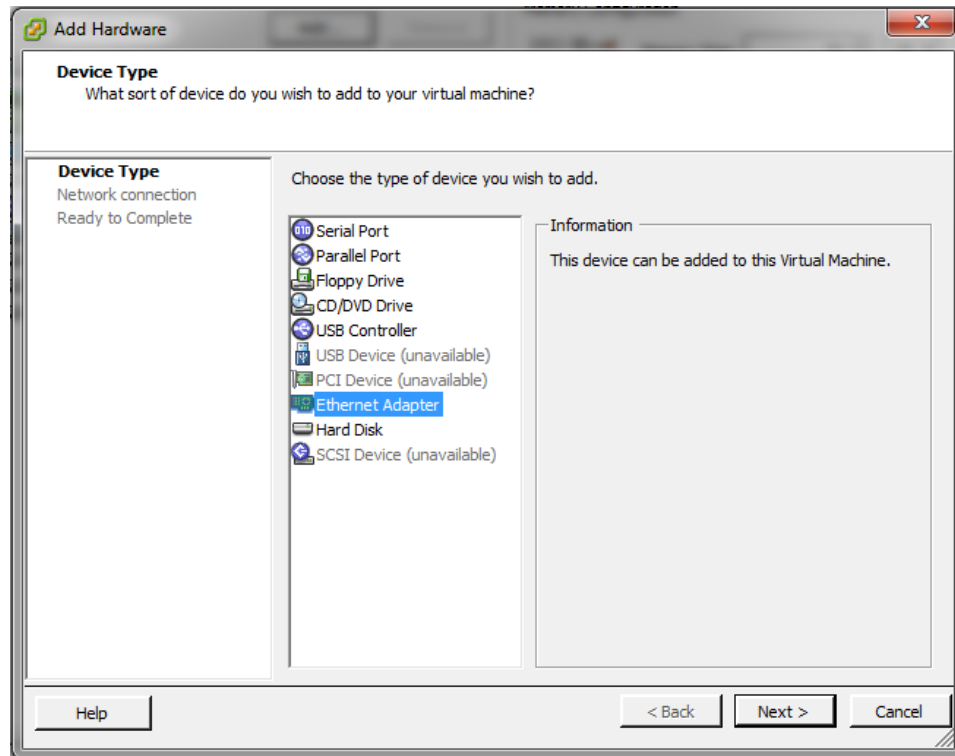
Step 4: Set the Front-End network adapters

Once the HCP-VM guest OS is configured, you need to set 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 **Next**.
6. On the **Network Connection** page in the **Adapter Type** panel **Type** dropdown field, take one of the following actions:
 - If you're switching to VMXNET3, select **VMXNET3**
 - If you're switching to e1000, select **e1000**
7. In the **Network connection** panel, select **Named network with specified label**.
8. Select **Front-End Network**.



9. Click on **Next**.

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

11. Back on the **Virtual Machine Properties** window, click **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 5: "Set the Back-End network adapters"](#) below

The first and third e1000 network adapters need to be connected to the Front-end network. The second and fourth e1000 network adapters need to be connect to the Back-end network. For more information on connecting the second and fourth e1000 network adapters to the Back-end network, see [Step 5: "Set the Back-End network adapters"](#) below

Step 5: 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.

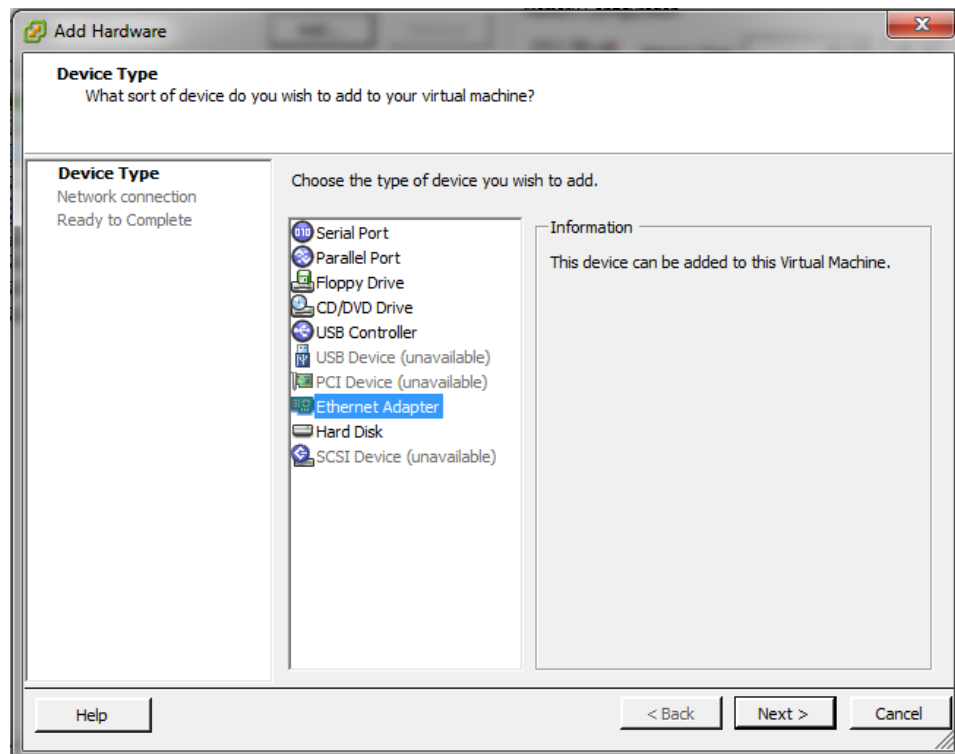
2. In the drop down menu that appears, click on **Edit Settings...**

A **Virtual Machine Properties** window appears.

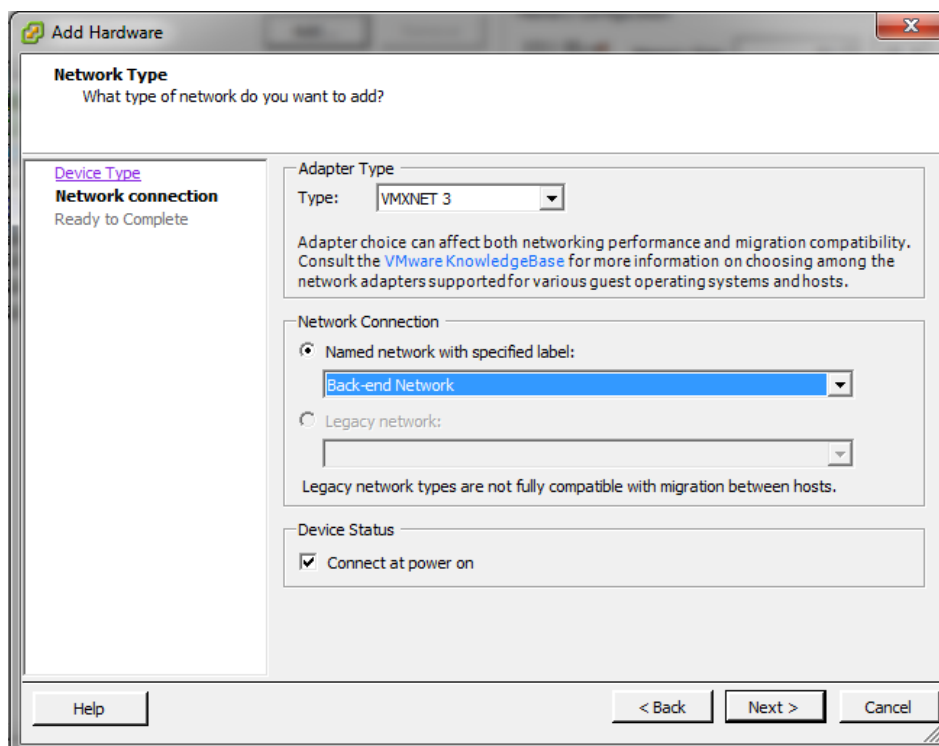
3. In the **Edit Settings** window **Hardware** tab, click on **Add**.

An **Add Hardware** window opens.

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



5. Click **Next**.
6. On the **Network Connection** page in the **Adapter Type** panel **Type** dropdown field, take one of the following actions:
 - If you're switching to VMXNET3, select **VMXNET3**
 - If you're switching to e1000, select **e1000**
7. In the **Network connection** panel, select **Named network with specified label**.
8. Select **Back-End Network**.



9. Click on **Next**.

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

11. Back on the **Virtual Machine Properties** window, click **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 4: "Set the Front-End network adapters"](#) on page 140

The first and third e1000 network adapters need to be connected to the Front-end network. The second and fourth e1000 network adapters need to be connect to the Back-end network. For more information on connecting the first and third e1000 network adapter to the Front-end network, see the previous step [Step 4: "Set the Front-End network adapters"](#) on page 140

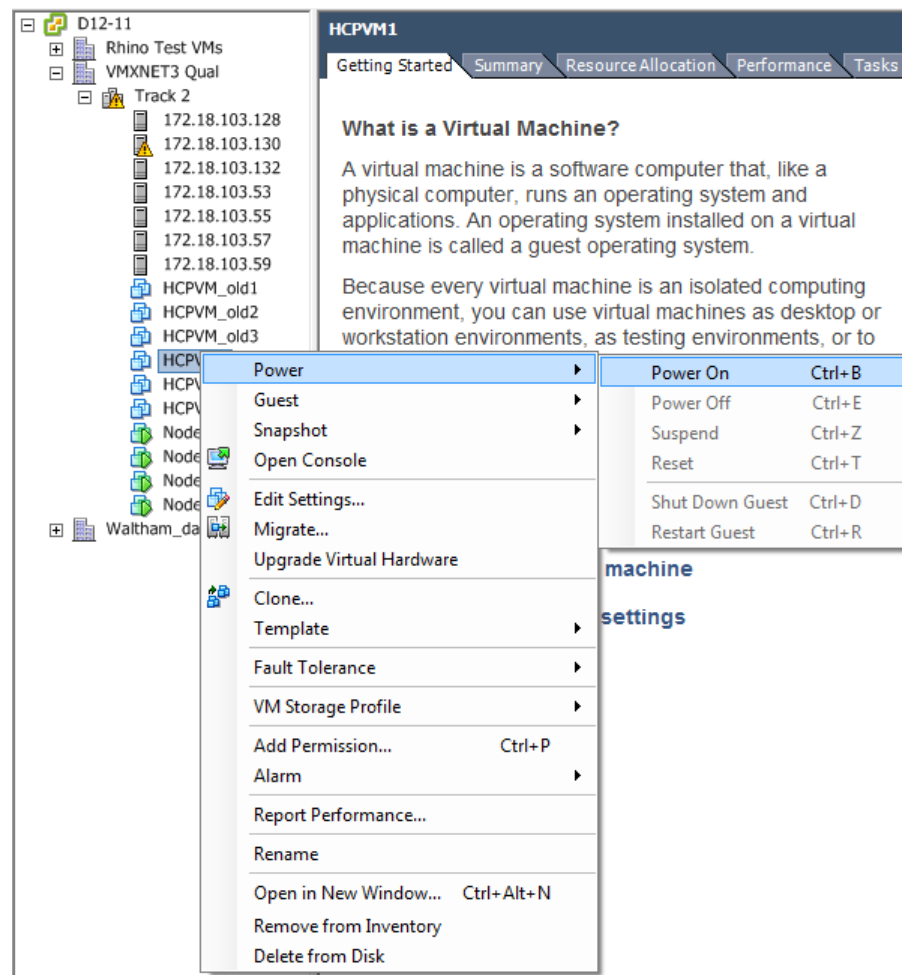
Power on the HCP-VM

Once you HCP-VM network adapters are configured, 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 drop down menu appears.

2. In the drop down 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.

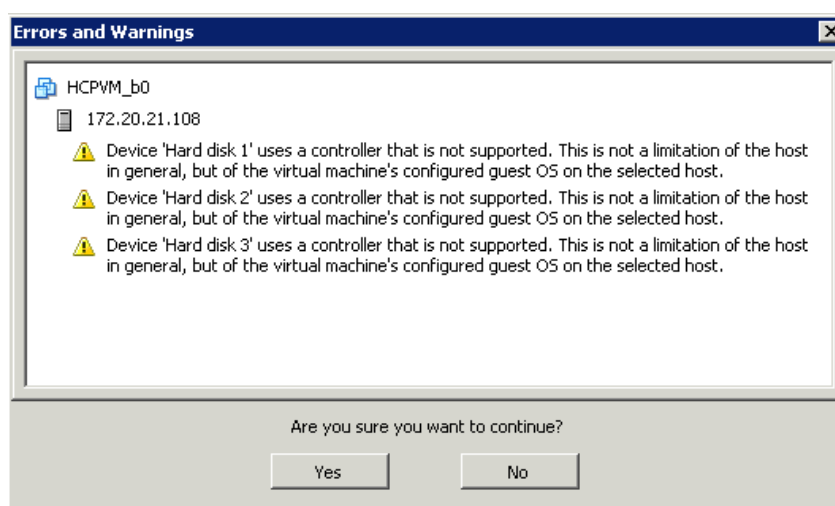
Failover management

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 the following error which can be safely ignored:



Maintenance procedures

The following sections outline ways to keep 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"](#), provision the LUNs to be added to each ESXi host in the system.
2. As described in ["Add datastores to vSphere HA cluster"](#), add the LUNs to new datastores.



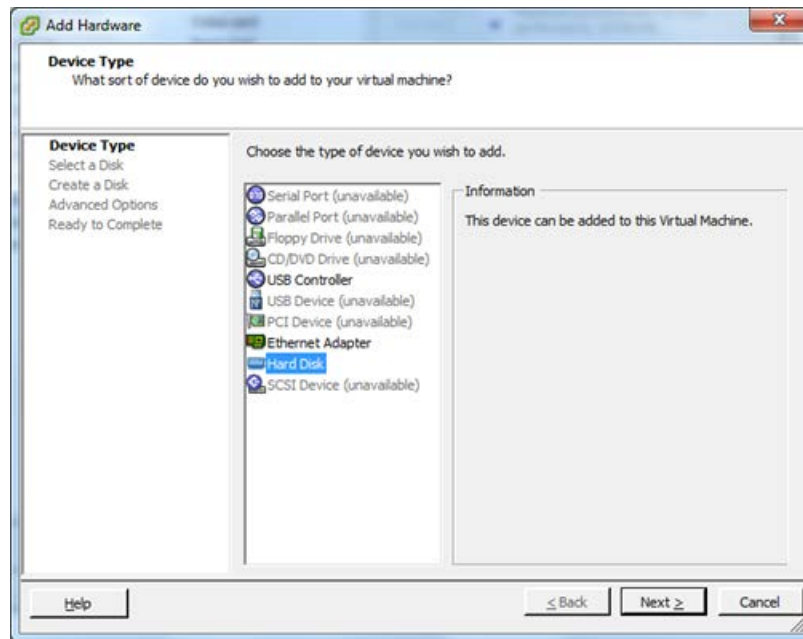
Important: This must be one LUN per datastore.

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 **Add**.



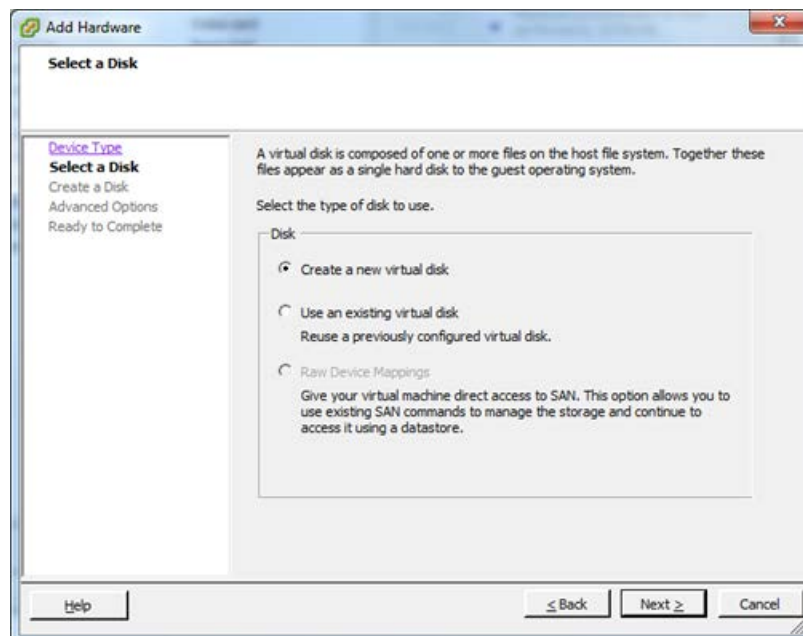
5. In the **Add Hardware** window, select **Hard Disk**.

6. Click **Next**.



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

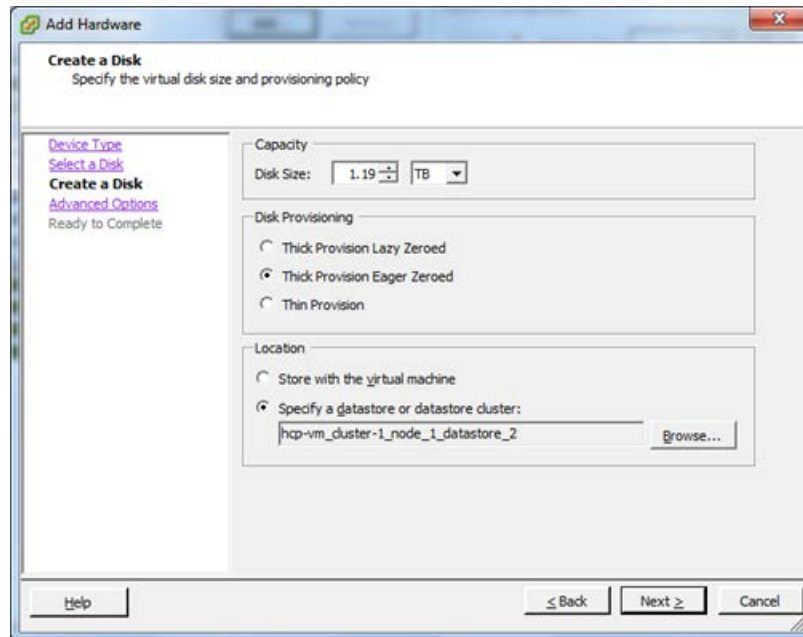
8. Click **Next**.



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

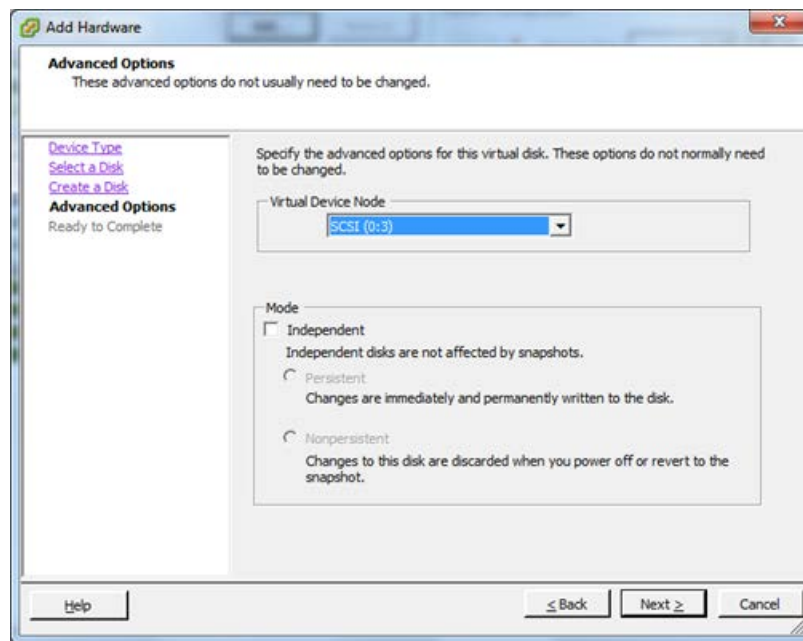
In the following example, the size of the LUN provisioned was 1.2TB.

- 10.** Select **Thick Provision Eager Zeroed**.
- 11.** Browse to the new datastore that will be added.
- 12.** Click **Next**.



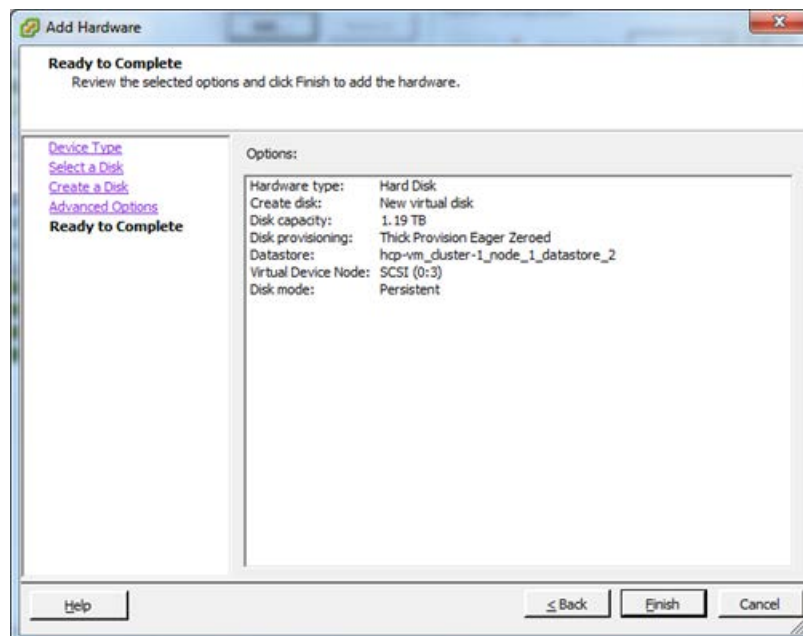
- 13.** Select the next available SCSI disk in the Virtual Device node section.

14. Click **Next**.



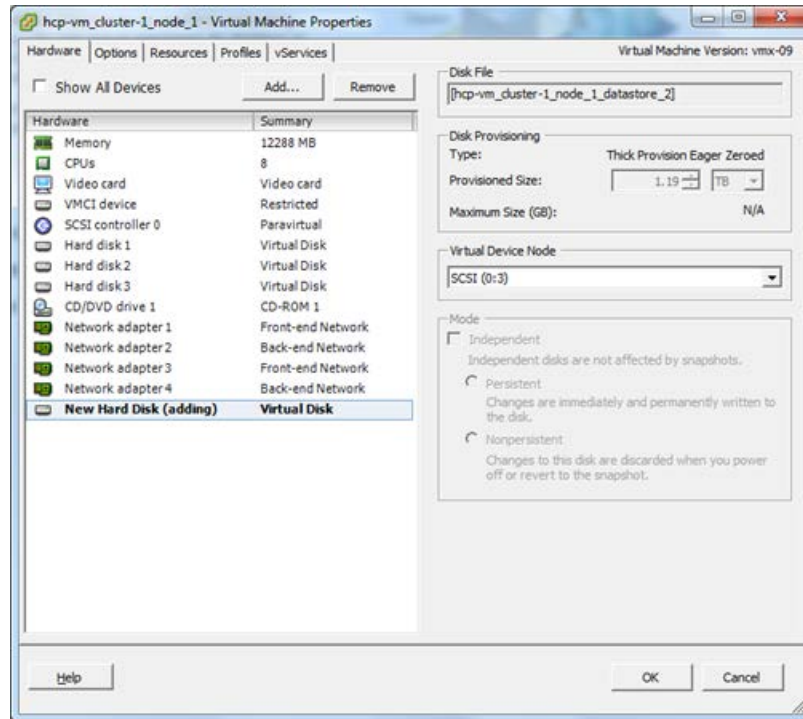
15. Verify the options selected.

16. Click **Finish**.



17. Back in the **Virtual Machine Properties** window, verify that the new Hard Disk listed.

18. Click **OK**.

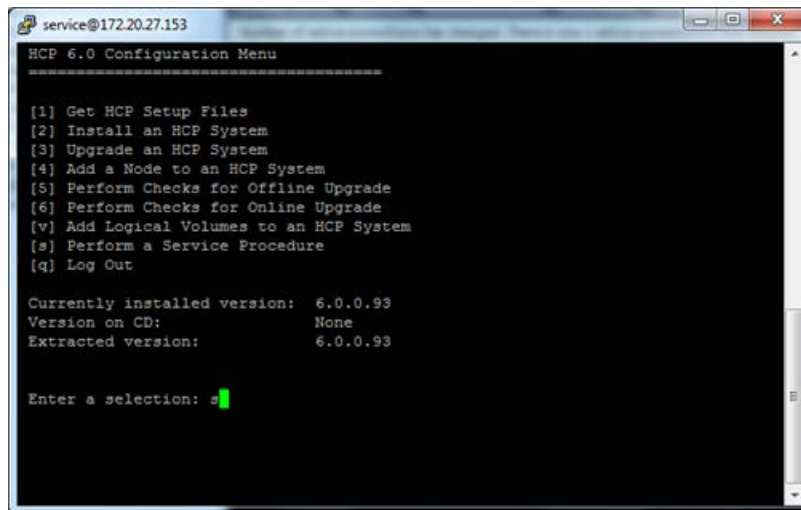


19. Using a tool, like PuTTY, load the appropriate service user ssh keys into your system.

20. SSH as the service user to the node with the highest IP.

21. Enter this command to open the HCP install shell:

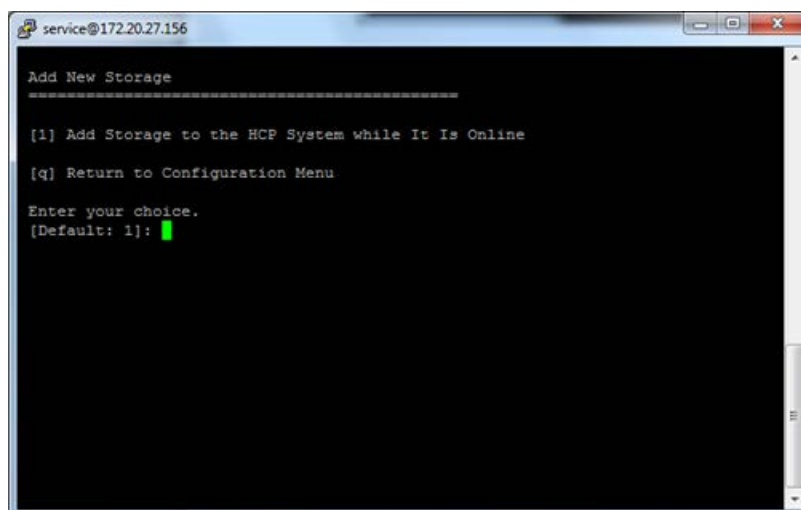
- `/home/service/bin/install`

22. Enter *v* to Add Logical Volumes to an HCP System.

```
service@172.20.27.153
HCP 6.0 Configuration Menu
=====
[1] Get HCP Setup Files
[2] Install an HCP System
[3] Upgrade an HCP System
[4] Add a Node to an HCP System
[5] Perform Checks for Offline Upgrade
[6] Perform Checks for Online Upgrade
[v] Add Logical Volumes to an HCP System
[s] Perform a Service Procedure
[q] Log Out

Currently installed version: 6.0.0.93
Version on CD: None
Extracted version: 6.0.0.93

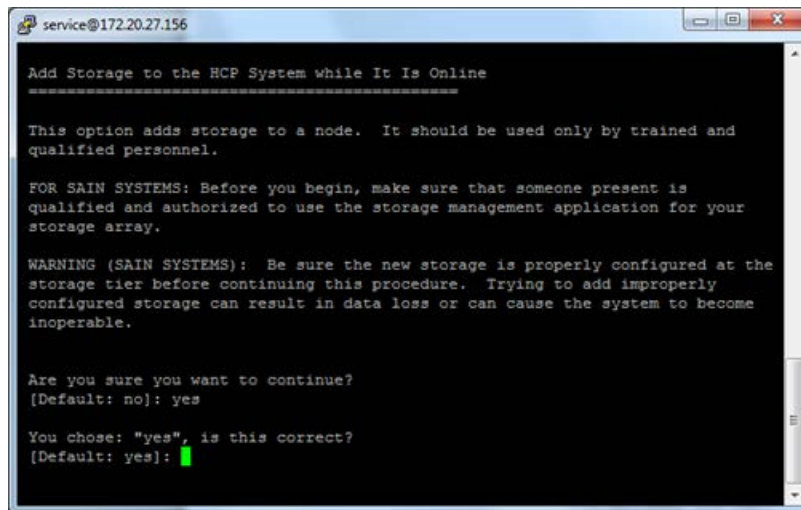
Enter a selection: s
```

23. Enter *1* to Add Storage to the HCP System while it is online.

```
service@172.20.27.156
Add New Storage
=====
[1] Add Storage to the HCP System while It Is Online
[q] Return to Configuration Menu

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

24. Enter **y** to verify that you want to add storage.



```

service@172.20.27.156
Add Storage to the HCP System while It Is Online
=====

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

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

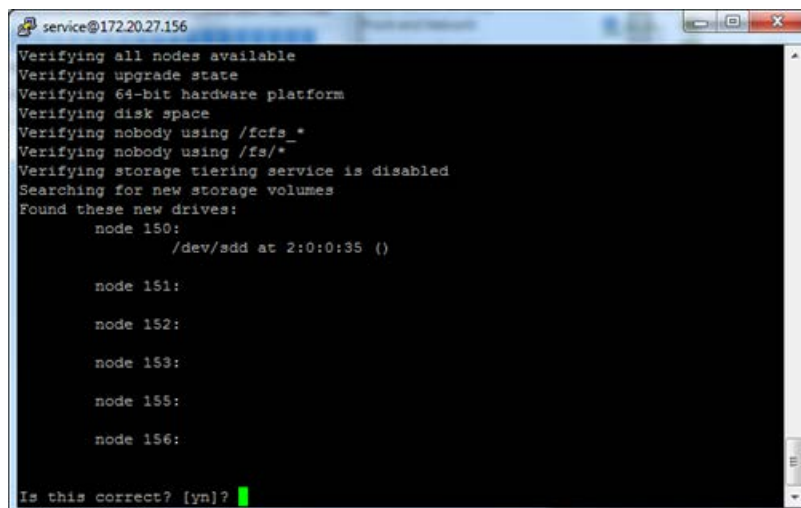
WARNING (SAIN SYSTEMS): Be sure the new storage is properly configured at the
storage tier before continuing this procedure.  Trying to add improperly
configured storage can result in data loss or can cause the system to become
inoperable.

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

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

25. Verify the new devices that were found. Typically this would show storage added to all nodes.

26. Enter **Y**.



```

service@172.20.27.156
Verifying all nodes available
Verifying upgrade state
Verifying 64-bit hardware platform
Verifying disk space
Verifying nobody using /fcfs_*
Verifying nobody using /fs/*
Verifying storage tiering service is disabled
Searching for new storage volumes
Found these new drives:
  node 150:
    /dev/sdd at 2:0:0:35 ()

  node 151:

  node 152:

  node 153:

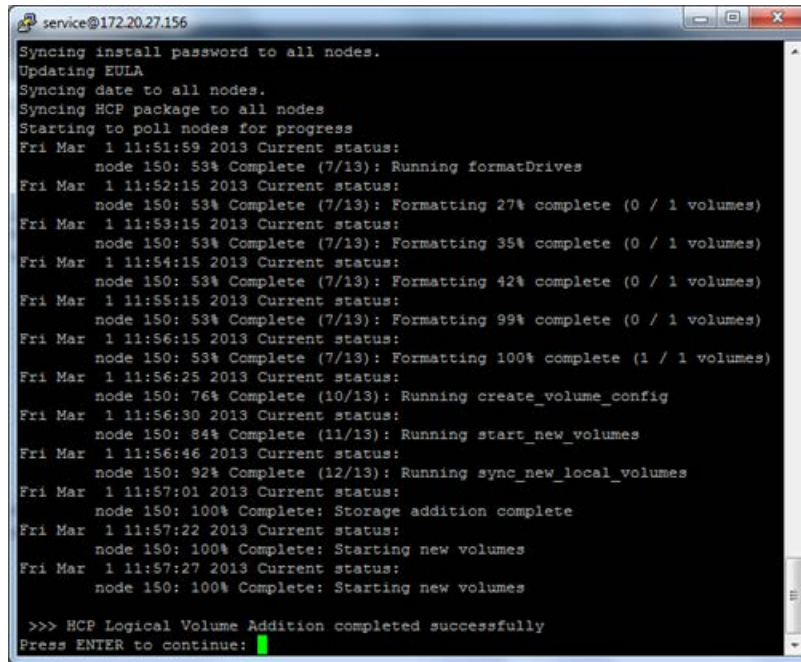
  node 155:

  node 156:

Is this correct? [yn]?
  
```

27. The new LUN will format for use by HCP.
28. When complete, hit **Enter** to continue.

29. Log into the HCP system management console to verify the newly added disks.



```

service@172.20.27.156
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:

```

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 27.
2. Deploy the OVF on the selected ESXi hosts. For more information about deploying the OVF see, ["Deploying the HCP-VM OVF VDMK"](#) on page 79 or ["Deploy the HCP-VM OVF RDM"](#) on page 91.
3. Change the network information on the newly deployed HCP-VM nodes. For more information about changing network information see, ["Configuring the HCP-VM network"](#) on page 108.
4. From the highest active HCP-VM node, run the add node service procedure found in the *Installing and Maintaining an HCP System* manual.

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.

The 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 in the System Management Console, mouse over **Monitoring** to display a secondary menu.
3. In the secondary menu, click on **SNMP**.
4. 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.
5. Click on the **Update Settings** button.
6. 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.
7. 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 161
4	Identify the HCP system to be monitored.	Step 4: "Identify the HCP system" on page 162

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.

Configuring networking for HCP virtual network management

Networks should be configured for particular switches in the system before the OVF is deployed.

HCP networking information

Make sure to review Administering HCP for the latest information on Network Administration in HCP.

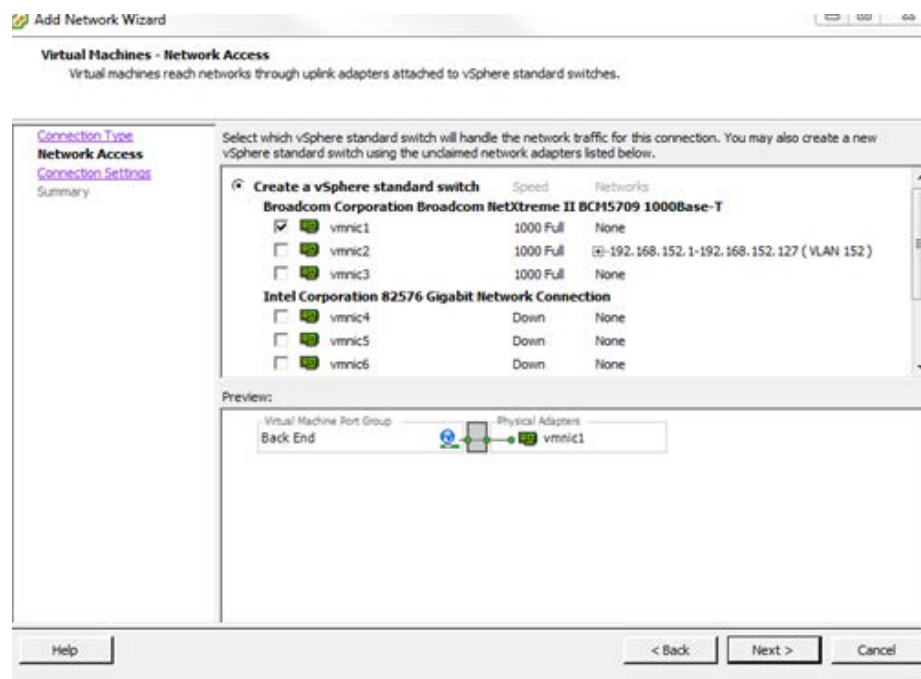
Configure networking for Back-end switching

To configure the network for Back-end switching:

1. Access the vSphere Client.
2. In the left side navigation window, select an ESXi host.
3. In the right side window, click on the **Configuration** tab.
4. Click on **Networking** under the **Hardware** section.
5. In the top right section of the right side window, Click on **Add Networking**.
6. In the **Add Network Wizard**, select **Virtual Machine**.
7. Click **Next**.

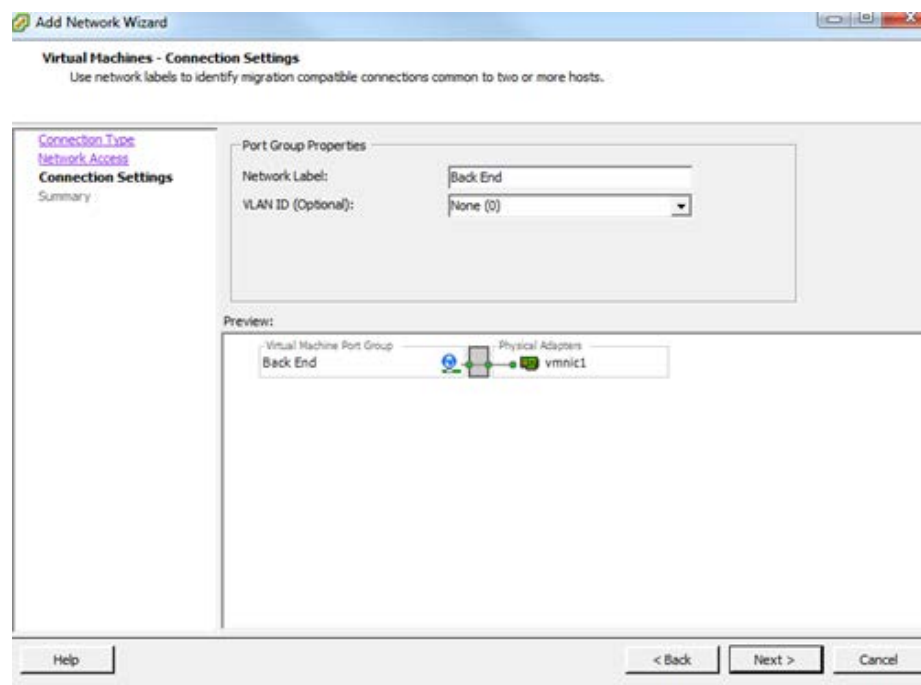
8. Select the target Physical NIC for the Back-end.

9. Click **Next**.



10. Enter a label for Back-end.

11. Click **Next**.



12. Review the newly configured switches.

13. Click **Finish**.


Configure the networking for Front-end switching

To configure the network for front-end switching:

- 1.** From the vSphere Client, in the left side navigation window, select an ESXi host.
- 2.** In the right side window, click on the **Configuration** tab.
- 3.** Click on **Networking** under the **Hardware** section.
- 4.** In the top right section of the right side window, click on **Add Networking**.
- 5.** In the **Add Network Wizard**, select **Virtual Machine**.
- 6.** Click **Next**.
- 7.** Select the target Physical NIC for the Front-end.

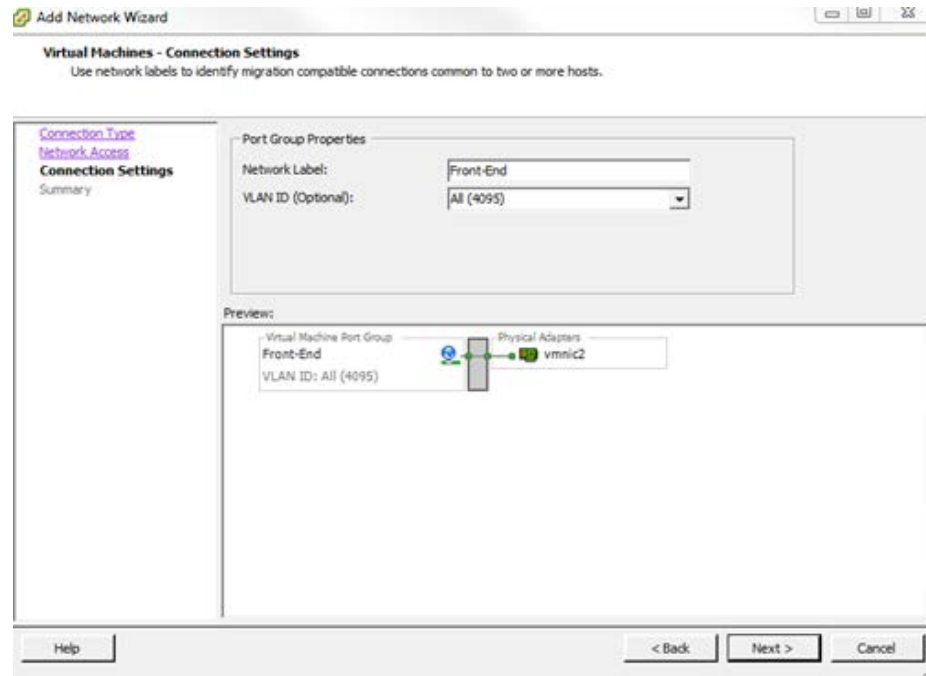
8. Click **Next.**

Note: This Physical Network adapter must be configured from the physical switch as a trunked interface or tagged on the VLANs that are planned to be used by HCP.



9. Label the **Network** and in the **VLAN ID** dropdown menu, select **All (4095)**.

This will allow HCP to configure virtual interfaces to talk to the switch on different networks (VLANs).

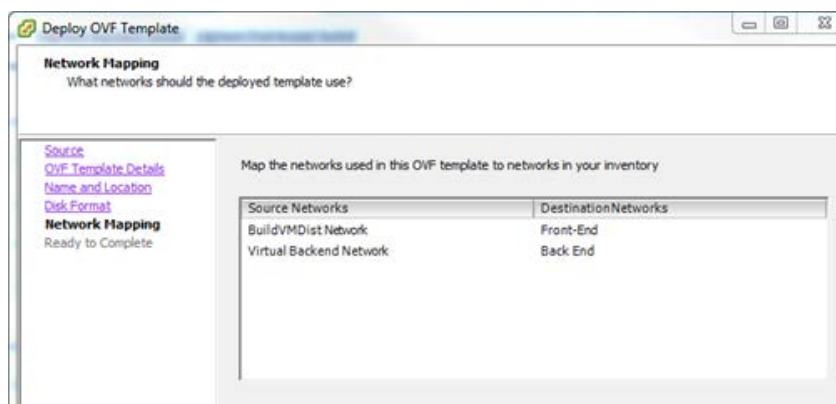


OVF deployment information

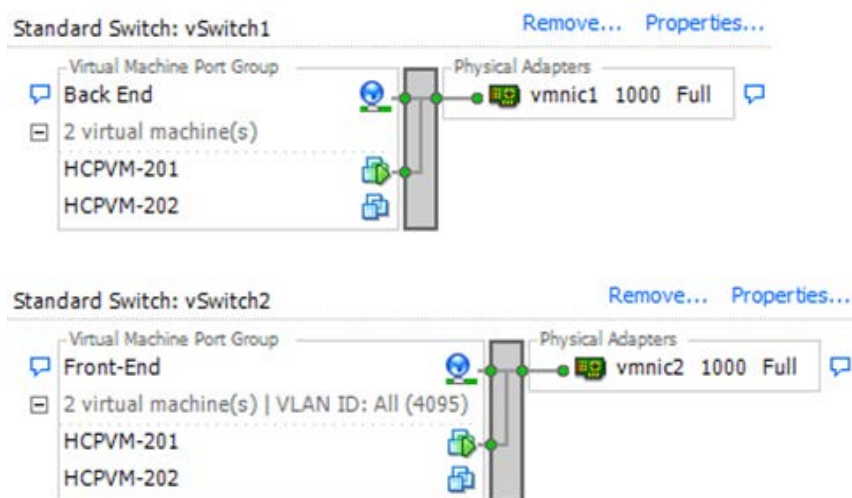
To deploy the OVF with the Networks you set up in the previous chapter:

1. After you have configured the Front and Back-end networks, deploy the HCP-VM OVF (For help see, ["Deploying the HCP-VM OVF VDMK"](#)) and follow the steps until you get to the **Network Mapping** menu.
2. Set the Destination Networks for the **BuildVMDist Network** to be the Front-End network you created in the previous chapter.
3. Set the Destination Networks for the **Virtual Back-end Network** to be the Back-end network you created in the previous chapter.

4. Finish the remainder of the deployment.



Your Virtual Port Groups will now look like this:

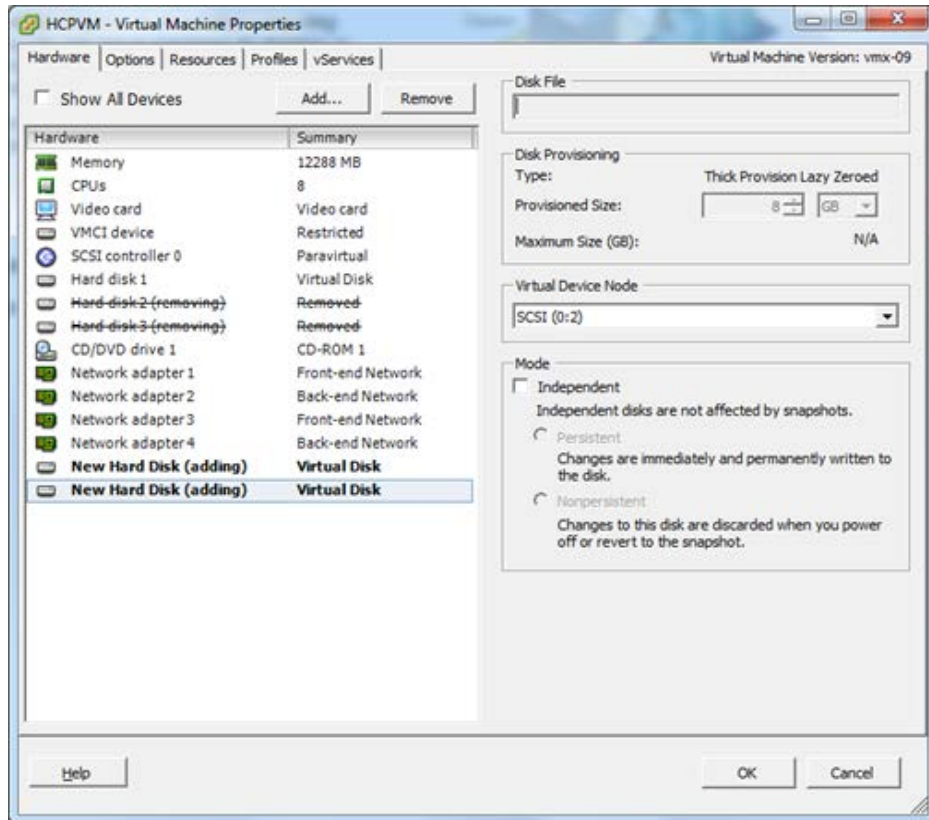


Changing the VMDK target size

It might be necessary to change the size of the VMDKs included with the HCP-VM OVF. To do this, allocate the appropriate storage to provision the datastore. If the HCP-VM is used for evaluation and **not** placed into production, the VMDKs can be made smaller than 500GB and/or be configured for thin provisioning.

To change the size of the VMDKs, perform the following steps on each HCP-VM node in the system:

1. Deploy the OVF as described in chapter, ["Deploying the HCP-VM OVF VDMK"](#).
2. Right-click on the HCP-VM node in the vSphere Client and select **Edit Settings**.
3. In the **Virtual Machine Properties** window, select **Hard disk 2** and click **Remove** (Do not delete the 32 GB OS vmdk).
4. Select **Hard disk 3** and click **Remove** (Do not delete the 32GB OS vmdk).
5. Click **Add**.
6. In the **Add Hardware** window, select **Hard Disk**.
7. Select the appropriate properties and size of the HCP data LUN.
8. Repeat the last three steps to create the second LUN.
9. In the Virtual Machine Properties window, click **OK**.



10. Right click on an inactive node and in the submenu hover over **Power** and click **Power On**.
11. Power on the rest of the HCP-VM nodes.
12. See, [Chapter 3: "Configuring the HCP-VM environment"](#) to change the network configuration on each node.
13. See, [Chapter 4: "Creating the HCP-VM system"](#) to install the HCP software.



DRS settings

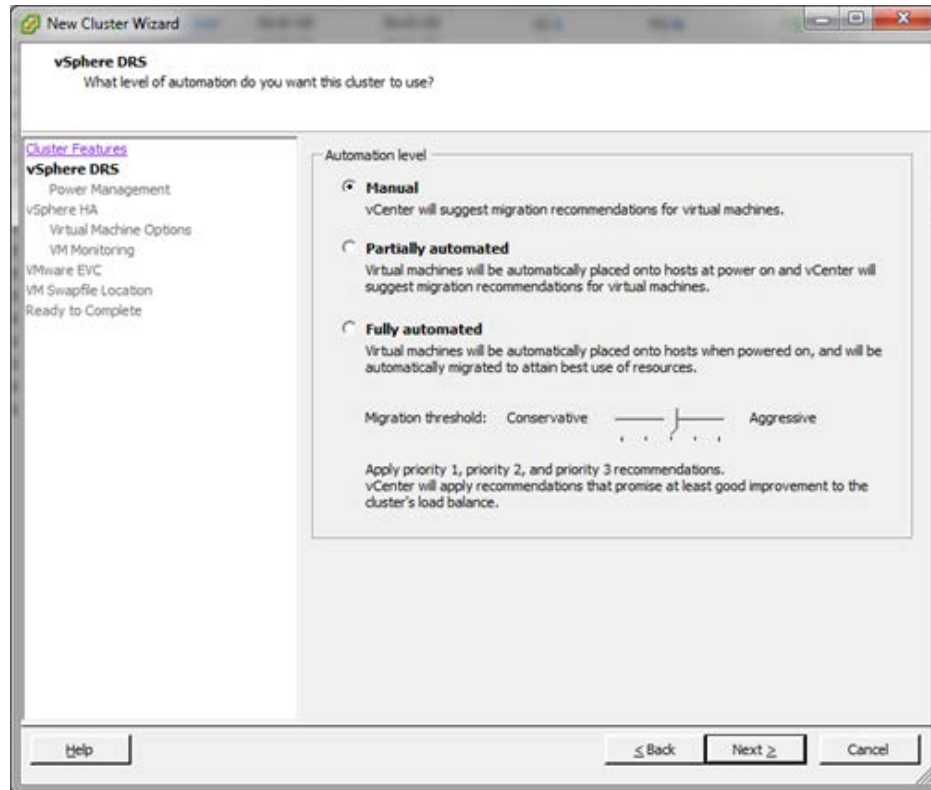
To modify the DRS setting:

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 **Next**.



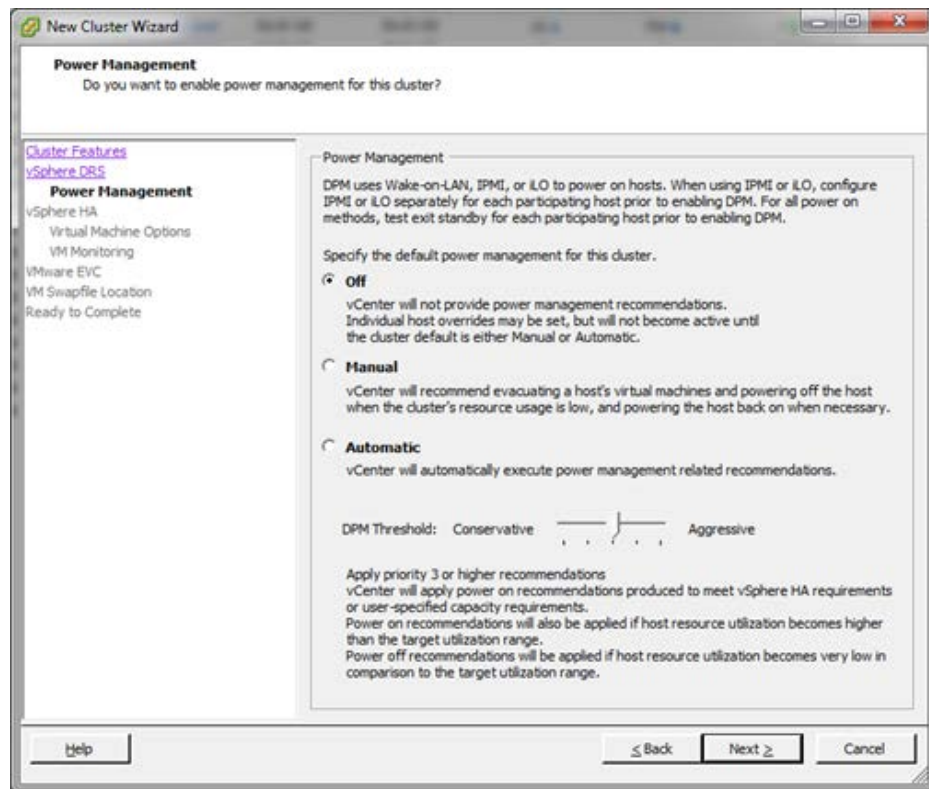
Important: Only turn on this feature if you feel your environment will benefit from it and you fully understand its functionality.

6. Select **Manual** for the DRS automation level in order to specify where VM guests should reside.
7. Click **Next**.



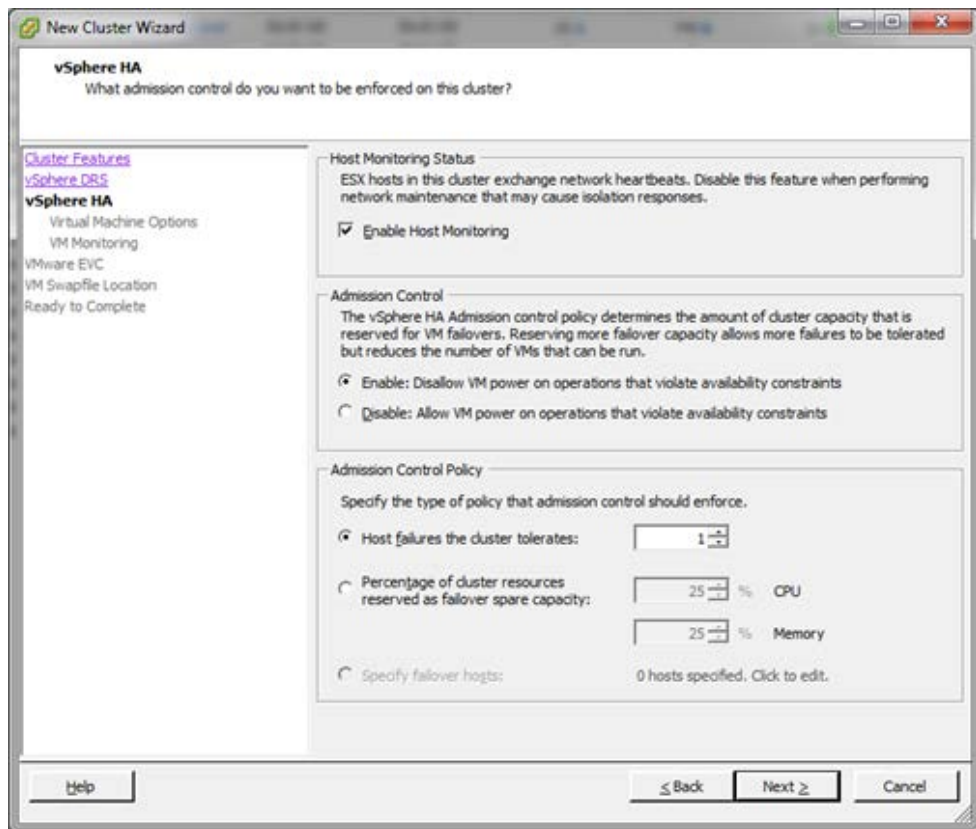
8. Select **Off** for the Power Management.

9. Click **Next**.

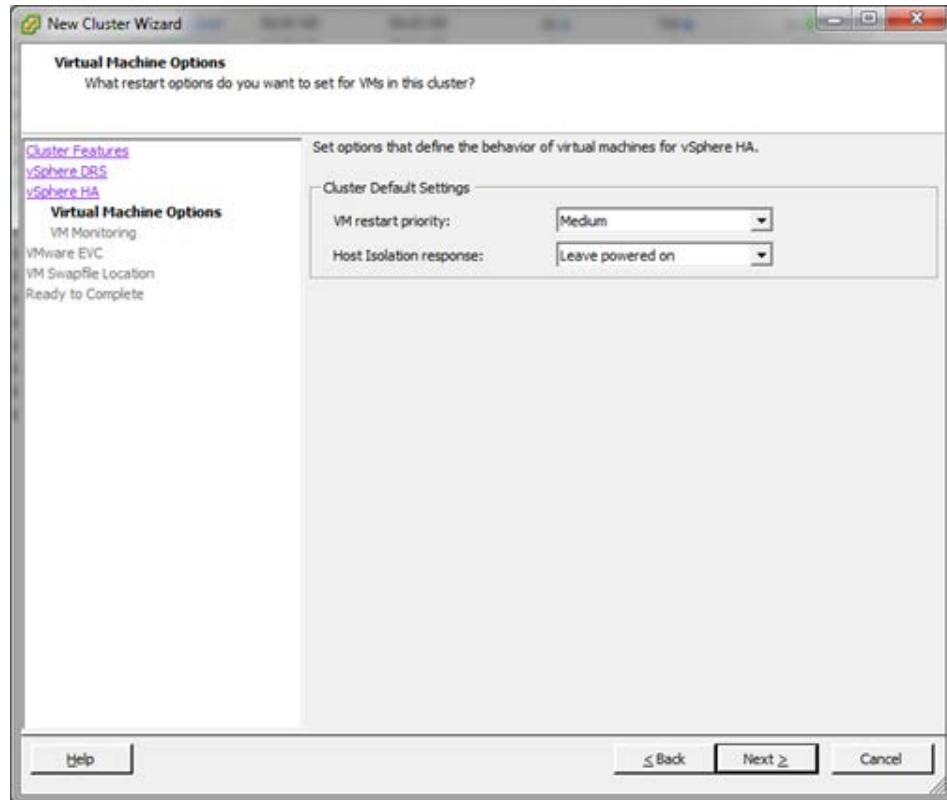


10. Select **Enable Host Monitoring** and keep the default settings.

11. Click **Next**.

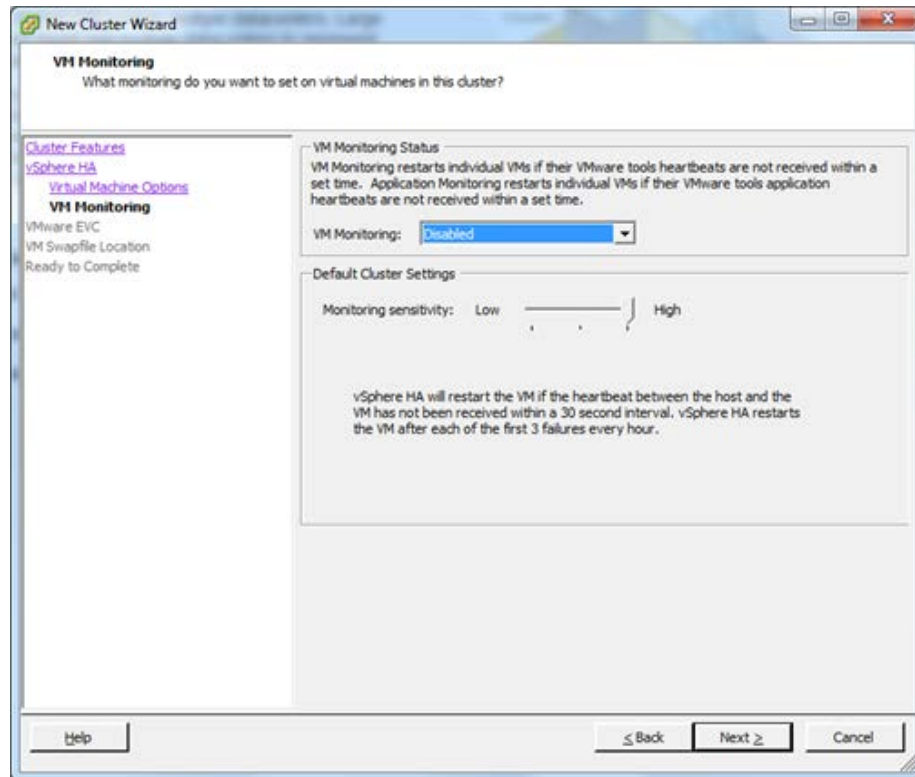


12. Leave the default settings and click **Next**.



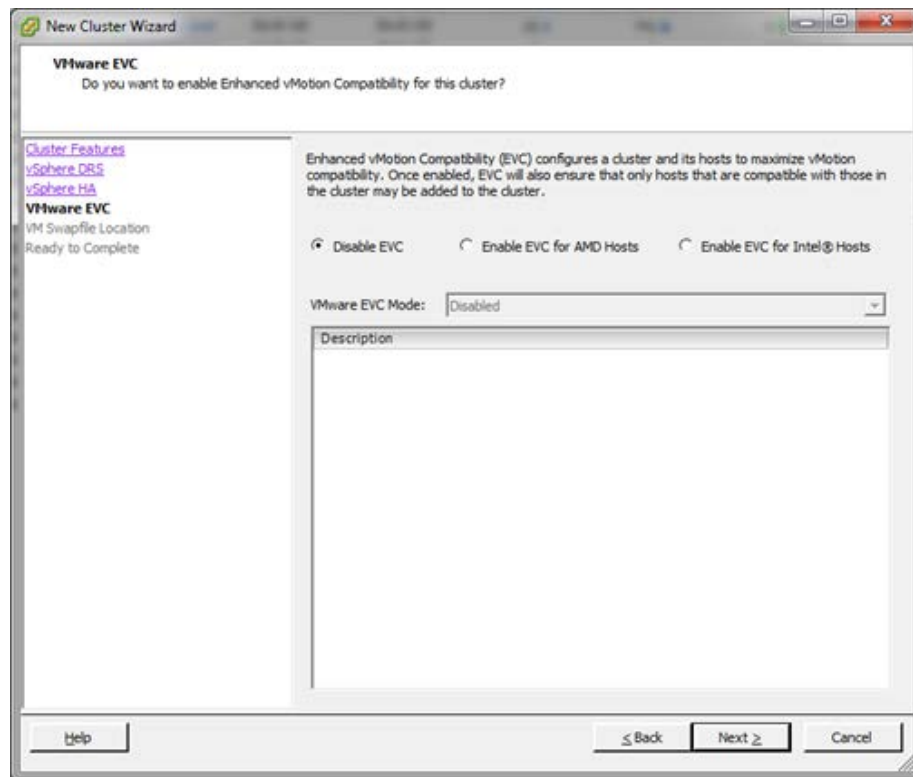
13. Set the **VM Monitoring** to **Disabled**.

14. Click **Next**.



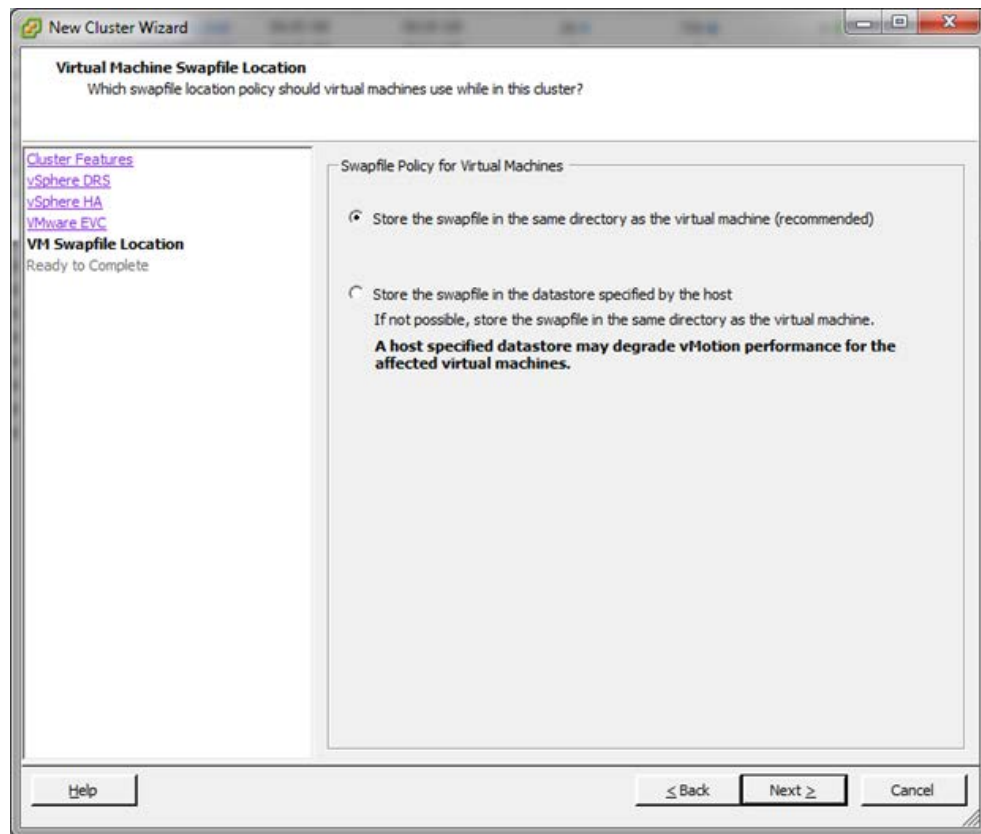
15. Select **Disable EVC**.

16. Click **Next**.



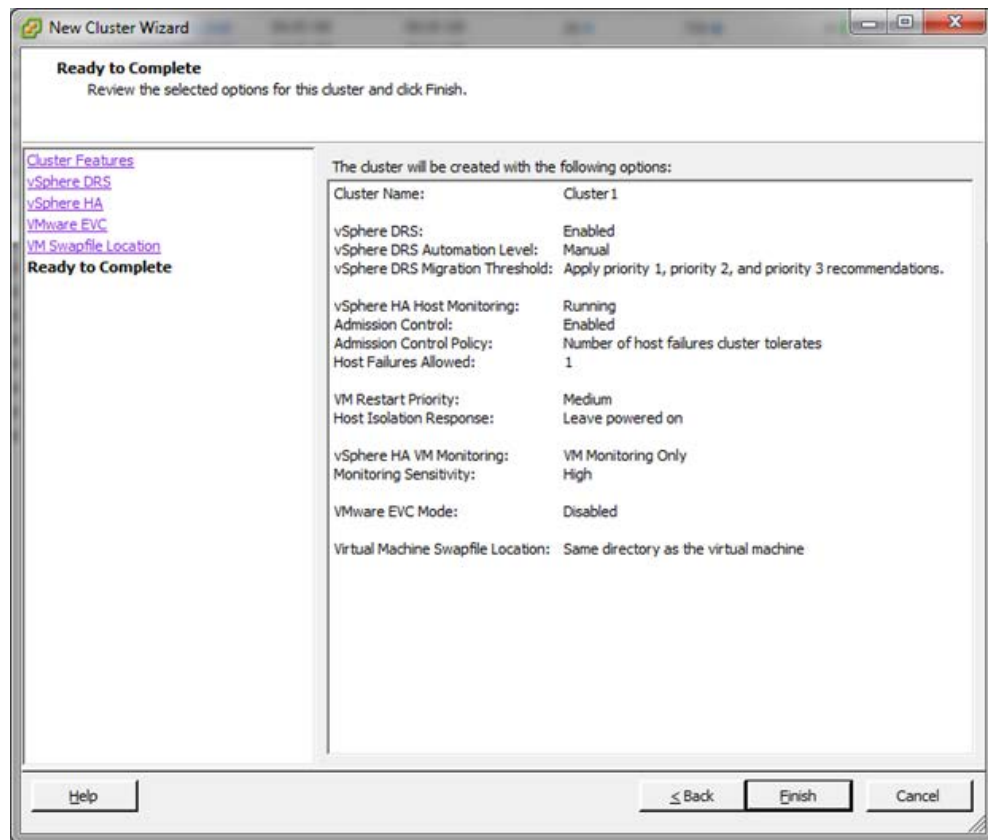
17. Select where you prefer to store your Swapfile location.

18. Click **Next**.

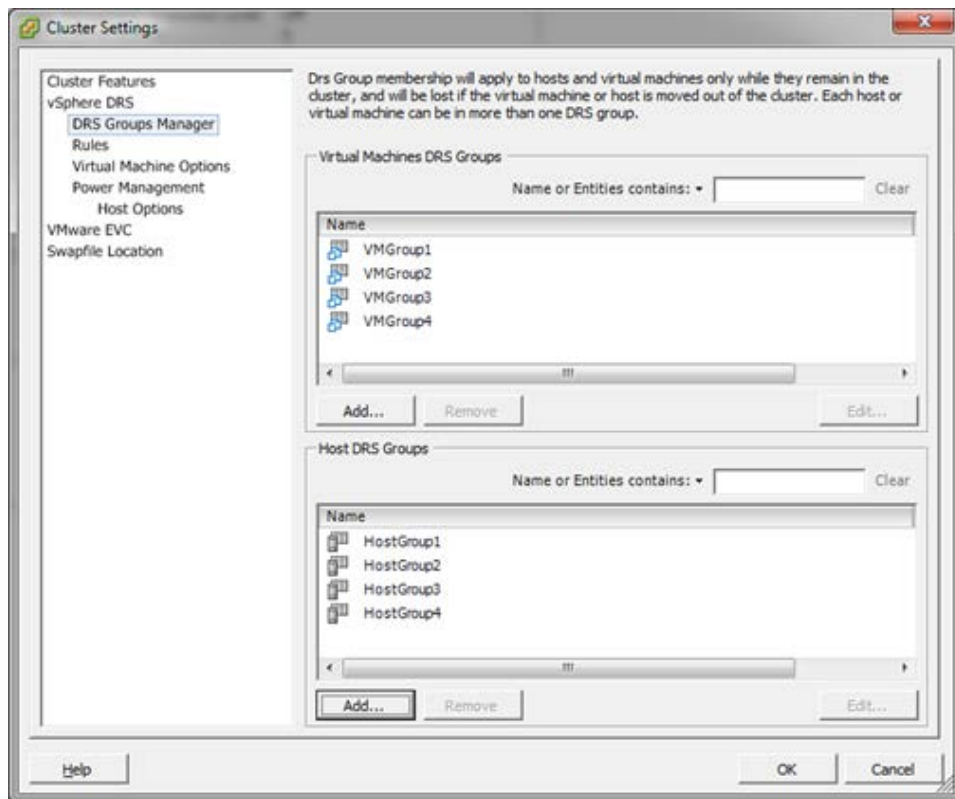


19. Review your settings.

20. Click **Finish**.



21. In the left side navigation bar, select a Cluster and right click it. Then click on **Edit Settings**.
22. On the left side navigation bar of the **Settings** window, click on **DRS Groups Manager**.
23. In the **DRS Groups Manager** , create a group and add the Virtual Machines that you would like to keep on a specific server.
24. Create one Virtual Machine DRS Group for each Host.
25. Click **Add** in the Host DRS Groups section and place one host from the cluster in each group.
26. Click **Next**.



27. On the left side navigation bar of the **Settings** window, click on **Rules**.
28. 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**.
29. Select **Should run on hosts in group**.
30. Click **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.

Rule

Rule | DRS Groups Manager |

Give the new rule a name and choose its type from the menu below.
Then, select the entities to which this rule will apply.

Name
Group1

Type
Virtual Machines to Hosts

DRS Groups

Cluster Vm Group:
VMGroup1

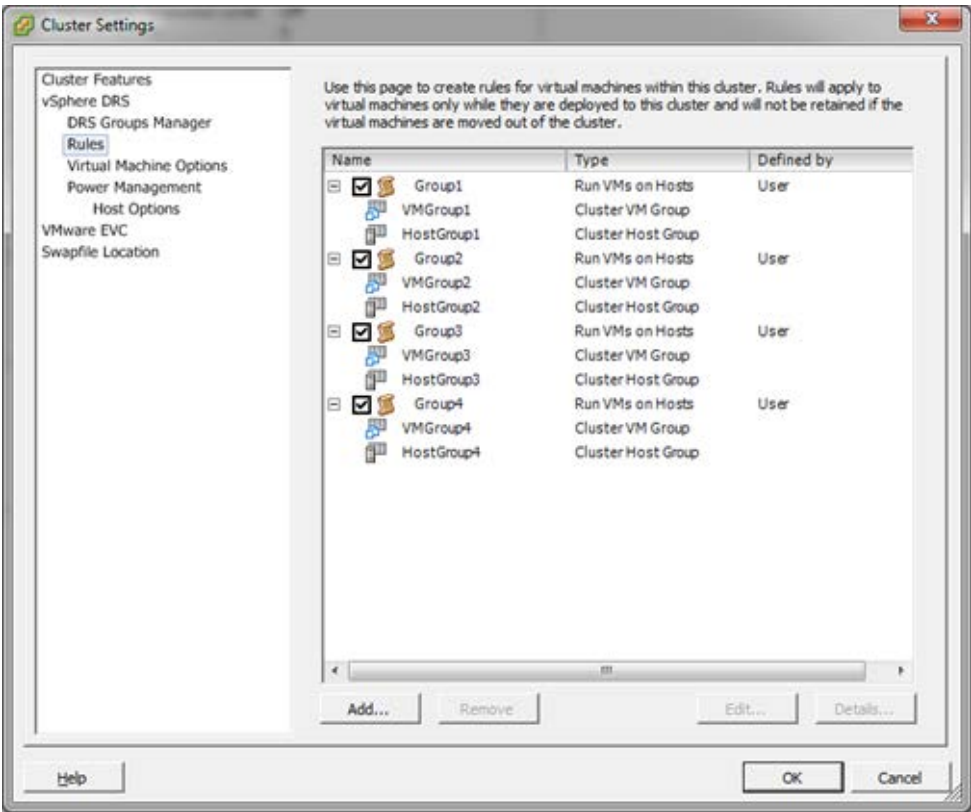
Should run on hosts in group

Cluster Host Group:
HostGroup1

Virtual machines that are members of the Cluster DRS VM Group
VMGroup1 Should run on hosts in group HostGroup1.

OK Cancel

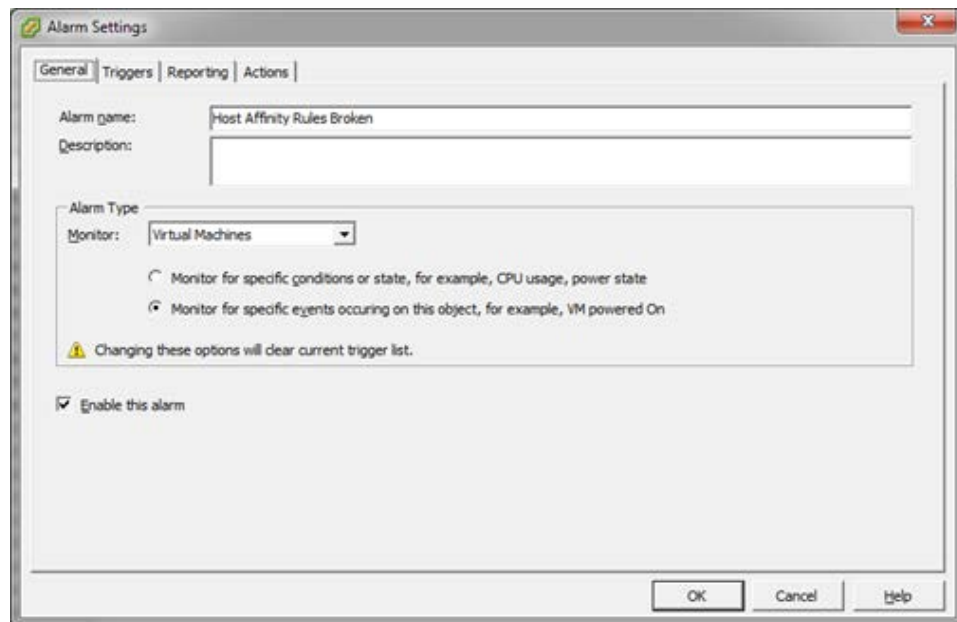
Now that you have created all the Rules your cluster settings should look something like this:



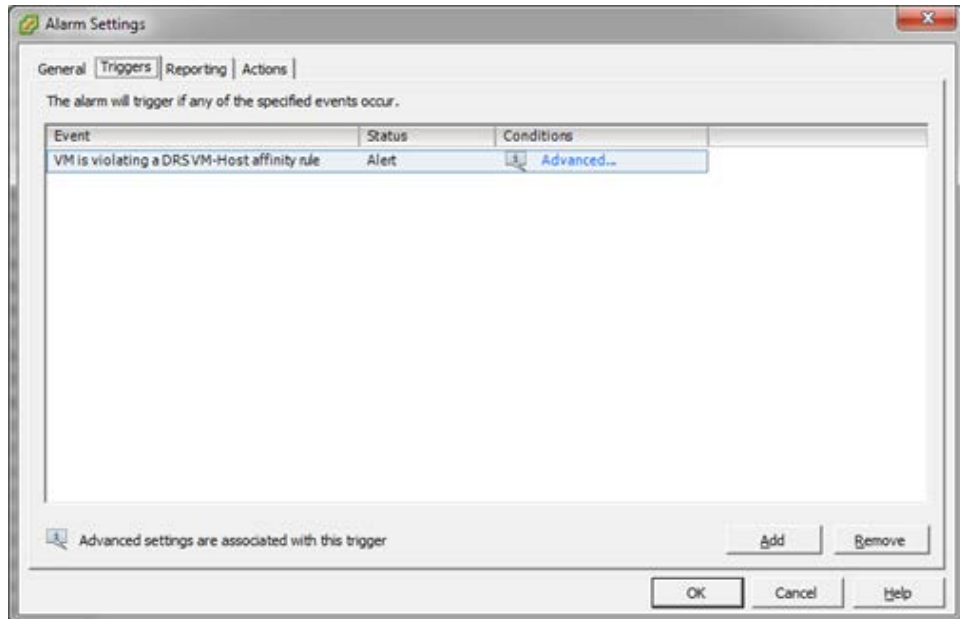
Setting an alarm

To set an alarm:


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

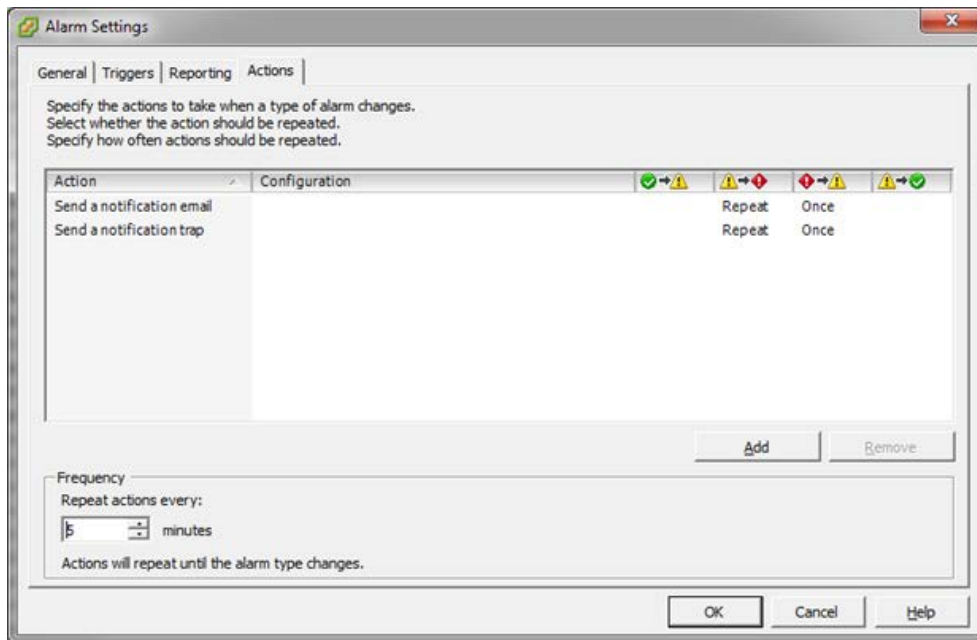


6. Under the **Trigger Conditions** select an **Argument of VM** name.
7. Set the Value equal to each VM you want to monitor.
8. Add one argument for each VM.



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

The image below shows either an email or SNMP trap, and what to do when it goes from Green to Warning or Warning to Alert. Since you selected Alert earlier you can set how often the alert repeats in Frequency.





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 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. Logical volumes can be local or external.

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.

HCP-DM treats HCAP 2.x archives and local file systems as namespaces.

network

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

network file system

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

network interface controller (NIC)

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

NFS

See ["network file system"](#).

NIC

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

node

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

O**object**

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

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

open virtualization format (OVF)

Standard file style for packaging and distributing virtual software.

OVF

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

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

The process of keeping selected HCP tenants and namespaces and selected default-namespace directories in two HCP systems in sync with each other. This entails copying object creations, deletions, and metadata changes from each system to the other or from one system to the other. HCP also replicates tenant and namespace configuration, tenant-level user and group accounts, retention classes, content classes, all compliance log messages, and all HCP tenant log messages.

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

•

.zip 79

B

Back-end network 17, 165

Back-end Network 71

base configuration, Hi-Track Monitor 160

C

Cluster 31

Compute 14

configuring Hi-Track Monitor 159

Console pages

SNMP 158

D

data center 30

Datastores 49

diagnostic 131

DRS 173

E

email, Hi-Track Monitor 160

enabling SNMP 158

ESXi 39

Evaluation 19

F

Failover 19, 147

Fibre Channel Connectivity 46

Front-end network 17, 167

Front-end Network 65

H

HCP 13

HCP-VM 13, 19, 25

HCP-VM network 108

HCP-VM nodes 155

HCP system 116

HCP systems

enabling SNMP 158

identifying in Hi-Track Monitor 162

HDDS 19

Heartbeat 62

Hi-Track Monitor

about 157

base configuration 160

configuring 159

email 160

identifying HCP systems 162

logging in 159

transport agents 161

HiTrack 18

HNAS 60, 64

I

IPMI 130

L

logging into Hi-Track Monitor 159

logical volumes 148

LUN 15

M

Metadata 13

multi-cast 17

N

Namespaces 13

Network Time Protocol 27

NFS 60

NFS datastore 61

NTP 27

OVF

O

OVF 79, 91, 169

P

pNIC 17-18

PVSCSI 15

R

RAID6 14

RDM 15, 91

Repository 13

S

SAR 131

SNMP 130

SNMP page 158

SNMP, enabling 158

switching 65, 165, 167

Switching 71

System management console 130

T

transport agents, Hi-Track Monitor 161

V

VMDK 15, 79, 171

VMFS 15

vmNICs 17

Vmware 13, 21

vNIC 17

vSphere 13

vSphere HA cluster 19, 30

Hitachi Vantara

Corporate Headquarters
2845 Lafayette 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

