vmware PRESS

FREE VMware Press 2013 eSampler

Includes excerpts from upcoming VMware Press titles

vmwarepress.com

ALWAYS LEARNING

PEARSON

vmware PRESS

CHAPTERS IN THIS eSAMPLER ARE DRAFT MANUSCRIPTS

NOT AVAILABLE FOR RESALE

This manuscript has been provided by Pearson Education at this early stage to create awareness for this upcoming book. It has not been copy edited or proofread yet; we trust that you will judge this book on technical merit, not on grammatical and punctuation errors that will be fixed at a later stage.

vmwarepress.com

vmware® PRESS

FREE

VMware Press 2013 eSampler

TABLE OF CONTENTS

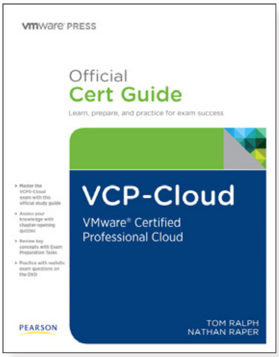
vmware® PRESS
Official
Cert Guide
Learn, prepare, and practice for exam success
VCP-Cloud
VMware® Certified Professional Cloud
PEARSON
TOM RALPH
NATHAN RAPER

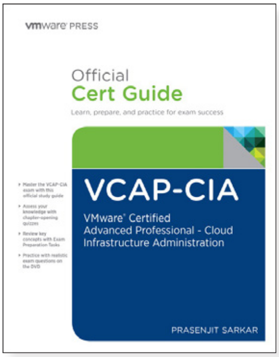
vmware® PRESS
Official
Cert Guide
Learn, prepare, and practice for exam success
VCAP-CIA
VMware® Certified Advanced Professional - Cloud Infrastructure Administration
PRASENJIT SARKAR

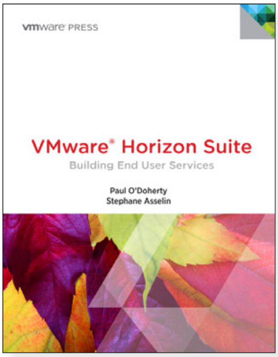

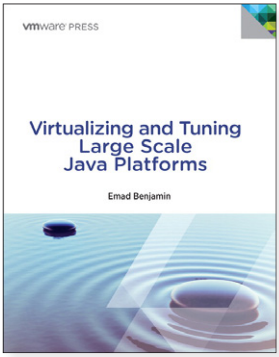

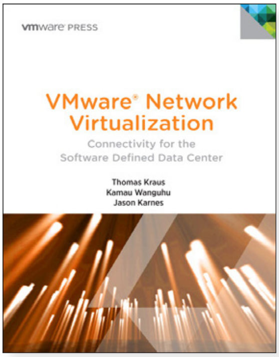
vmware® PRESS
VMware® Network Virtualization
Connectivity for the Software Defined Data Center
Thomas Kraus
Kamau Wanguhu
Jason Karnes

vmwarepress.com

vmware® PRESS

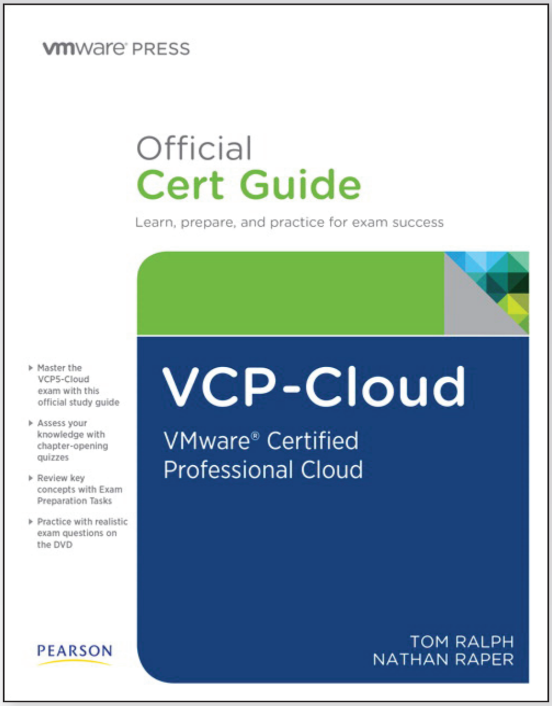

vmware® PRESS

Official
Cert Guide

Learn, prepare, and practice for exam success

- Master the VCP5-Cloud exam with this official study guide
- Assess your knowledge with chapter-opening quizzes
- Review key concepts with Exam Preparation Tasks
- Practice with realistic exam questions on the DVD

VCP-Cloud

VMware® Certified Professional Cloud

PEARSON

TOM RALPH
NATHAN RAPER

CHAPTER 8

Allocate and Manage vCloud Resources

AVAILABLE – NOVEMBER 2013

vmwarepress.com

This chapter covers the following subjects:

- **What is a vCloud Resource?**—This section explains vCloud resources, what they are, and how to define them in vCloud Director.
- **Create and Administer Provider vDCs**—This section explains Provider vDCs, their creation and administration.
- **Create and Administer Organization vDCs**—This section covers the creation of Organization vDCs and their administration.
- **Catalog Management**—In this section you will learn about catalog management, including the population of the vApp Templates and media items.

This chapter covers a portion of the Exam Title objective 1.1 and Exam Title objective 1.2.

CHAPTER 8

Allocate and Manage vCloud Resources

Good news, with the introduction of vCloud Director you now have the ability to control the exact amount of resources your end-users can consume, while still allowing for multi-tenancy. This new ability to manage the allocation and consumption of user resources will enable you to accurately assign and monitor resources in the cloud. However, this ability does not come without a price, and the management overhead introduced by vCloud Director must be accounted for. In this chapter we will cover not only how to configure and use these new *allocation models*, we will also cover their proper usage. Finally, we will explore the relationship between allocation models and vCenter Server configuration changes.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter or simply jump to the "Exam Preparation Tasks" section for review. If you are in doubt, read the entire chapter. Table 8-1 outlines the major headings in this chapter and the corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Troubleshooting Scenarios."

Table 8-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping

Foundations Topics Section	Questions Covered in This Section
What is a vCloud Resource	9
Create and Administer Provider vDCs	1-2
Create and Administer Organization vDCs	3, 4, 6, 7
Catalog Management	5, 8

1. It is possible for more than 1 resource pool or cluster to be used by a single provider vDC.

   a. True

   b. False

2. When adding storage to vCloud Director, data stores are directly added to the provider vDC.

   a. True

   b. False

3. Where is the limit on CPU resources placed in an Allocation Pool organization vDC?

   a. On individual Virtual Machines

   b. On the Provider vDC

   c. On the Organization vDC

   d. On the vCenter Server Cluster

4. When defining a Pay as You Go organization vDC, which three options are defined? (Choose three.)

   a. CPU Limit per Virtual Machine

   b. Memory Limit per Virtual Machine

   c. vCPU Count per Virtual Machine

   d. Maximum number of VMs that can be deployed

   e. Maximum RAM allocated to the Organization vDC

5. When sharing an item from the catalog, what items are allowed to be shared?

   a. ISO images

   b. vApp Templates

   c. Floppy Images

   d. vApps

6. A reservation organization vDC, allows for the configuration of which two options? (Choose two)

   a. CPU Limit

   b. Memory per VM limit

c. CPU Reservation

d. Memory limit

7. How many organizations can share an organization vDC?

a.

b.

c.

d. nlimited

8. vCloud Director shares a published catalog to how many organizations?

a. 0

b. 1

c. User definable

d. All Organizations

9. vCloud Director provisions resources from which hyper-visor?

a. VMware vSphere

b. Microsoft Windows Hyper-V

c. Citrix XEN Server

Foundation Topics

What is a vCloud Resource?

VMware defines cloud resources in two sections, compute and network resources. vCenter Server clusters and vSphere hosts provide compute resources, while vCloud Networking and Security (in conjunction with vCenter Server) provides the network resources. We covered network resource configuration and allocation in Chapter 6 – Configure and Administer vCloud Networking.

vCloud Director allows for compute resources to be provisioned using a provider virtual datacenter, assigned to organizations through organization virtual datacenters, and finally consumed by users through containers called vApps.

In this chapter we discuss compute resources and how vCloud Director presents those resources to you, the Cloud Administrator. Then, we will explain how to define the consumption model for those resources, and ultimately how the end user consumes the resources provided.

Create and Administer Provider vDCs

vCloud Director's first abstraction of resources is the Provider Virtual Datacenter, commonly referred to as a Provider vDC. A Provider vDC takes the compute and memory resources from a vCenter Server resource pool and combines them with one or more available datastores to create a group of resources available within the cloud. These resources are then provisioned to one or more Org vDCs, or Organization Virtual Datacenters. We will cover Org vDCs in the next section.

##KeyTopic

A Provider vDC is a combination of one or more resource pools or clusters defined by vCenter Server. If multiple resource pools are provisioned to a single Provider vDC, that Provider vDC is considered an elastic Provider vDC.

What is an Elastic Provider vDC?

When VMware released vCloud Director 1.5, they incorporated the ability to include more than 1 resource pool per Provider vDC for a Pay as You Go allocation model. With the release of 5.1, VMware expanded this capability to include the Allocation Pool allocation model. When configured with multiple resource pools, these Elastic Provider vDCs provide the flexibility for cloud resources to be scaled as needed.

NOTE: A cluster in vCenter Server is considered the root resource pool. It is a VMware recommended practice to place Provider vDCs at the root level (i.e. the cluster) during initial configuration.

Several things must be in place and working for an elastic Provider vDC to function properly. First, the resource pools must exist in the same vCenter Server and the same vCenter Server Datacenter. Second, the network pool that is backing the workloads must be capable of being extended to all resource pools used by the Provider vDC. This extension is necessary to avoid issues with network communication between VMs. Finally, the storage between the resource pools should be shared. If the storage is not shared, deployment times will be greatly extended. This is due to the fact that vCloud Director must export the VM using the 'Export OVF' process, then re-import the VM to the other resource pool's storage.

NOTE: Elastic Provider vDCs are only available when using the Allocation Pool model or the Pay as You Go model. A Provider vDC configured using the Reservation Pool model cannot be used as an elastic Provider vDC.

##Key Topic

ACTIVITY 8-1 CREATE A PROVIDER VDC

1. Login to vCloud Director, click on the Manage & Monitor tab

2. Click on the Provider vDCs option in the left pane, as shown in Figure 8-1

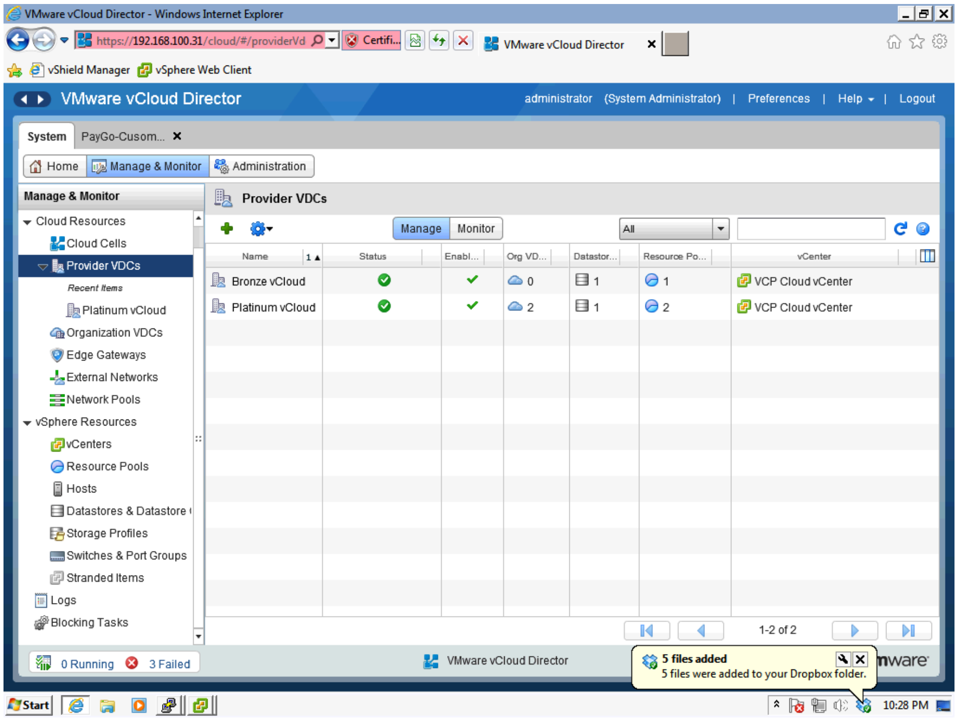

Name	Status	Enabl...	Org VD...	Datastor...	Resource Po...	vCenter
Bronze vCloud	✓	✓	0	1	1	VCP Cloud vCenter
Platinum vCloud	✓	✓	2	1	2	VCP Cloud vCenter

5 files added
5 files were added to your Dropbox folder.

Figure 8-1 Provider vDC Selection

3. Click the green + or the blue gear symbol and select **New Provider vDC,** shown in Figure 8-2

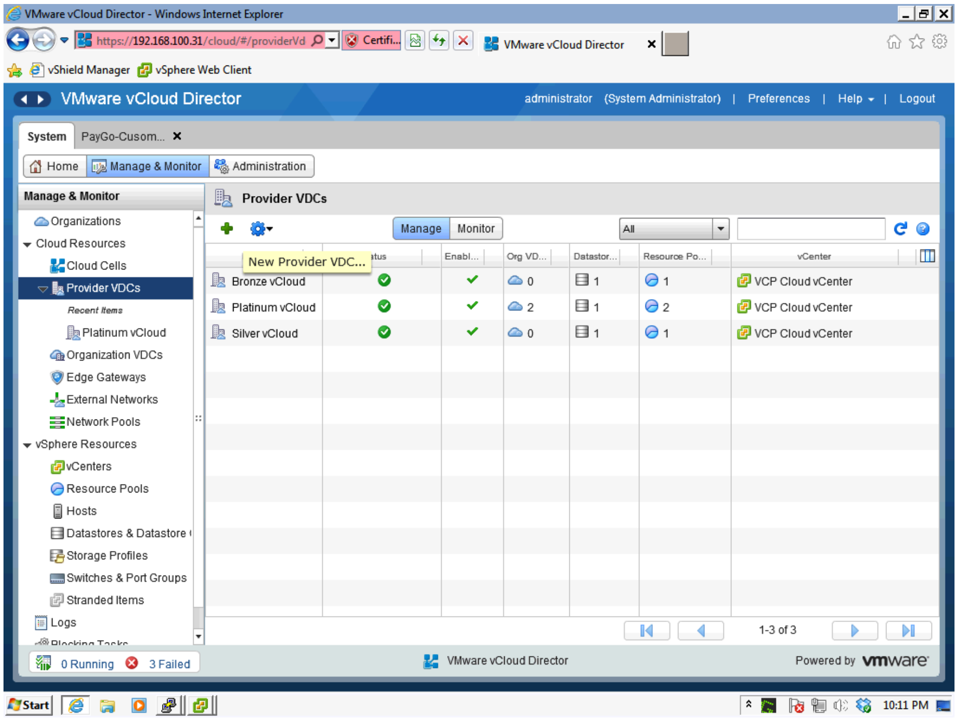

	...tus	Enabl...	Org VD...	Datastor...	Resource Po...	vCenter
Bronze vCloud	✓	✓	0	1	1	VCP Cloud vCenter
Platinum vCloud	✓	✓	2	1	2	VCP Cloud vCenter
Silver vCloud	✓	✓	0	1	1	VCP Cloud vCenter

Figure 8-2 Add new Provider vDC

4. Type in a name and description

NOTE: A good use of the description field is to indicate what the resource pool provides as shown in Figure 8-3

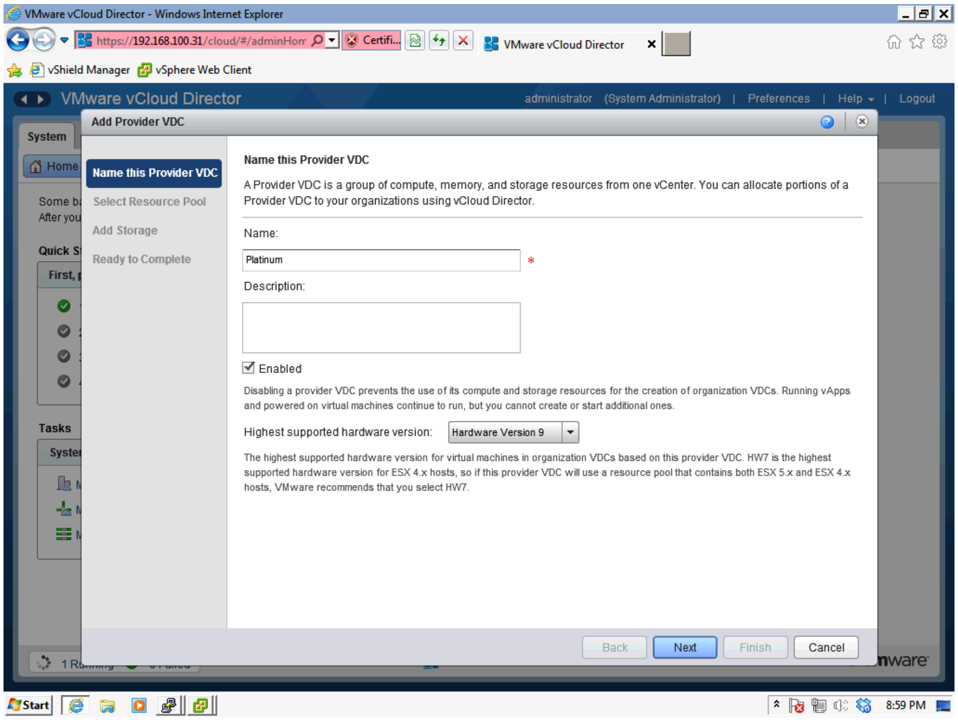

Name this Provider VDC

A Provider VDC is a group of compute, memory, and storage resources from one vCenter. You can allocate portions of a Provider VDC to your organizations using vCloud Director.

Name: Platinum *

Description:

☑ Enabled

Disabling a provider VDC prevents the use of its compute and storage resources for the creation of organization VDCs. Running vApps and powered on virtual machines continue to run, but you cannot create or start additional ones.

Highest supported hardware version: Hardware Version 9

The highest supported hardware version for virtual machines in organization VDCs based on this provider VDC. HW7 is the highest supported hardware version for ESX 4.x hosts, so if this provider VDC will use a resource pool that contains both ESX 5.x and ESX 4.x hosts, VMware recommends that you select HW7.

Figure 8-3 Provider vDC naming

5. Select the highest VM Hardware version that your vSphere installation supports, as shown in Figure 8-4

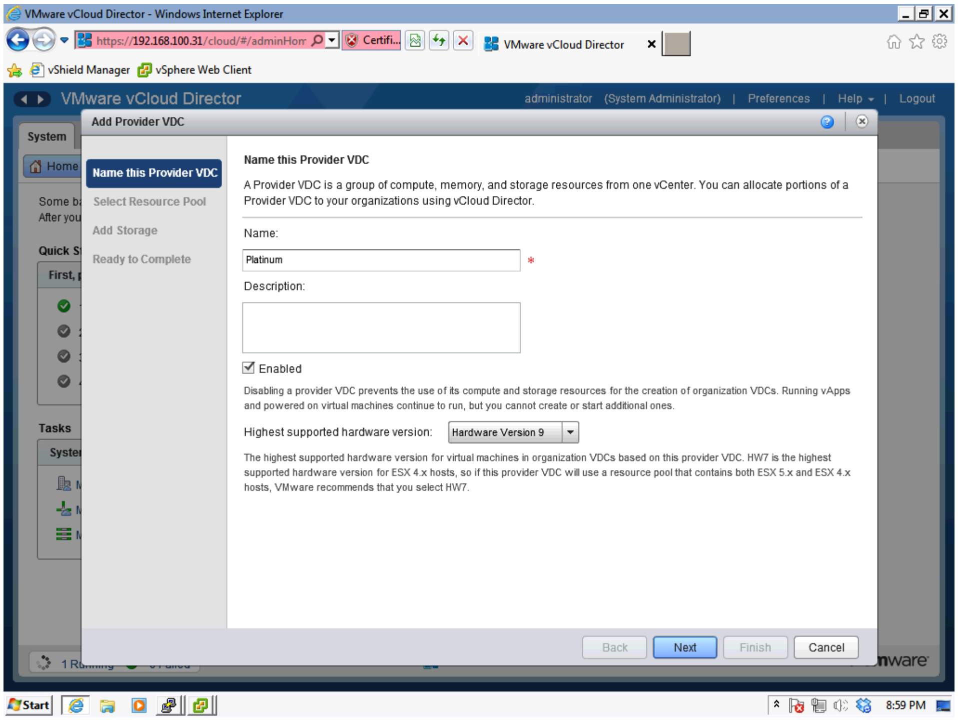

Name this Provider VDC

A Provider VDC is a group of compute, memory, and storage resources from one vCenter. You can allocate portions of a Provider VDC to your organizations using vCloud Director.

Name: Platinum *

Description:

☑ Enabled

Disabling a provider VDC prevents the use of its compute and storage resources for the creation of organization VDCs. Running vApps and powered on virtual machines continue to run, but you cannot create or start additional ones.

Highest supported hardware version: Hardware Version 9

The highest supported hardware version for virtual machines in organization VDCs based on this provider VDC. HW7 is the highest supported hardware version for ESX 4.x hosts, so if this provider VDC will use a resource pool that contains both ESX 5.x and ESX 4.x hosts, VMware recommends that you select HW7.

Figure 8-4 VM Hardware selection

NOTE: Here is a quick reference chart of VM Hardware versions and the corresponding vSphere Version

- vSphere 4 – Hardware Version 7
- vSphere 5 – Hardware Version 8
- vSphere 5.1 – Hardware Version 9

***end note

6. Next, as shown in Figure 8-5, select the vCenter Server and the resource pool that this Provider vDC will consume.

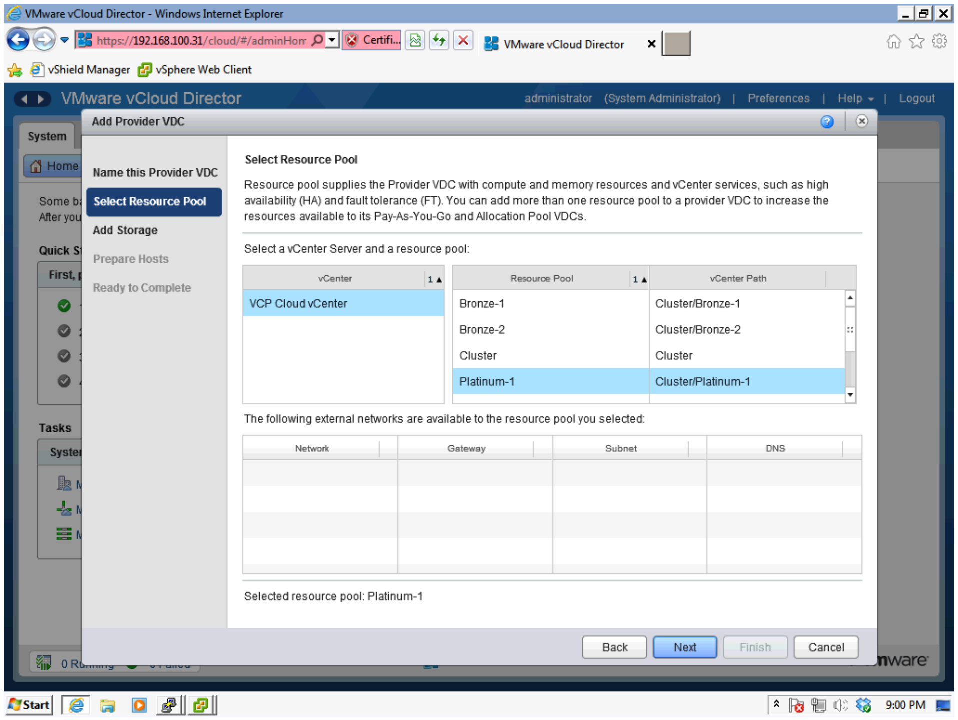

Select Resource Pool

Resource pool supplies the Provider VDC with compute and memory resources and vCenter services, such as high availability (HA) and fault tolerance (FT). You can add more than one resource pool to a provider VDC to increase the resources available to its Pay-As-You-Go and Allocation Pool VDCs.

Select a vCenter Server and a resource pool:

vCenter
VCP Cloud vCenter

Resource Pool	vCenter Path
Bronze-1	Cluster/Bronze-1
Bronze-2	Cluster/Bronze-2
Cluster	Cluster
Platinum-1	Cluster/Platinum-1

The following external networks are available to the resource pool you selected:

Network	Gateway	Subnet	DNS

Selected resource pool: Platinum-1

Figure 8-5 vCenter and Resource Pool selection

NOTE: vCloud Director assumes that there will be no other users of the resource pool allocated to the Provider vDC. If there are other workloads utilizing this resource pool, they may be starved of resources as vCloud Director will consume and allocate all resources in the pool to the Provider vDC.

7. Select the appropriate storage profile for this Provider vDC. Figure 8-6 shows the ‘Any’ profile selected.

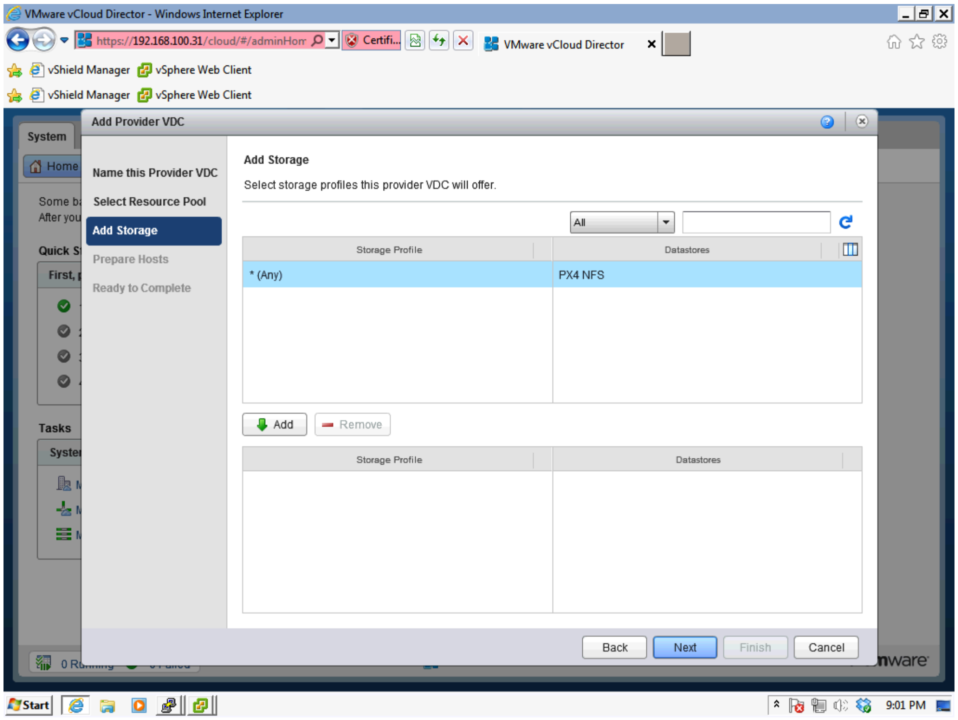

Storage Profile	Datastores
* (Any)	PX4 NFS

Figure 8-6 Storage Profile Selection

8. Click **Finish**.

Preparing a Provider vDC

After you have created your first Provider vDC, you must prepare the hosts that will provide the physical resources to that Provider vDC. Until the hosts are prepared by vCloud Director, they cannot be used to host a vCloud workload.

Key Topic

ACTIVITY 8-2 PREPARE A HOST

1. Click on the Manage & Monitor tab, then select the Provider vDCs option in the left pane
2. Find the Provider vDC with hosts that need prepared, then right click on it, and select **Open** , as shown in Figure 8-7

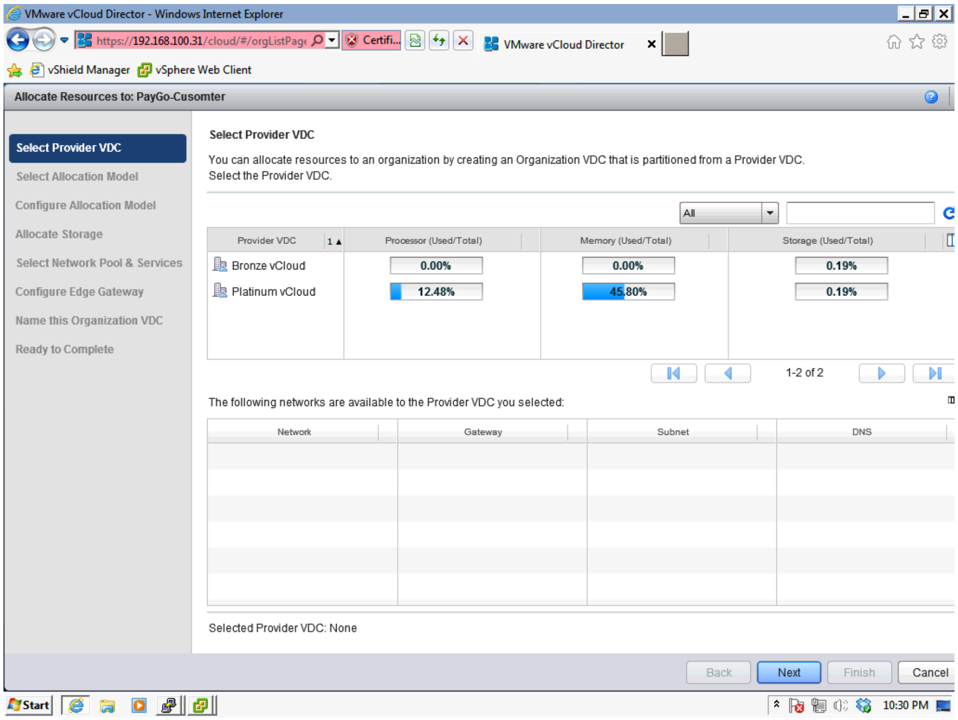
VMware vCloud Director - Windows Internet Explorer
https://192.168.100.31/cloud/#/orgListPage
Certifi...
VMware vCloud Director
vShield Manager
vSphere Web Client
Start
10:30 PM

Allocate Resources to: PayGo-Cusomter

Select Provider VDC

You can allocate resources to an organization by creating an Organization VDC that is partitioned from a Provider VDC.
Select the Provider VDC.

Provider VDC	Processor (Used/Total)	Memory (Used/Total)	Storage (Used/Total)
Bronze vCloud	0.00%	0.00%	0.19%
Platinum vCloud	12.48%	45.80%	0.19%

1-2 of 2

The following networks are available to the Provider VDC you selected:

Network	Gateway	Subnet	DNS

Selected Provider VDC: None

Figure 8-7 Provider vDC selection

3. As shown in Figure 8-8, Select the Hosts tab

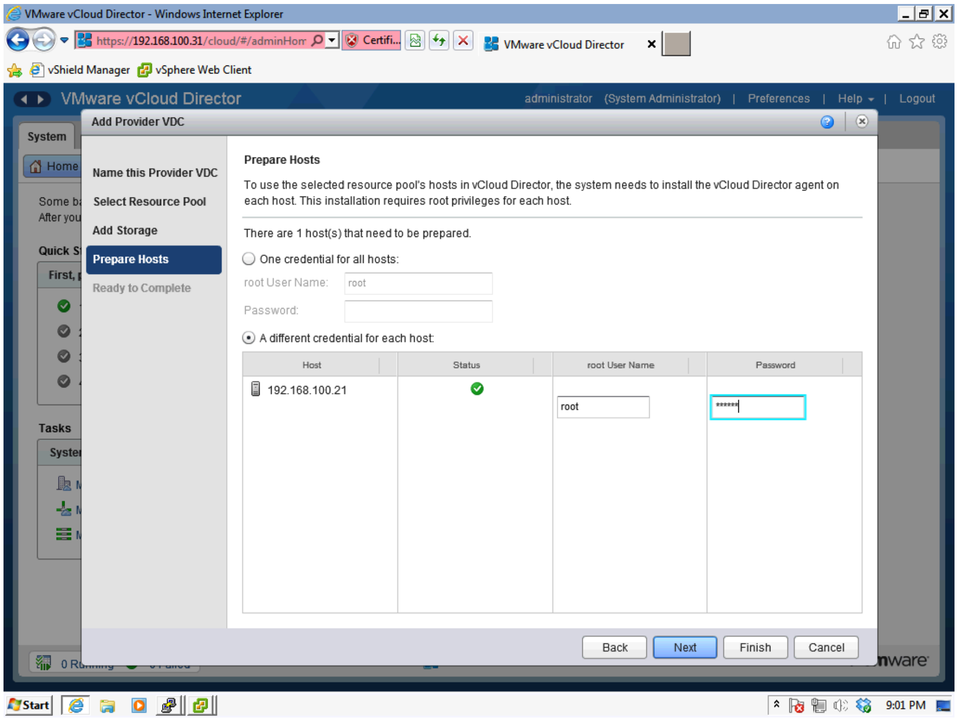

VMware vCloud Director - Windows Internet Explorer

https://192.168.100.31/cloud/#/adminHom

VMware vCloud Director

administrator (System Administrator) | Preferences | Help | Logout

Add Provider VDC

- Name this Provider VDC
- Select Resource Pool
- Add Storage
- **Prepare Hosts**
- Ready to Complete

Prepare Hosts

To use the selected resource pool's hosts in vCloud Director, the system needs to install the vCloud Director agent on each host. This installation requires root privileges for each host.

There are 1 host(s) that need to be prepared.

○ One credential for all hosts:

root User Name: root

Password:

● A different credential for each host:

Host	Status	root User Name	Password
192.168.100.21	✓	root	******

Back | Next | Finish | Cancel

Figure 8-8 Host Selection

4. Select the hosts to prepare, then right click and select **Prepare Host**, as shown in Figure 8-9

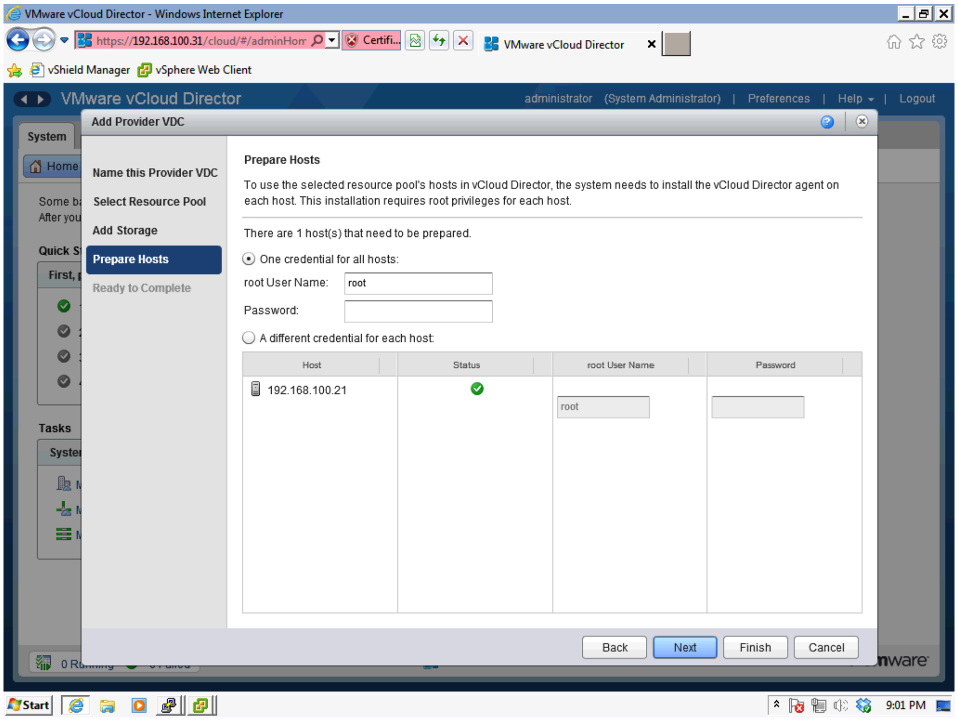

Prepare Hosts

To use the selected resource pool's hosts in vCloud Director, the system needs to install the vCloud Director agent on each host. This installation requires root privileges for each host.

There are 1 host(s) that need to be prepared.

● One credential for all hosts:

root User Name: root

Password:

○ A different credential for each host:

Host	Status	root User Name	Password
192.168.100.21	✓	root	

Figure 8-9 Prepare host selection

5. Input a username and password of a user that has administrator privileges on the vSphere host. Figure 8-10 shows an example using root, which is the most common user used.

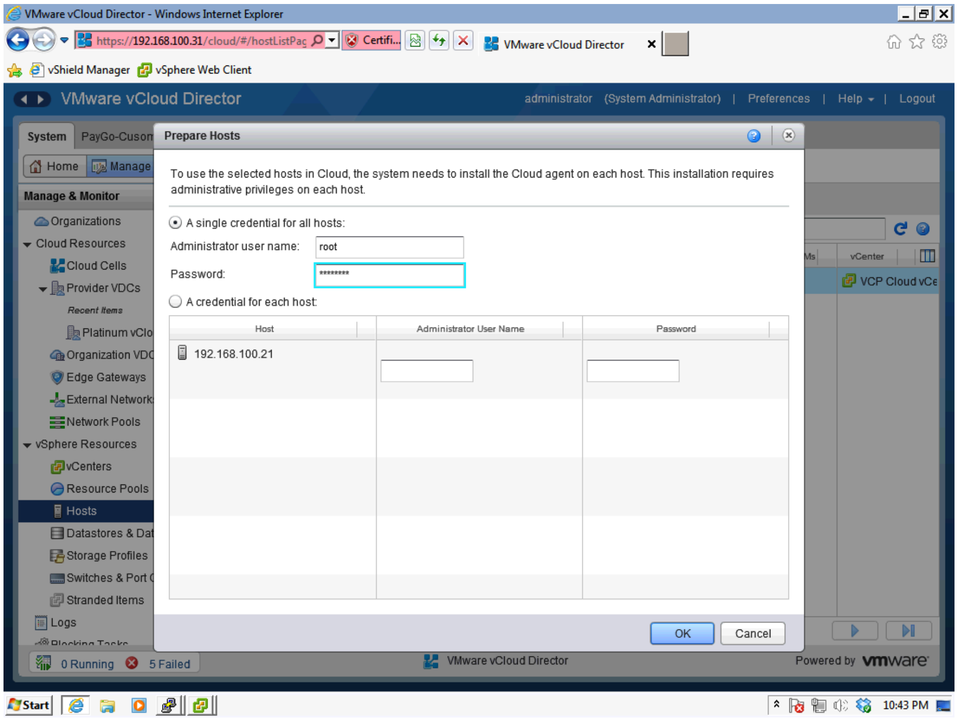

Prepare Hosts

To use the selected hosts in Cloud, the system needs to install the Cloud agent on each host. This installation requires administrative privileges on each host.

● A single credential for all hosts:

Administrator user name: root

Password: ********

○ A credential for each host:

Host	Administrator User Name	Password
192.168.100.21		

Figure 8-10 Username input

6. Verify that the hosts have been prepared properly. The hosts in Figure 8-11 have been successfully provisioned.

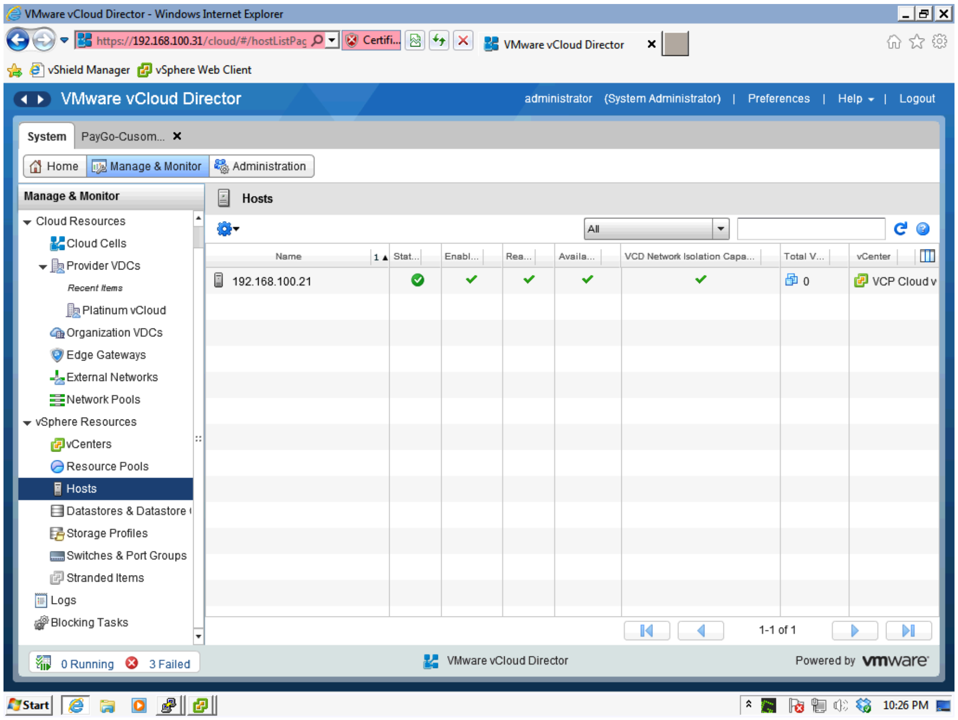

Name	Stat...	Enabl...	Rea...	Availa...	VCD Network Isolation Capa...	Total V...	vCenter
192.168.100.21	✓	✓	✓	✓	✓	0	VCP Cloud v

Figure 8-11 Hosts successfully provisioned

Enable Provider Storage

After the hosts have been enabled for the cloud, you must verify that the storage presented to those hosts is available for usage. This requires a Storage Profile to be created in vCenter Server and attached to the datastores on the hosts. Once the storage profile is created and presented, you can enable it in the Provider vDC inside of vCloud Director.

Key Topic

vCloud Director 5.1 utilizes the vCenter Server Storage Profile service for combining and utilizing datastores. A storage profile must be created and defined by a vCenter Server administrator. If the administrator neglects this step the *(Any) storage profile will be utilized, which includes any datastore that is not assigned to a storage profile. A storage profile consists of storage capabilities that can be user defined or vendor defined. If you have a large infrastructure and require the ability to segment your storage, you may need to create user defined storage capabilities and selectively assign them to storage profiles.

We will cover how to check the storage profile and how to attach it to a Provider vDC in Activity 8-3.

Key Topic

ACTIVITY 8-3 ASSIGN A STORAGE PROFILE TO A DATASTORE

1. Login to the vCenter Server that vCloud Director is attached to
2. Select **Management > VM Storage Profiles > [Datacenter Name]**
3. Click on **Create VM Storage Profile**
4. Select a name that is unique to this vCenter Server or Cluster to reduce confusion later, similar to the profile being created in Figure 8-12.

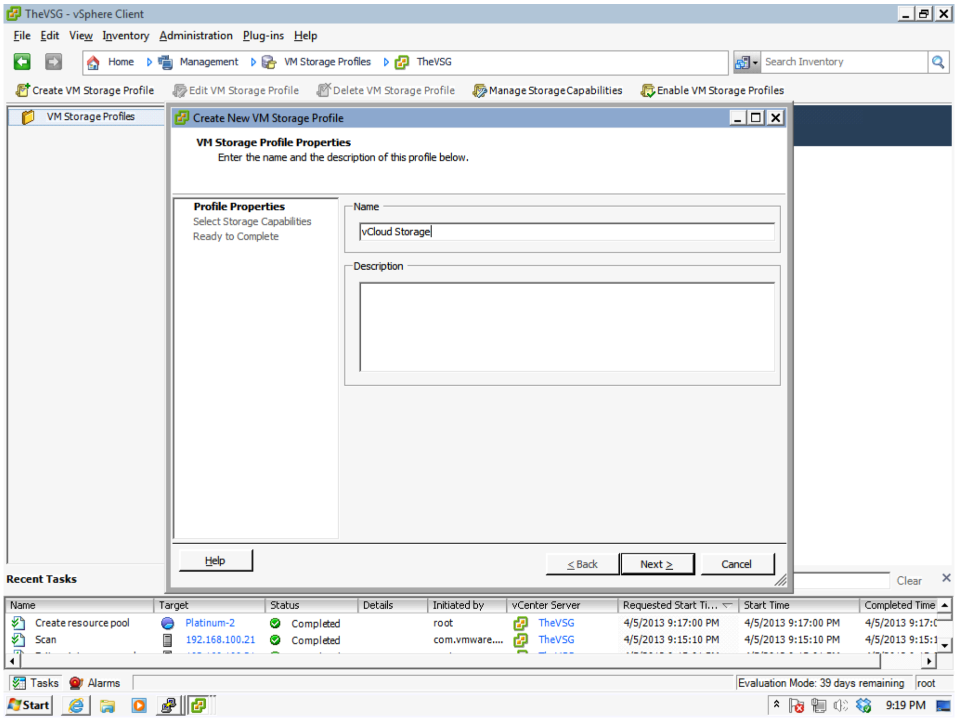

Create New VM Storage Profile

VM Storage Profile Properties

Enter the name and the description of this profile below.

Name: vCloud Storage

Description:

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Start Time	Completed Time
Create resource pool	Platinum-2	Completed		root	TheVSG	4/5/2013 9:17:00 PM	4/5/2013 9:17:00 PM	4/5/2013 9:17:0
Scan	192.168.100.21	Completed		com.vmware....	TheVSG	4/5/2013 9:15:10 PM	4/5/2013 9:15:10 PM	4/5/2013 9:15:1

Evaluation Mode: 39 days remaining | root

Figure 8-12 Storage Profile Creation in vCenter

5. Select the vendor or user defined storage capabilities for this profile. Figure 8-13 shows an NFS profile being selected.

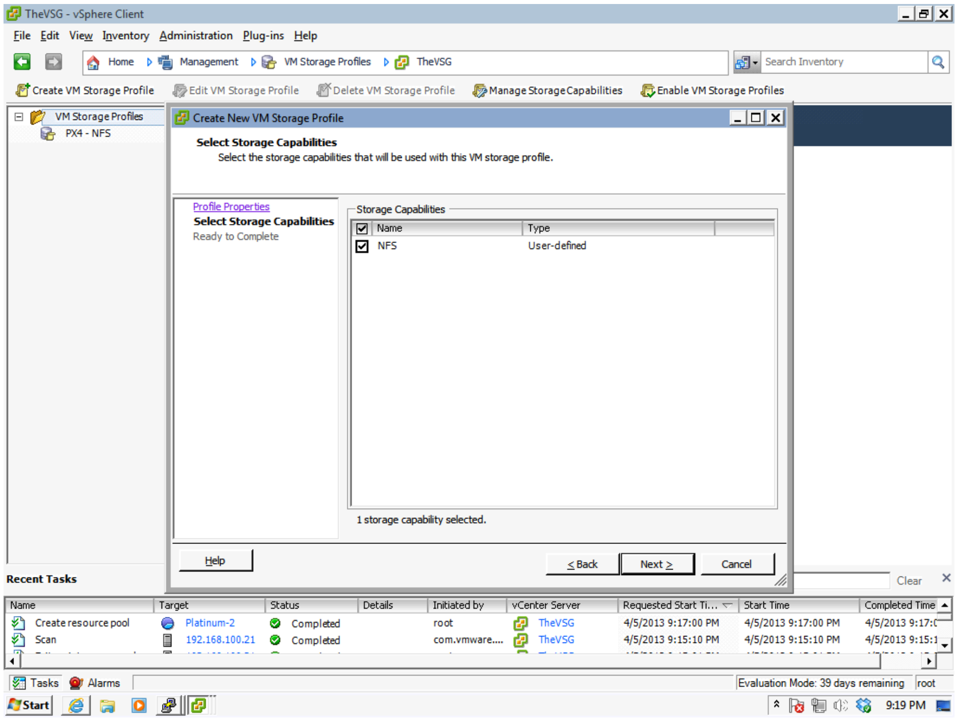

Create New VM Storage Profile

Select Storage Capabilities

Select the storage capabilities that will be used with this VM storage profile.

Storage Capabilities

☑	Name	Type
☑	NFS	User-defined

1 storage capability selected.

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Start Time	Completed Time
Create resource pool	Platinum-2	Completed		root	TheVSG	4/5/2013 9:17:00 PM	4/5/2013 9:17:00 PM	4/5/2013 9:17:0
Scan	192.168.100.21	Completed		com.vmware....	TheVSG	4/5/2013 9:15:10 PM	4/5/2013 9:15:10 PM	4/5/2013 9:15:1

Figure 8-13 Storage Capabilities in vCenter

6. Select the **Inventory** > **Datastores** and **Datastore Cluster View**. Now select each datastore and verify that the proper user-defined storage capability has been assigned to the datastores as shown in Figure 8-14.

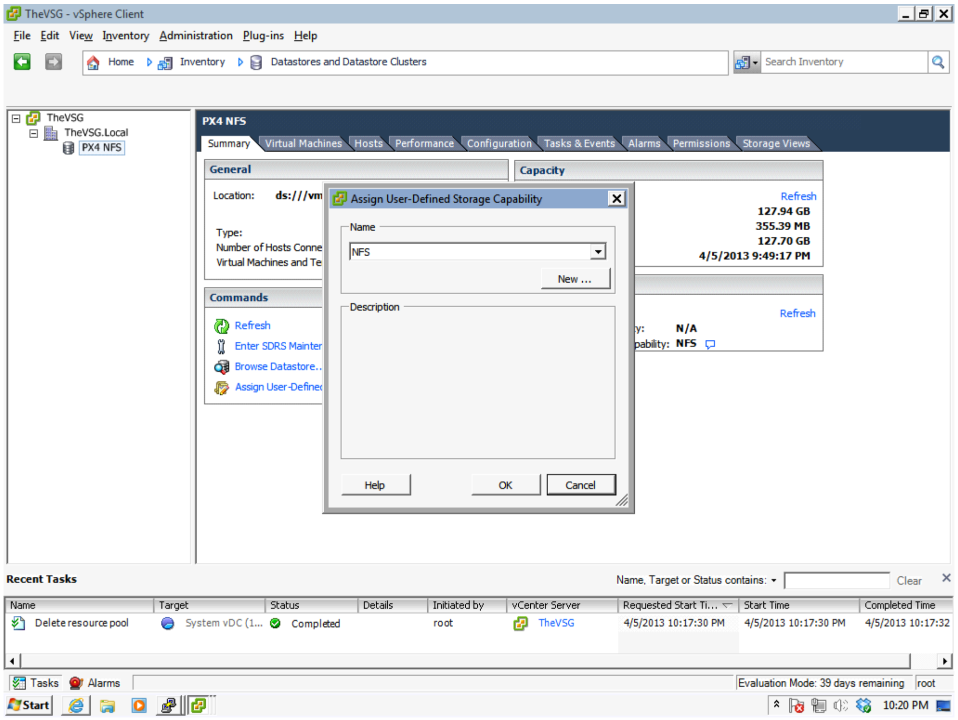

PX4 NFS

Location: ds:///vm

Type:
Number of Hosts Conne
Virtual Machines and Te

127.94 GB
355.39 MB
127.70 GB
4/5/2013 9:49:17 PM

y: N/A
pability: NFS

Assign User-Defined Storage Capability

Name
NFS

Description

Recent Tasks

Name, Target or Status contains:

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Start Time	Completed Time
Delete resource pool	System vDC (1...	Completed		root	TheVSG	4/5/2013 10:17:30 PM	4/5/2013 10:17:30 PM	4/5/2013 10:17:32

Evaluation Mode: 39 days remaining root

10:20 PM

Figure 8-14 Storage Profile Assignment

vCloud Director should now display the available storage profiles in the system administrator view. If the profiles are not visible, a problem has occurred. Troubleshooting storage profiles is covered in the next section.

Decommissioning a Provider vDC

In the life cycle of a vCloud deployment, there may come a time when you need to decommission a Provider vDC, or possibly a storage profile. This task is not as simple as just powering on the new Provider vDC and powering down the old infrastructure. The process can vary based on the hardware and the use of the hardware that is being replaced. For the purpose of the VCP-Cloud and VCP-IaaS exams, we will only cover the disabling and deleting of a Provider vDC, and the addition of a storage profile to a Provider vDC.

Disabling a Provider vDC does not power off the existing vApps, stop access to the associated network, or delete files from the storage profile. Disabling the Provider vDC stops vCloud Director from adding new workloads to the Provider vDC, and prevents the powering on of any vApp that may be powered down. It is still possible to migrate vApps or workloads off of the Provider vDC.

ACTIVITY 8-4 DISABLING A PROVIDER VDC

##key Topic

1. Open the vCloud Director UI, as shown in Figure 8-15

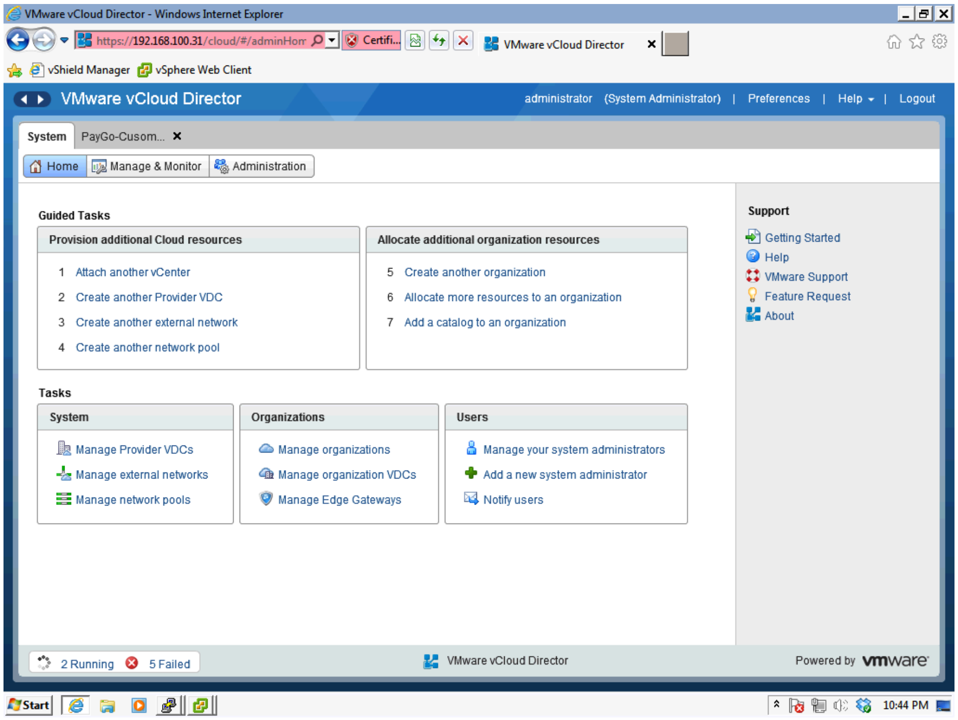

VMware vCloud Director - Windows Internet Explorer

https://192.168.100.31/cloud/#/adminHom

Certifi...

VMware vCloud Director

vShield Manager

vSphere Web Client

VMware vCloud Director

administrator (System Administrator) | Preferences | Help | Logout

System

PayGo-Cusom...

Home

Manage & Monitor

Administration

Guided Tasks

Provision additional Cloud resources

1 Attach another vCenter

2 Create another Provider VDC

3 Create another external network

4 Create another network pool

Allocate additional organization resources

5 Create another organization

6 Allocate more resources to an organization

7 Add a catalog to an organization

Tasks

System

Manage Provider VDCs

Manage external networks

Manage network pools

Organizations

Manage organizations

Manage organization VDCs

Manage Edge Gateways

Users

Manage your system administrators

Add a new system administrator

Notify users

Support

Getting Started

Help

VMware Support

Feature Request

About

2 Running

5 Failed

VMware vCloud Director

Powered by vmware

Start

10:44 PM

Figure 8-15 vCloud Director UI

2. Select the Manage & Monitor tab, then select the Provider vDCs option in the left pane as shown in Figure 8-16

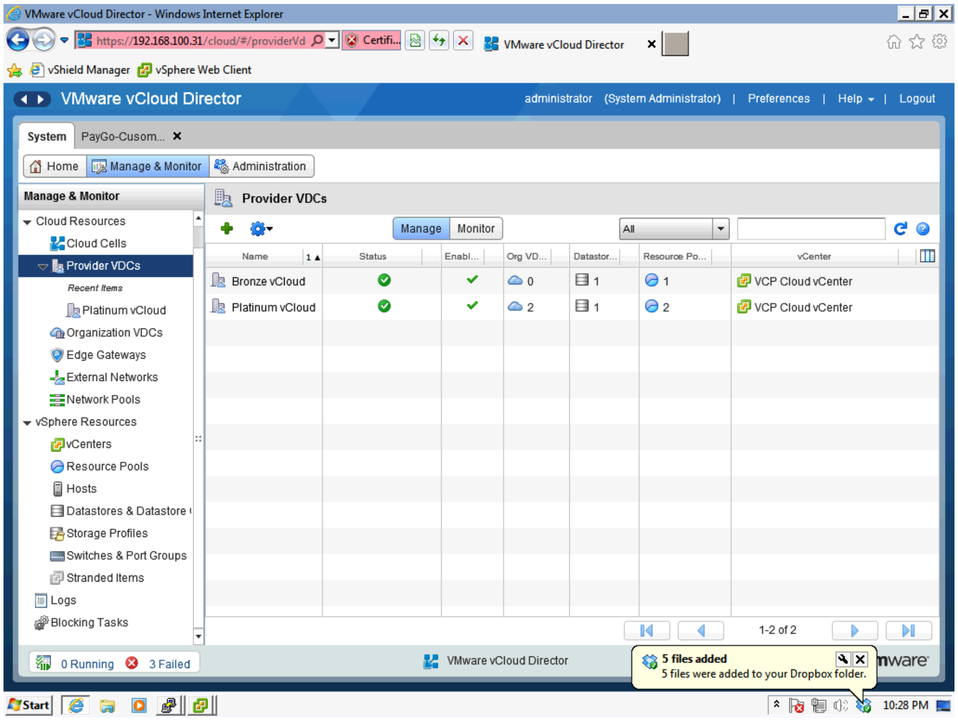

Name	Status	Enabl...	Org VD...	Datastor...	Resource Po...	vCenter
Bronze vCloud	✓	✓	0	1	1	VCP Cloud vCenter
Platinum vCloud	✓	✓	2	1	2	VCP Cloud vCenter

5 files added
5 files were added to your Dropbox folder.

Figure 8-16 Manage & Monitor tab

3. Right click on the Provider vDC that is to be disabled, such as the one shown in Figure 8-17

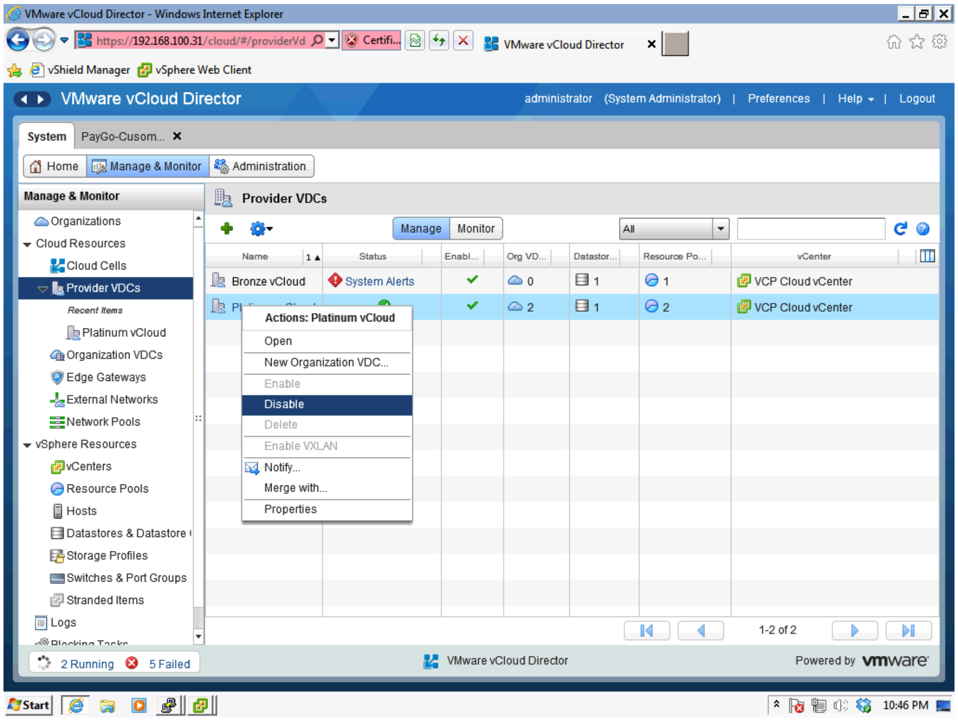

Name	Status	Enabl...	Org VD...	Datastor...	Resource Po...	vCenter
Bronze vCloud	System Alerts	✓	0	1	1	VCP Cloud vCenter
Pl	✓	✓	2	1	2	VCP Cloud vCenter

Figure 8-17 Disabling a Provider vDC

4. Select **Disable**

After the Provider vDC is disabled, you must remove the resources allocated to the Provider vDC before it can be removed from vCloud Director. These resources can include networks, vShield Edge devices, storage profiles, Org vDCs, catalogs, vApps, and Media. We will cover the removal of these items in their respective chapters/sections.

ACTIVITY 8-5 DELETE A PROVIDER VDC

1. Open the vCloud Director UI, as shown in Figure 8-18

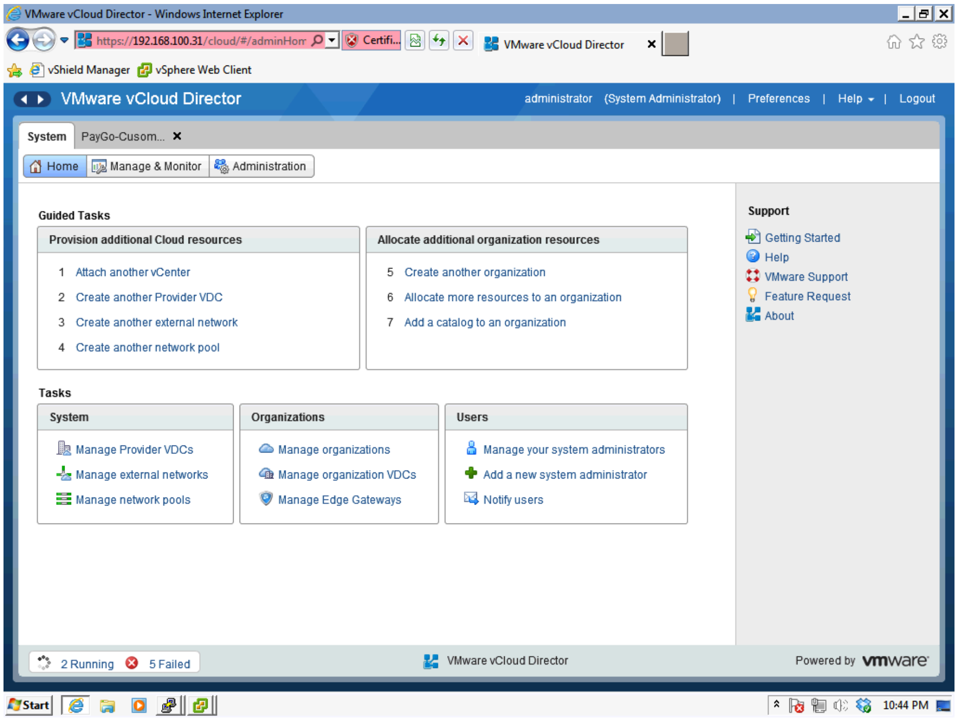
VMware vCloud Director - Windows Internet Explorer
https://192.168.100.31/cloud/#/adminHom
Certifi...
VMware vCloud Director
vShield Manager
vSphere Web Client
VMware vCloud Director
administrator (System Administrator) | Preferences | Help | Logout
System
PayGo-Cusom...
Home
Manage & Monitor
Administration
Guided Tasks
Provision additional Cloud resources
1 Attach another vCenter
2 Create another Provider VDC
3 Create another external network
4 Create another network pool
Allocate additional organization resources
5 Create another organization
6 Allocate more resources to an organization
7 Add a catalog to an organization
Tasks
System
Manage Provider VDCs
Manage external networks
Manage network pools
Organizations
Manage organizations
Manage organization VDCs
Manage Edge Gateways
Users
Manage your system administrators
Add a new system administrator
Notify users
Support
Getting Started
Help
VMware Support
Feature Request
About
2 Running
5 Failed
VMware vCloud Director
Powered by vmware
Start
10:44 PM

Figure 8-18 vCloud Director Main page

2. Select the Manage & Monitor tab, then select the Provider vDCs option in the left pane as shown in 8-19

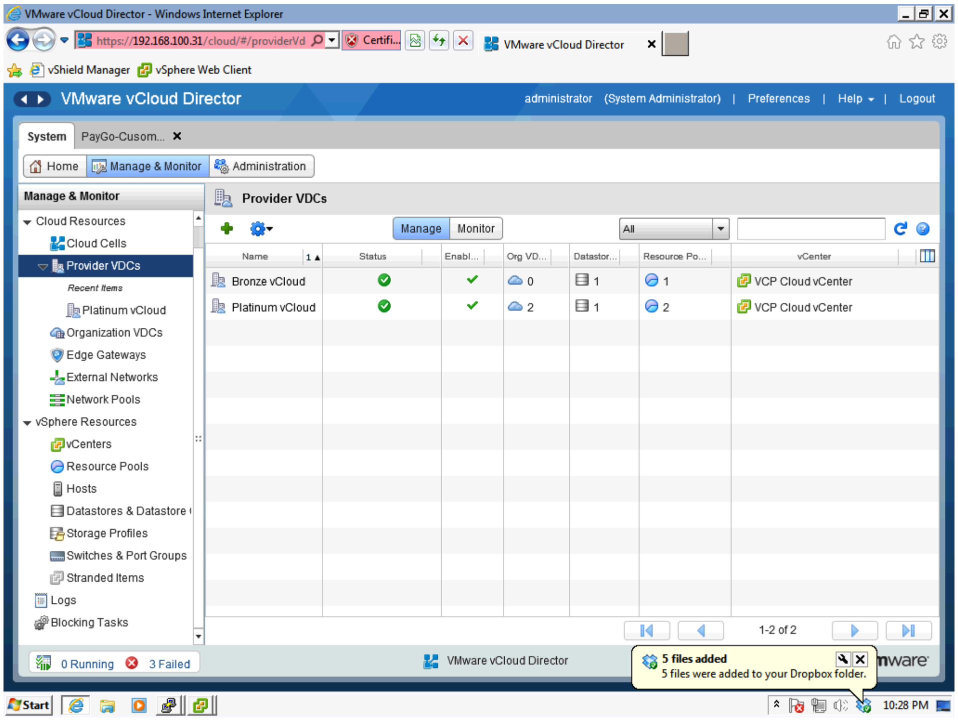

VMware vCloud Director - Windows Internet Explorer

https://192.168.100.31/cloud/#/providerVd

VMware vCloud Director

administrator (System Administrator) | Preferences | Help | Logout

System | PayGo-Cusom...

Home | Manage & Monitor | Administration

Provider VDCs

Manage | Monitor | All

Name	Status	Enabl...	Org VD...	Datastor...	Resource Po...	vCenter
Bronze vCloud	✓	✓	0	1	1	VCP Cloud vCenter
Platinum vCloud	✓	✓	2	1	2	VCP Cloud vCenter

1-2 of 2

0 Running | 3 Failed

5 files added
5 files were added to your Dropbox folder.

10:28 PM

Figure 8-19 Manage & Monitor tab

3. Right click on the Provider vDC that is to be deleted, as shown in Figure 8-20

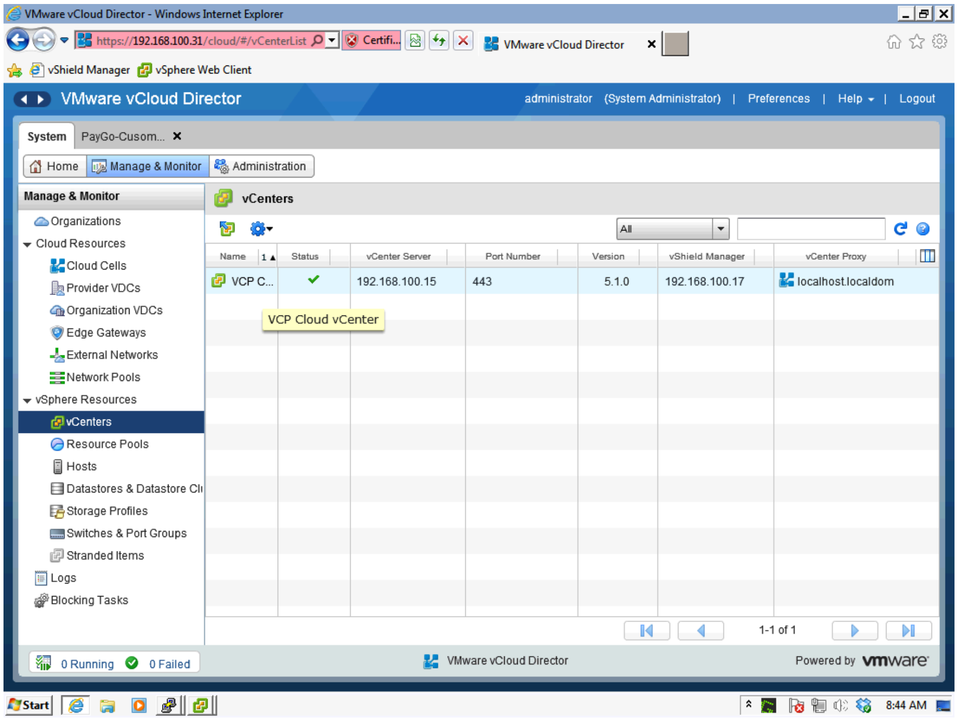

Name	Status	vCenter Server	Port Number	Version	vShield Manager	vCenter Proxy
VCP C...	✓	192.168.100.15	443	5.1.0	192.168.100.17	localhost.localdom

Figure 8-20 Deleting a Provider vDC

4. Select **Delete**

If the Provider vDC still has a resource assigned to it you will receive an error similar to that shown in Figure 8-21:

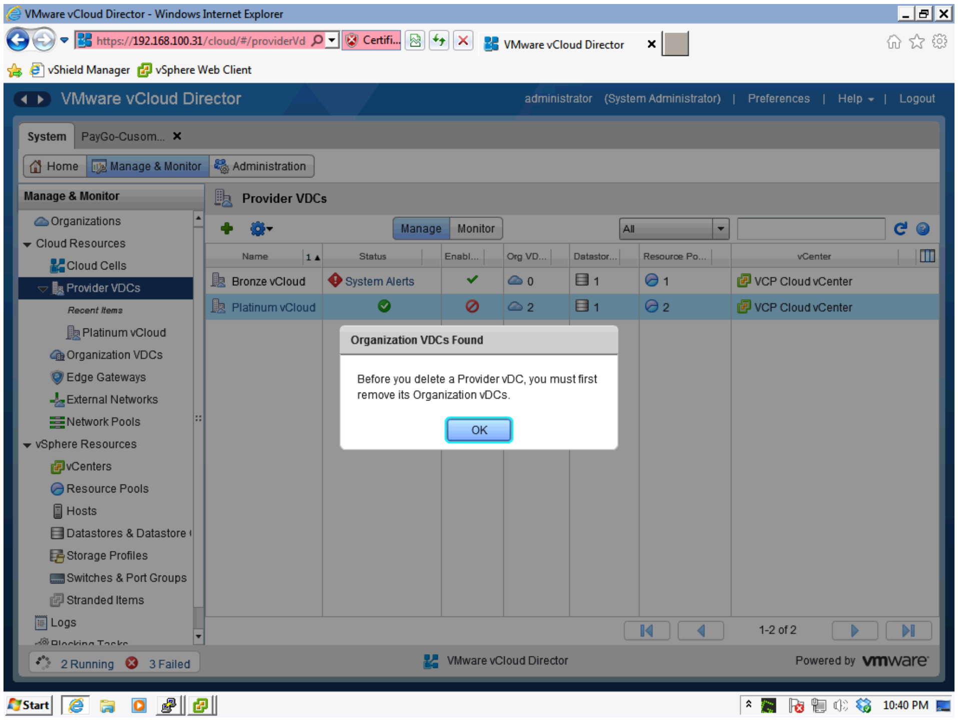

Name	Status	Enabl...	Org VD...	Datastor...	Resource Po...	vCenter
Bronze vCloud	System Alerts	✓	0	1	1	VCP Cloud vCenter
Platinum vCloud	✓	⊘	2	1	2	VCP Cloud vCenter

Organization VDCs Found

Before you delete a Provider vDC, you must first remove its Organization vDCs.

OK

Figure 8-21 Resource deletion error

If this happens, find the resource that the error mentions, remove it from the Provider vDC, then attempt the delete operation again. For example, as you can see in the error shown in Figure 8-22 there is a left over Org vDC that must be removed before the Provider vDC can be deleted.

Once all items are removed, the delete process will remove the binding of vCloud Director to the resource pool in vCenter Server. The operation does not remove the vCloud Agents from the hosts, this must be done separately. To do so, navigate to the host tab and select **disable** and **unprepare host**, as shown in Figure 8-22

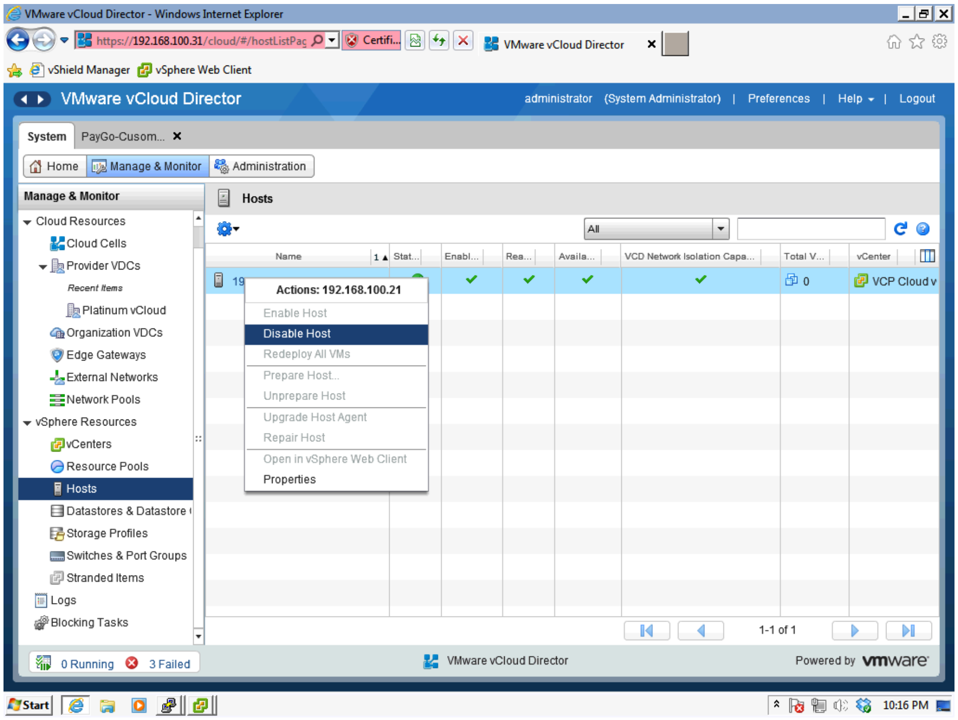

Figure 8-22 Disable and Unprepare a host

Storage Profile in vCenter

It is common for vCenter Server to take a few minutes to sync newly created and defined storage profiles. To speed this process up you can force a re-sync operation, which is performed from the system administrator view.

ACTIVITY 8-6 SYNC VCENTER STORAGE PROFILES

1. Open the vCenter Servers view from the left hand pane of the Manage & Monitor tab, as shown in Figure 8-23

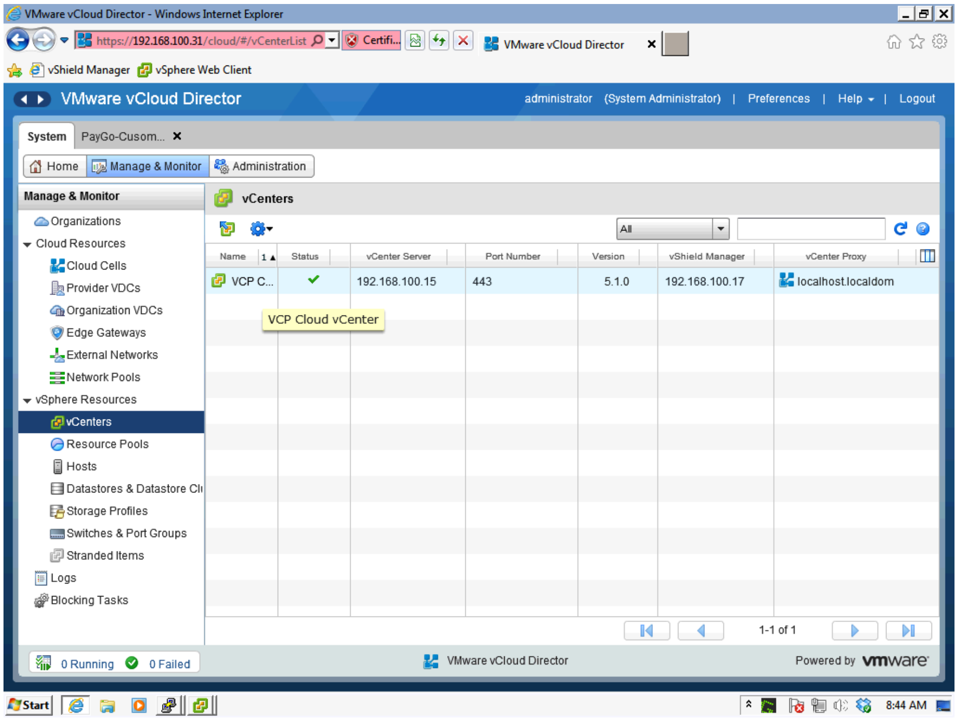

Name	Status	vCenter Server	Port Number	Version	vShield Manager	vCenter Proxy
VCP C...	✓	192.168.100.15	443	5.1.0	192.168.100.17	localhost.localdom

Figure 8-23 Manage & Monitor Tab

2. Right click on the vCenter Server where the storage profile exists and select **Refresh Storage Profiles** as shown in Figure 8-24

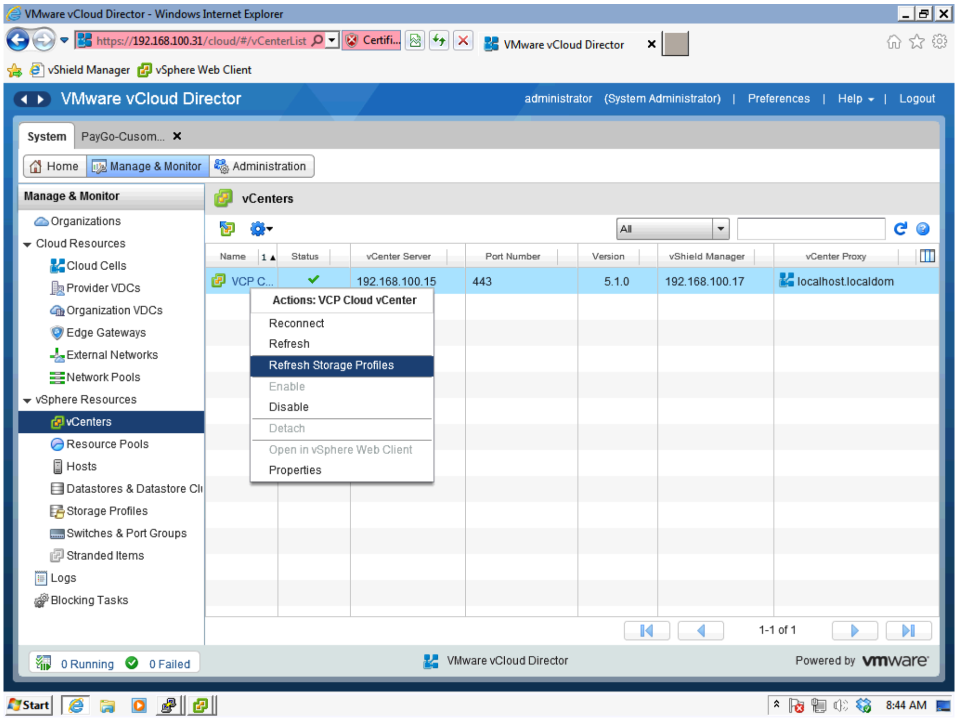

Name	1 ▲	Status	vCenter Server	Port Number	Version	vShield Manager	vCenter Proxy
VCP C...		✓	192.168.100.15	443	5.1.0	192.168.100.17	localhost.localdom

Figure 8-24 Refresh Storage Profiles

3. Verify that the task completed without error by clicking on the **Logs** menu option on the left hand pane. View the log status to ensure that the Refresh Storage Profile operation completed without error. The log should look similar to the one shown in Figure 8-25

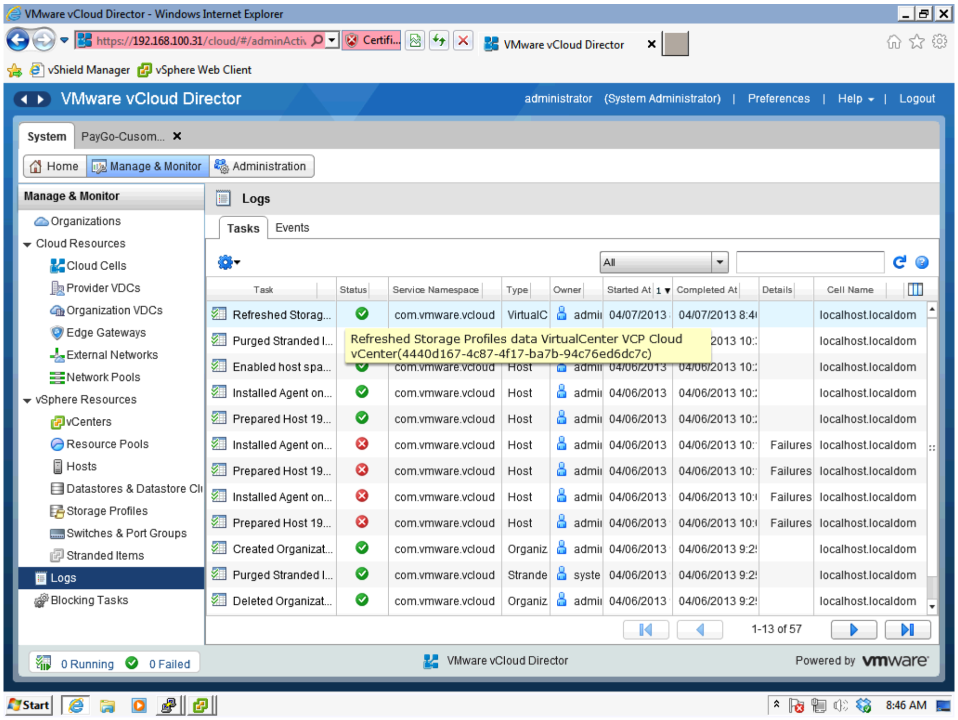

Task	Status	Service Namespace	Type	Owner	Started At 1 ▾	Completed At	Details	Cell Name
Refreshed Storag...	✓	com.vmware.vcloud	VirtualC	admi	04/07/2013	04/07/2013 8:4		localhost.localdom
Purged Stranded I...						2013 10:		localhost.localdom
Enabled host spa...	✓	com.vmware.vcloud	Host	admi	04/06/2013	04/06/2013 10:		localhost.localdom
Installed Agent on...	✓	com.vmware.vcloud	Host	admi	04/06/2013	04/06/2013 10:		localhost.localdom
Prepared Host 19...	✓	com.vmware.vcloud	Host	admi	04/06/2013	04/06/2013 10:		localhost.localdom
Installed Agent on...	✗	com.vmware.vcloud	Host	admi	04/06/2013	04/06/2013 10:	Failures	localhost.localdom
Prepared Host 19...	✗	com.vmware.vcloud	Host	admi	04/06/2013	04/06/2013 10:	Failures	localhost.localdom
Installed Agent on...	✗	com.vmware.vcloud	Host	admi	04/06/2013	04/06/2013 10:	Failures	localhost.localdom
Prepared Host 19...	✗	com.vmware.vcloud	Host	admi	04/06/2013	04/06/2013 10:	Failures	localhost.localdom
Created Organizat...	✓	com.vmware.vcloud	Organiz	admi	04/06/2013	04/06/2013 9:2		localhost.localdom
Purged Stranded I...	✓	com.vmware.vcloud	Strande	syste	04/06/2013	04/06/2013 9:2		localhost.localdom
Deleted Organizat...	✓	com.vmware.vcloud	Organiz	admi	04/06/2013	04/06/2013 9:2		localhost.localdom

Refreshed Storage Profiles data VirtualCenter VCP Cloud vCenter(4440d167-4c87-4f17-ba7b-94c76ed6dc7c)

1-13 of 57

Figure 8-25 vCloud Director Event log

Create and Administer Organization vDCs

Creating a Provider vDC provides resources that can be provisioned to the organizations in the cloud. This is accomplished through the creation of organization virtual datacenters (Org vDCs). An Org vDC is how vCloud Director defines a space for vApps to be stored and ran, as well as a space for CD-ROM (.iso) and floppy image (.flp) files to be stored. The Org vDC provisions resources to an organization (or group within an organization) by defining the CPU, RAM, storage and network resources that will be available to vApps within the Org vDC.

When creating an Org vDC, an important step is choosing the appropriate allocation model. The allocation model determines how resources are allocated by the Org vDC. There are three different allocation models, Pay as You Go, Allocation Pool, and Reservation Pool. An organization’s SLAs, OLAs and other agreements typically determine the allocation model that will be used by the Org vDC. The three allocation models provision resources as follows:

- PAYG
  - Resources are not allocated in advance

  - Resources are reserved as vApps are powered on
  - A percentage of resources can be reserved
  - vCPU speed rating is adjustable

- Allocation Pool

  - A percentage of resources is allocated in advance
  - This percentage of resources is reserved
  - Cloud administrators control the amount of over commit
  - Resources are sharable between Org vDCs

- Reservation Pool

  - All resources are allocated and reserved in advance
  - Users are able to adjust individual VM priorities and resources
  - No sharing of resources with other Org vDCs

##Key Topic

Pay as You Go

The Pay as You Go (commonly referred to as 'PAYG') allocation model most closely resembles the Amazon EC2 model, where the compute resources are defined on a per VM basis. PAYG allows for performance guarantees on each VM as well as limits for each VM. The benefit of PAYG is that the provider vDC's resources are only consumed when a VM or vApp is powered on.

NOTE PAYG is comparable to how an electric company charges for power. When the lights are on you are being charged, when they are off there is no charge.

When a VM or a vApp is powered on, vCloud Director will place the vApp into the provider vDC that is assigned to the organization vDC. Once the VM is placed into a resource pool it is configured with the limits and guarantees that are defined by the organization vDC.

PAYG allocation models are unique in that they allow you to place limits on a single VM and not the entire resource pool, while still allowing for a % reservation. A PAYG model is ideal for use cases where noisy neighbors are a concern. PAYG will allow you, the cloud administrator, to limit the amount of resources a single VM can consume.

It should be noted that there is some risk with using a PAYG allocation model. Because there are no pre-reserved or committed resources, it is possible that a VM will be unable to power on if the provider vDC is out of resources.

At some point an Org vDC configured with the PAYG allocation model may need to be re-sized. This can be done, but it does require all vApps to be powered down and backed up before the adjustment can be made. Figure 8-26 shows the Pay as You Go Org vDC properties.

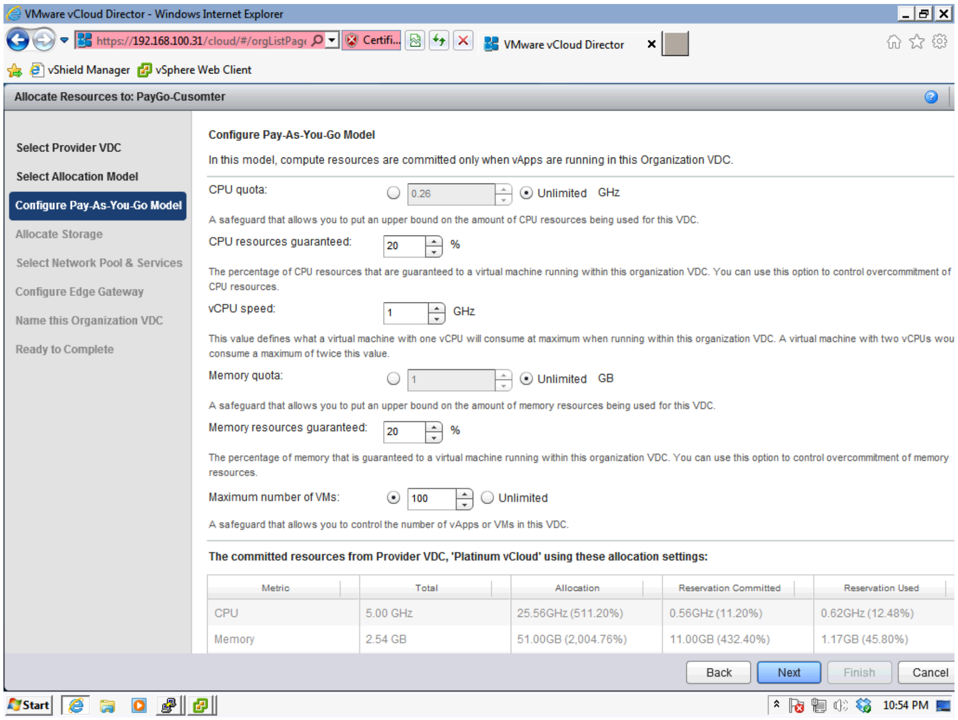

Allocate Resources to: PayGo-Cusomter

Configure Pay-As-You-Go Model

In this model, compute resources are committed only when vApps are running in this Organization VDC.

CPU quota: ○ 0.26 ● Unlimited GHz

A safeguard that allows you to put an upper bound on the amount of CPU resources being used for this VDC.

CPU resources guaranteed: 20 %

The percentage of CPU resources that are guaranteed to a virtual machine running within this organization VDC. You can use this option to control overcommitment of CPU resources.

vCPU speed: 1 GHz

This value defines what a virtual machine with one vCPU will consume at maximum when running within this organization VDC. A virtual machine with two vCPUs wou consume a maximum of twice this value.

Memory quota: ○ 1 ● Unlimited GB

A safeguard that allows you to put an upper bound on the amount of memory resources being used for this VDC.

Memory resources guaranteed: 20 %

The percentage of memory that is guaranteed to a virtual machine running within this organization VDC. You can use this option to control overcommitment of memory resources.

Maximum number of VMs: ● 100 ○ Unlimited

A safeguard that allows you to control the number of vApps or VMs in this VDC.

The committed resources from Provider VDC, 'Platinum vCloud' using these allocation settings:

Metric	Total	Allocation	Reservation Committed	Reservation Used
CPU	5.00 GHz	25.56GHz (511.20%)	0.56GHz (11.20%)	0.62GHz (12.48%)
Memory	2.54 GB	51.00GB (2,004.76%)	11.00GB (432.40%)	1.17GB (45.80%)

Figure 8-26 Pay as You Go Org vDC properties

During the creation of a PAYG allocation model you will define:

- **CPU Quota** – This quota is the limit of CPU resources that an organization can consume at any time. This limit can be higher than the actual resources presented, which gives a cloud administrator the ability to over commit CPU resources provided by the Org vDC.
- **CPU Resources Guaranteed** – vCenter Server defines this as the reservation on the virtual machine.

NOTE: This may affect the slot sizing for HA calculations in vCenter Server.

- **vCPU Speed** - In vCenter Server this is defined as the limit of a virtual machine’s CPU resource. This is a per-CPU value, so the number placed in this field will become the limit on the CPU of the virtual machine multiplied by the number of vCPUs assigned.
- **Memory Quota** – Much like the CPU quota, this is the limit of memory for an Org vDC. This limit will most likely be higher than the physical memory in the

backing Provider vDCs. Features built in to vSphere such as transparent page sharing, compression, swapping, and ballooning make over allocating memory not only possible, but highly efficient.

- **Memory Resources Guaranteed –** The guarantee here is placed on the virtual machine itself, meaning that a virtual machine will always have X amount of RAM available.

NOTE: This may affect the slot sizing for HA calculations in vCenter.

- **Maximum number of VMs** – It may seem that this field would only apply to powered on virtual machines or deployed virtual machines, but this is not the case. For the purpose of this value, a virtual machine is counted as ANY virtual machine in the Org vDC. This included catalog items, powered off and powered on virtual machines.

To bring all of these concepts together, consider the following example. An Org vDC with a CPU reservation of 20% and a vCPU speed of 2.5GHZ will produce a limit of 5GHz with 1GHz being reserved on a 2 vCPU virtual machine. That same VM with 12GB of assigned ram and a 10% guarantee will have a 1.2GB reservation for RAM (plus the virtualization overhead).

The PAYG allocation method can also present a billing problem to end users in a service provider model. The billing process for PAYG varies and can be unpredictable from month to month. That being said, PAYG does provide the best method of accounting for the actual resources utilized by an organization, as both of the other allocation models require a pre-purchased amount of resources.

##Key Topic

Allocation Pool

An allocation pool would be the second most commonly used allocation model in vCloud Director deployments. An allocation pool model gives you the ability to overcommit resources while still limiting the amount of resources that an organization consumes.

Much like PAYG, an allocation pool model can be ‘elastic’. This means that it can span multiple resource pools inside a single provider vDC. While providing a great deal of flexibility, this can complicate the troubleshooting of resources should a resource related problem arise..

For example, assume you have a provider vDC with multiple resource pools, with multiple organization vDCs provisioned. It is very possible that a vApp could span resource pools, which could result in decent performance for virtual machines in one resource pool but heavily overcommitted virtual machines in another resource pool. Since DRS does not balance virtual machines between resource pools, and vCloud Director's placement engine only runs on power on, this is a very real concern and possibility.

Unlike a PAYG model, the allocation pool model allows you to adjust resources and their commitments without causing a need to redeploy virtual machines. Figure 8-27 shows the Allocation Org vDC properties page.

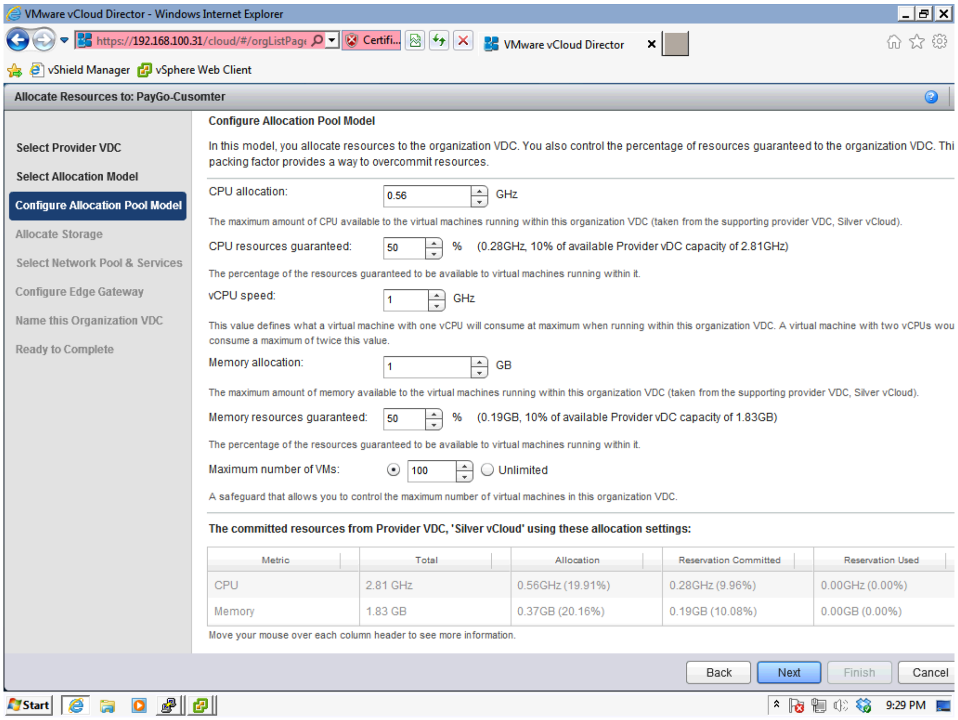
VMware vCloud Director - Windows Internet Explorer
https://192.168.100.31/cloud/#/orgListPage
Certifi...
VMware vCloud Director
vShield Manager
vSphere Web Client
Start
9:29 PM

Allocate Resources to: PayGo-Cusomter

Select Provider VDC
Select Allocation Model
Configure Allocation Pool Model
Allocate Storage
Select Network Pool & Services
Configure Edge Gateway
Name this Organization VDC
Ready to Complete

Configure Allocation Pool Model

In this model, you allocate resources to the organization VDC. You also control the percentage of resources guaranteed to the organization VDC. Thi packing factor provides a way to overcommit resources.

CPU allocation: 0.56 GHz

The maximum amount of CPU available to the virtual machines running within this organization VDC (taken from the supporting provider VDC, Silver vCloud).

CPU resources guaranteed: 50 % (0.28GHz, 10% of available Provider vDC capacity of 2.81GHz)

The percentage of the resources guaranteed to be available to virtual machines running within it.

vCPU speed: 1 GHz

This value defines what a virtual machine with one vCPU will consume at maximum when running within this organization VDC. A virtual machine with two vCPUs wou consume a maximum of twice this value.

Memory allocation: 1 GB

The maximum amount of memory available to the virtual machines running within this organization VDC (taken from the supporting provider VDC, Silver vCloud).

Memory resources guaranteed: 50 % (0.19GB, 10% of available Provider vDC capacity of 1.83GB)

The percentage of the resources guaranteed to be available to virtual machines running within it.

Maximum number of VMs: ● 100 ○ Unlimited

A safeguard that allows you to control the maximum number of virtual machines in this organization VDC.

The committed resources from Provider VDC, 'Silver vCloud' using these allocation settings:

Metric	Total	Allocation	Reservation Committed	Reservation Used
CPU	2.81 GHz	0.56GHz (19.91%)	0.28GHz (9.96%)	0.00GHz (0.00%)
Memory	1.83 GB	0.37GB (20.16%)	0.19GB (10.08%)	0.00GB (0.00%)

Move your mouse over each column header to see more information.

Figure 8-27 Allocation Model Org vDC properties

During the creation of an Org vDC with the allocation pool model, you define the following:

- **CPU Allocation** – Like the Pay as You Go allocation model, this is the limit of resources for the entire Org vDC. This limit is placed on the resource pool that was created by vCloud Director in vCenter Server. By placing a number higher than the physical resources, you are able to over commit the CPU resources in the Org vDC.
- **CPU Resources Guaranteed** – Unlike the Pay as You Go model, the CPU resource guarantee is not defined on the individual virtual machine. This reservation is placed on the Org vDC's resource pool in vCenter Server. What an allocation pool model does not protect against is the possibility of a single virtual machine compromising the rest of the machines in the Org vDC. That

said, this model will always guarantee resources to the Org vDC even if other tenants or Org vDCs in the Provider vDC are consuming all of their resources.

This guarantee is calculated using the vCPU speed that defined by the CPU allocation. For example, a vCPU speed of .5GHZ and a 10% reservation will result in a 50MHz reservation on the Org vDC resource pool.

- **vCPU Speed** – Also unlike the vCPU speed in the Pay as You Go model, the vCPU speed for an allocation pool Org vDC defines a value used to calculate the amount of resources to reserve. This does not place a limit on the VM or the resource pool. The vCPU speed is only used to calculate the reservation that should be placed on the Org vDC resource pool.
- **Memory Allocation** – Much like the Pay as You Go model this defines the maximum amount of RAM that can be allocated to virtual machines. Even if a virtual machine is only using 1GB of RAM, if it is configured with 100GB of RAM, the virtual machine will consume 100GB of RAM out of the allocation pool.
- **Memory Resources Guaranteed** – Like the CPU resources guaranteed, this value is used to calculate how much RAM is reserved for the virtual machine. Unlike the Pay as you Go model, this reservation is not placed on the virtual machine itself, but on the Org vDC's resource pool in vCenter Server.
- **Maximum number of VMs** – This field is the same as Pay as You Go. Any VM that is defined, including catalog items, powered off VMs and powered on VMs are included in this value.

To bring all of the concepts together for an allocation pool model Org vDC, consider the following example. Assume you have an allocation pool Org vDC defined with an allocation of 100GHz, vCPU Speed of 2GHz, and a CPU guarantee of 50%. For each 2vCPU virtual machine that is powered on in the Org vDC, vCloud Director will reserve 2GHz of compute power for that virtual machine.

The same applies to the memory allocation. Let's say the allocation pool Org vDC is configured with 200GB of memory allocation and 50% guaranteed. A virtual machine with a memory allocation of 50GB will receive an allocation from vCloud Director of 50GB from the 200GB available to the Org vDC, and set a reservation on the vCenter Server resource pool of 25GB.

The allocation pool model allows for more predictable billing for customers of service providers. It also allows for a service provider's charge in an overage billing model, much like a 95th percentile-billing model.

A benefit of the Allocation Pool and Pay as You Go allocation models is that they are dynamic. For each virtual machine that is powered on, vCloud Director calculates the settings for the resource pool and re-applies the proper settings. This is different than the process used by the final allocation model, the Reservation Pool model, which we will discuss next.

##Key Topic

Reservation Pool Model

A reservation pool model is the easiest to define in terms of physical hardware. It is also the easiest to explain to someone new to virtualization or to the over commitment of hardware.

A reservation pool allocation model is exactly what it sounds like, reserved capacity for the Org vDC. This model allows a tenant to pay for 10GHz and 100GB of compute power, and that is exactly what they get, no more and no less. The tenant can then control which VMs get access to those resources through shares and limits on the individual VMs.

While this allows for easier accounting to the provider on resources (since there is no sharing of resources between tenants), it does present the possibility of wasted resources since an idle tenant's resources cannot be shared with other tenants that may need them.

An Org vDC configured with the allocation pool model cannot be used with elastic provider vDCs. If the Org vDC is assigned to a Provider vDC with multiple resource pools, the Org vDC will consume the first resource pool listed and will not expand or move to other resource pools listed. Figure 8-28 shows the Reservation Pool model Org vDC properties.

For billing purposes when using vCenter Chargeback, the tenant's bill will be the same every month. With the reservation pool allocation model, the resources are reserved for the tenant whether they use them or not.

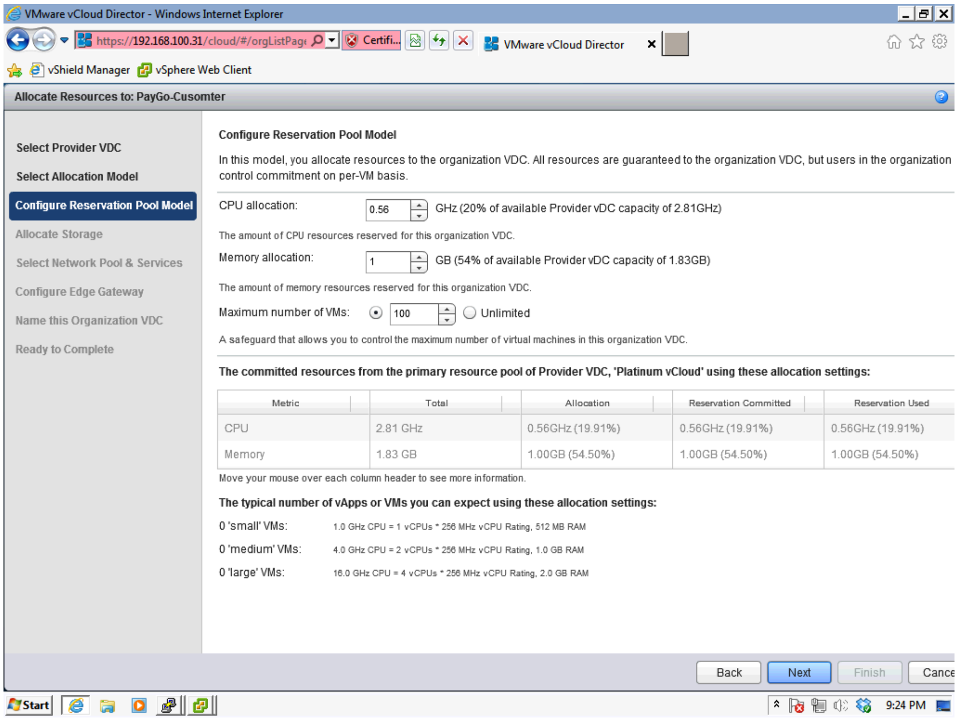
VMware vCloud Director - Windows Internet Explorer
https://192.168.100.31/cloud/#/orgListPage
Certifi...
VMware vCloud Director
vShield Manager
vSphere Web Client

Allocate Resources to: PayGo-Cusomter

Select Provider VDC
Select Allocation Model
Configure Reservation Pool Model
Allocate Storage
Select Network Pool & Services
Configure Edge Gateway
Name this Organization VDC
Ready to Complete

Configure Reservation Pool Model

In this model, you allocate resources to the organization VDC. All resources are guaranteed to the organization VDC, but users in the organization control commitment on per-VM basis.

CPU allocation: 0.56 GHz (20% of available Provider vDC capacity of 2.81GHz)

The amount of CPU resources reserved for this organization VDC.

Memory allocation: 1 GB (54% of available Provider vDC capacity of 1.83GB)

The amount of memory resources reserved for this organization VDC.

Maximum number of VMs: ● 100 ○ Unlimited

A safeguard that allows you to control the maximum number of virtual machines in this organization VDC.

The committed resources from the primary resource pool of Provider VDC, 'Platinum vCloud' using these allocation settings:

Metric	Total	Allocation	Reservation Committed	Reservation Used
CPU	2.81 GHz	0.56GHz (19.91%)	0.56GHz (19.91%)	0.56GHz (19.91%)
Memory	1.83 GB	1.00GB (54.50%)	1.00GB (54.50%)	1.00GB (54.50%)

Move your mouse over each column header to see more information.

The typical number of vApps or VMs you can expect using these allocation settings:

0 'small' VMs:	1.0 GHz CPU = 1 vCPUs * 256 MHz vCPU Rating, 512 MB RAM
0 'medium' VMs:	4.0 GHz CPU = 2 vCPUs * 256 MHz vCPU Rating, 1.0 GB RAM
0 'large' VMs:	16.0 GHz CPU = 4 vCPUs * 256 MHz vCPU Rating, 2.0 GB RAM

Figure 8-28 Reservation model Org vDC properties

When configuring a reservation pool Org vDC you will define the following:

- **CPU Allocation** – The allocation number that is defined here becomes the limit and the reservation of CPU resources in the vCenter Server Resource Pool.
- **Memory Allocation** – Like the CPU allocation, the allocation provided here becomes the limit and reservation for memory resources in the vCenter Server Resource Pool.
- **Maximum number of VMs** – In the Allocation Pool model and Pay as You Go model, this is the maximum number of virtual machines that can belong to the Org vDC. Running VMs, catalog items, and VMs deployed but not powered on all count.

NOTE: Unlike the allocation pool and pay as you go models, the reservation pool model reserves all of its resources at creation, even if there are no powered on virtual machines. Use caution when specifying numbers, as it is possible to reserve all resources in a vCenter Server resource pool, potentially starving other Org vDCs of their resources.

Selecting the Right Allocation Model

There is not a perfect allocation model for every use case. In fact, most use cases fit well into more than one possible allocation model.

PAYG is traditionally used for transient work or virtual machines that have short life expectancies. In a hosting environment the PAYG allocation model can cause variable billing. The allocation pool model is also commonly used in an environment where workloads come and go, but also works well for environments where the tenant is unsure of their true demand.

The reservation pool model is typically used as premium service, since there is no sharing of resources. Compute power that is assigned to a reservation pool Org vDC is reserved for the tenant even if they do not use it. The reservation pool model can be considered the same as dedicating hardware to a tenant.

Create an Organization vDC

Now that we have discussed the allocation models in vCloud Director and their intended uses, we will use the information that you gathered previously to create an Org vDC.

ACTIVITY 8-7 CREATING AN ORG VDC

1. Login to vCloud Director as a System Administrator

NOTE: When configuring the consumption of compute and storage resources, you must be logged in as a system administrator. Organization users do not have the permissions to create or alter these settings.

2. Select the organization for the new organization virtual datacenter. In this example, ‘OrgA’ has been selected, as shown in Figure 8-29.

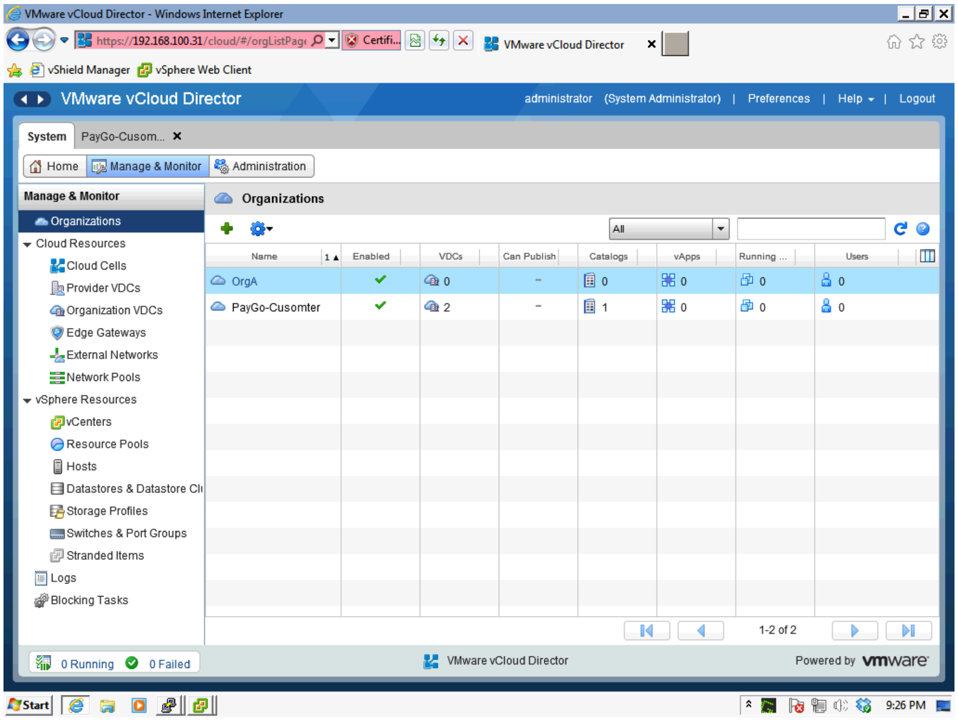

Name	Enabled	VDCs	Can Publish	Catalogs	vApps	Running ...	Users
OrgA	✓	0	–	0	0	0	0
PayGo-Cusomter	✓	2	–	1	0	0	0

Figure 8-29 Organization Selection

3. Select the Administration tab for the organization, as shown in Figure 8-30

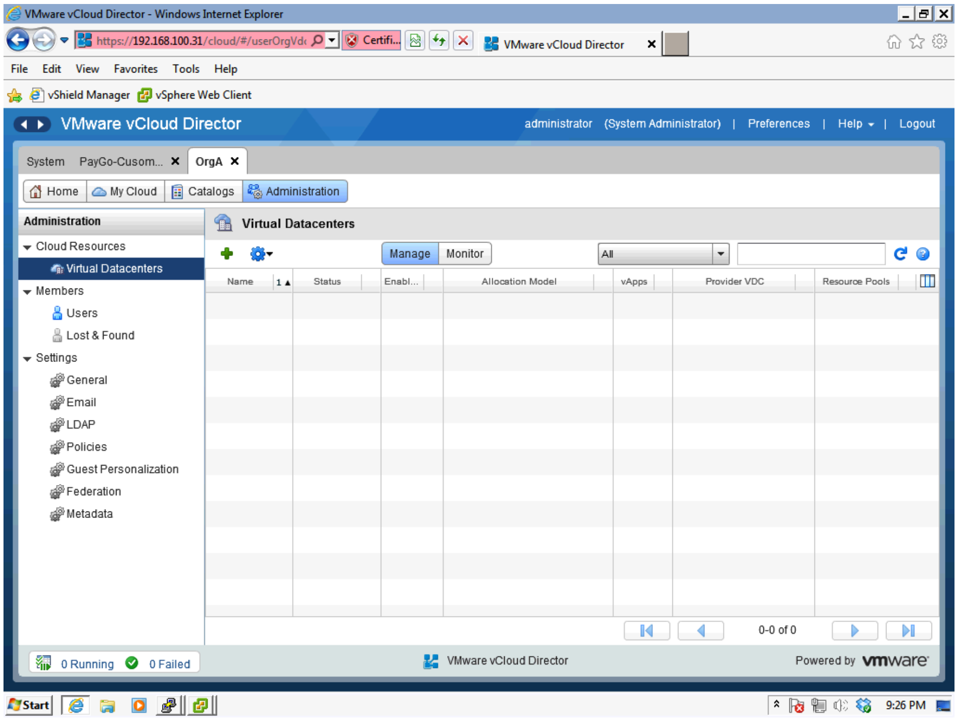

Figure 8-30 Administration Tab

4. Click on the **green** + symbol highlighted in Figure 8-31. This will start the Add Resources wizard.

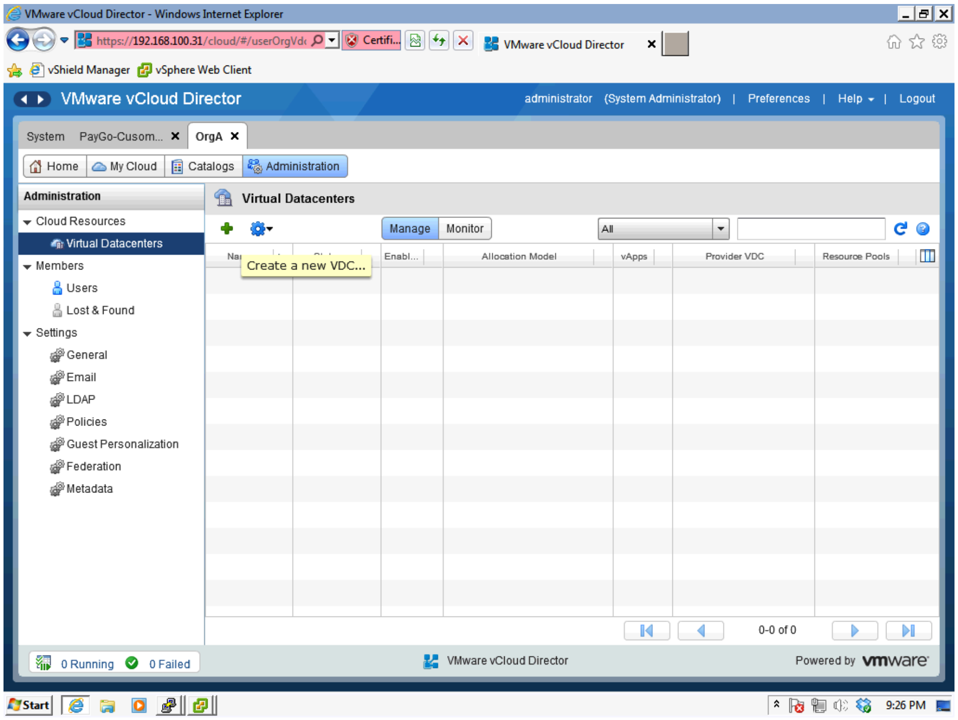

Figure 8-31 Starting the Add Resources Wizard

5. Select the provider virtual datacenter that this organization virtual datacenter will be created in. Figure 8-32 shows the Platinum Provider vDC being selected.

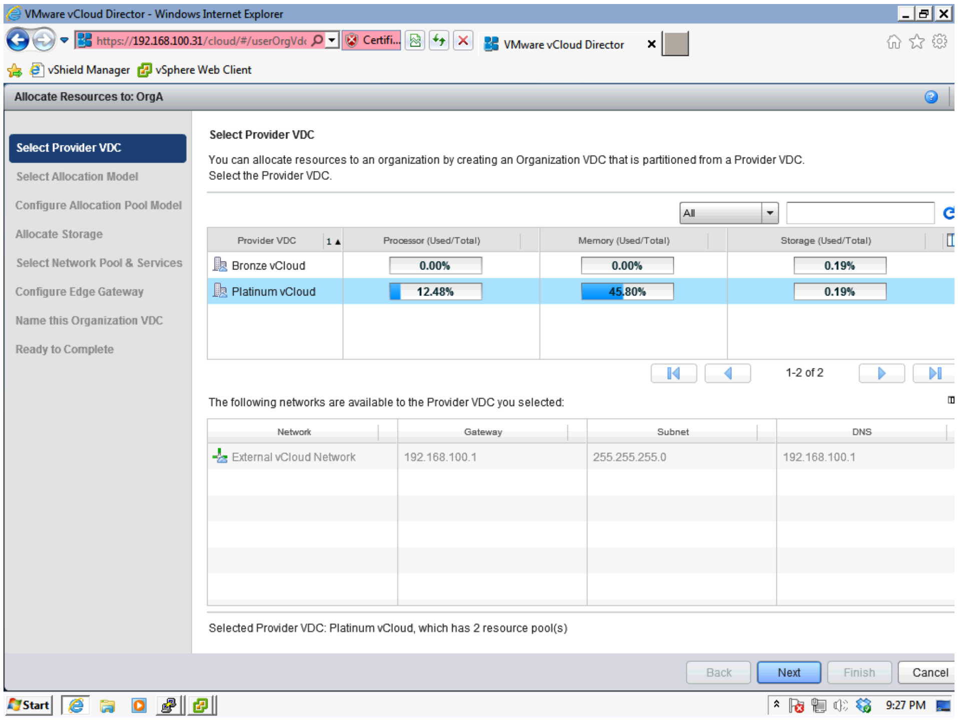
VMware vCloud Director - Windows Internet Explorer
https://192.168.100.31/cloud/#/userOrgVdc
Certifi...
VMware vCloud Director
vShield Manager
vSphere Web Client
Allocate Resources to: OrgA
Select Provider VDC
Select Allocation Model
Configure Allocation Pool Model
Allocate Storage
Select Network Pool & Services
Configure Edge Gateway
Name this Organization VDC
Ready to Complete
All
1-2 of 2
Back
Next
Finish
Cancel
Start
9:27 PM

Select Provider VDC

You can allocate resources to an organization by creating an Organization VDC that is partitioned from a Provider VDC.
Select the Provider VDC.

Provider VDC	Processor (Used/Total)	Memory (Used/Total)	Storage (Used/Total)
Bronze vCloud	0.00%	0.00%	0.19%
Platinum vCloud	12.48%	45.80%	0.19%

The following networks are available to the Provider VDC you selected:

Network	Gateway	Subnet	DNS
External vCloud Network	192.168.100.1	255.255.255.0	192.168.100.1

Selected Provider VDC: Platinum vCloud, which has 2 resource pool(s)

Figure 8-32 Selected the source Provider vDC

6. Select the allocation model that you have chosen for this Org vDC, as shown in Figure 8-33

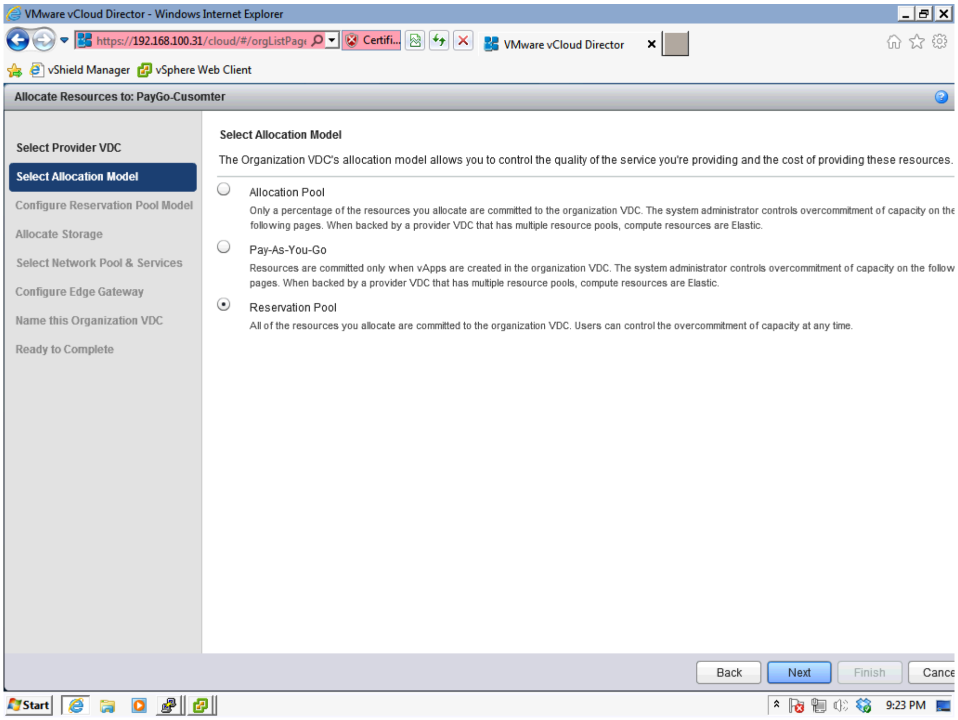
VMware vCloud Director - Windows Internet Explorer
https://192.168.100.31/cloud/#/orgListPage
Certifi...
VMware vCloud Director
vShield Manager
vSphere Web Client
Start
9:23 PM

Allocate Resources to: PayGo-Cusomter

Select Provider VDC
Select Allocation Model
Configure Reservation Pool Model
Allocate Storage
Select Network Pool & Services
Configure Edge Gateway
Name this Organization VDC
Ready to Complete

Select Allocation Model

The Organization VDC's allocation model allows you to control the quality of the service you're providing and the cost of providing these resources.

○ Allocation Pool

Only a percentage of the resources you allocate are committed to the organization VDC. The system administrator controls overcommitment of capacity on the following pages. When backed by a provider VDC that has multiple resource pools, compute resources are Elastic.

○ Pay-As-You-Go

Resources are committed only when vApps are created in the organization VDC. The system administrator controls overcommitment of capacity on the follow pages. When backed by a provider VDC that has multiple resource pools, compute resources are Elastic.

● Reservation Pool

All of the resources you allocate are committed to the organization VDC. Users can control the overcommitment of capacity at any time.

Figure 8-33 Org vDC Type Selection

7. Populate the allocation model fields as required for the allocation model that you have selected. As a guide, refer to the previous section for information on each of the Org vDC models and their required fields.
8. Select the storage profile that the Org vDC will use for the workloads.

NOTE: The any profile in vCenter Server will select Any storage available on the cluster, including local storage. If local storage is selected, vMotion and HA will not function for those virtual machines placed on local storage. To avoid this, simply create a storage profile and assign the proper datastores to the storage profile. Figure 8-34 shows the Any profile being selected.

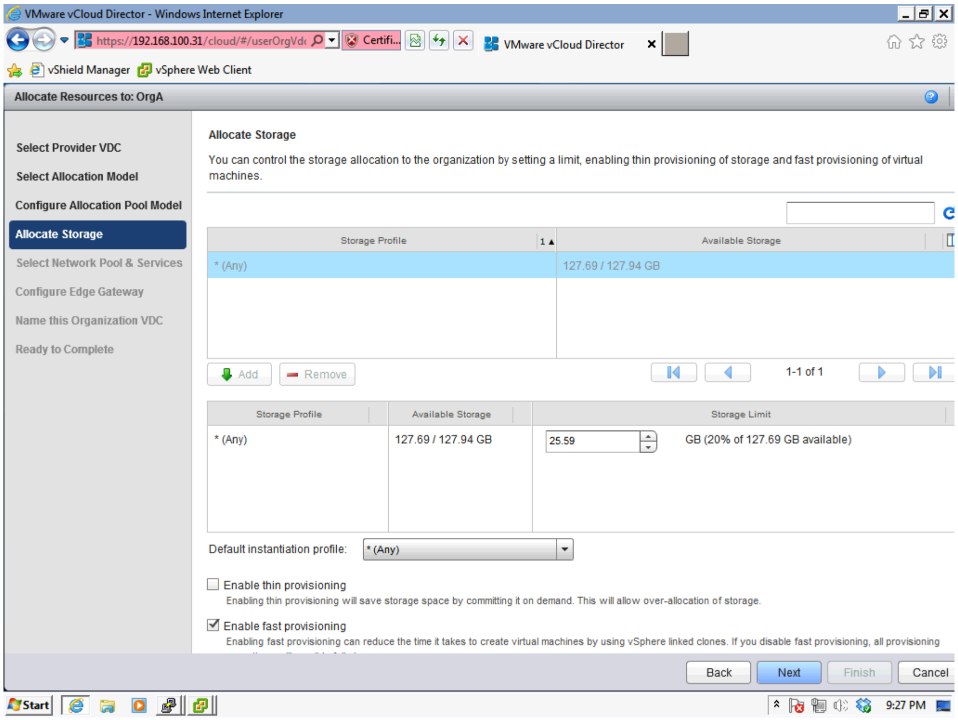

Allocate Resources to: OrgA

Allocate Storage

You can control the storage allocation to the organization by setting a limit, enabling thin provisioning of storage and fast provisioning of virtual machines.

Storage Profile	Available Storage
* (Any)	127.69 / 127.94 GB

Storage Profile	Available Storage	Storage Limit	
* (Any)	127.69 / 127.94 GB	25.59	GB (20% of 127.69 GB available)

Default instantiation profile: * (Any)

☐ Enable thin provisioning
Enabling thin provisioning will save storage space by committing it on demand. This will allow over-allocation of storage.

☑ Enable fast provisioning
Enabling fast provisioning can reduce the time it takes to create virtual machines by using vSphere linked clones. If you disable fast provisioning, all provisioning

Figure 8-34 Storage Profile Selection

After you have created the compute resources for the Org vDC, the next step is to configure the networking options for the Org vDC. These options were covered in Chapter 7. For configuration information on organization virtual datacenters, refer back to that chapter.

Once the networking options have been configured, specify a name for the Org vDC. You have now completed the configuration of an Org vDC. We will now cover how to remove and Org vDC from vCloud Director.

Delete an Organization vDC

Deleting an organization virtual datacenter is similar to the deletion of a provider virtual datacenter. All contents of the Org vDC must be removed before it can be deleted. Contents can include catalog items (media and vApp templates), vApps, networks, and edge gateways.

Here are the steps to check to see if items still exist in an organization virtual datacenter.

ACTIVITY 8-8 PREPARE TO REMOVE AN ORGANIZATION VIRTUAL DATACENTER

1. Login to vCloud Director as a System Administrator

2. Select the organization virtual datacenter that will be removed. The Allocation Org vDC has been selected in Figure 8-35

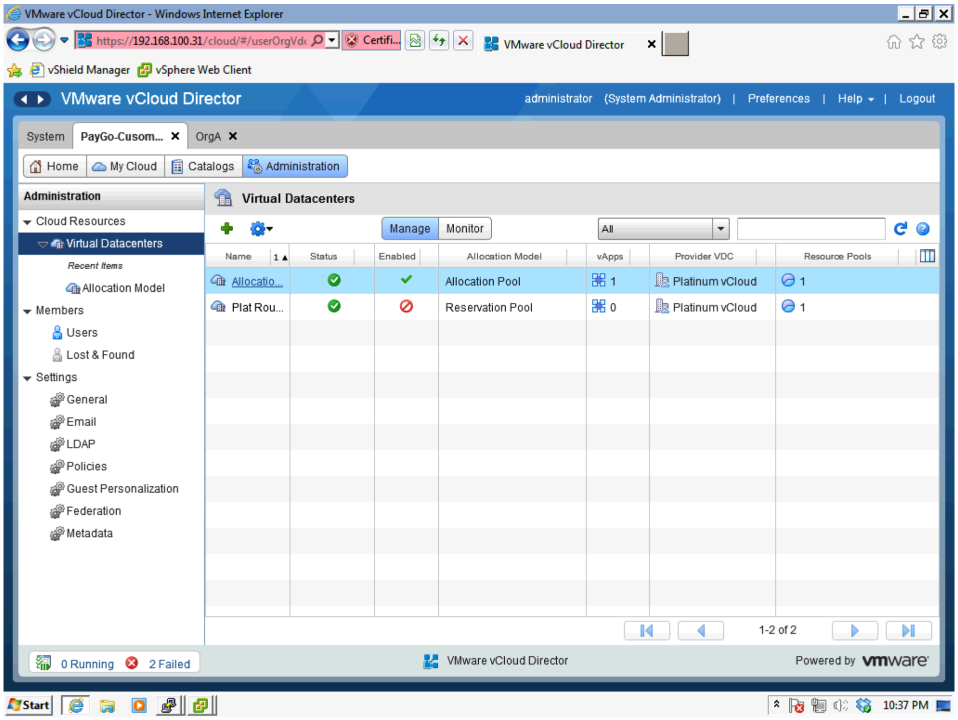

Name	Status	Enabled	Allocation Model	vApps	Provider VDC	Resource Pools
Allocatio...	✓	✓	Allocation Pool	1	Platinum vCloud	1
Plat Rou...	✓	⊘	Reservation Pool	0	Platinum vCloud	1

Figure 8-35 Organization vDC Removal

3. Verify that all vApps have been migrated off or deleted from the organization virtual datacenter. Figure 8-36 shows an empty vApp.

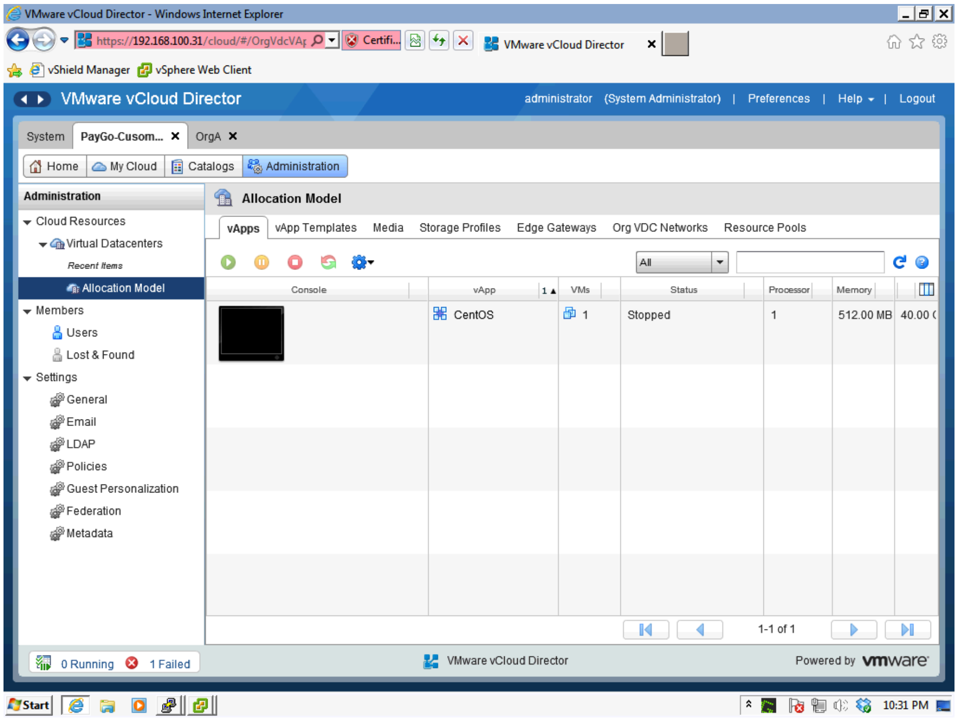

Console	vApp	VMs	Status	Processor	Memory	
	CentOS	1	Stopped	1	512.00 MB	40.00 (

Figure 8-36 Org vDC vApps

4. Verify that all vApp templates, also referred to as Catalog Items, are removed from the organization virtual datacenter. Figure 8-37 shows that all vApp Templates have been removed.

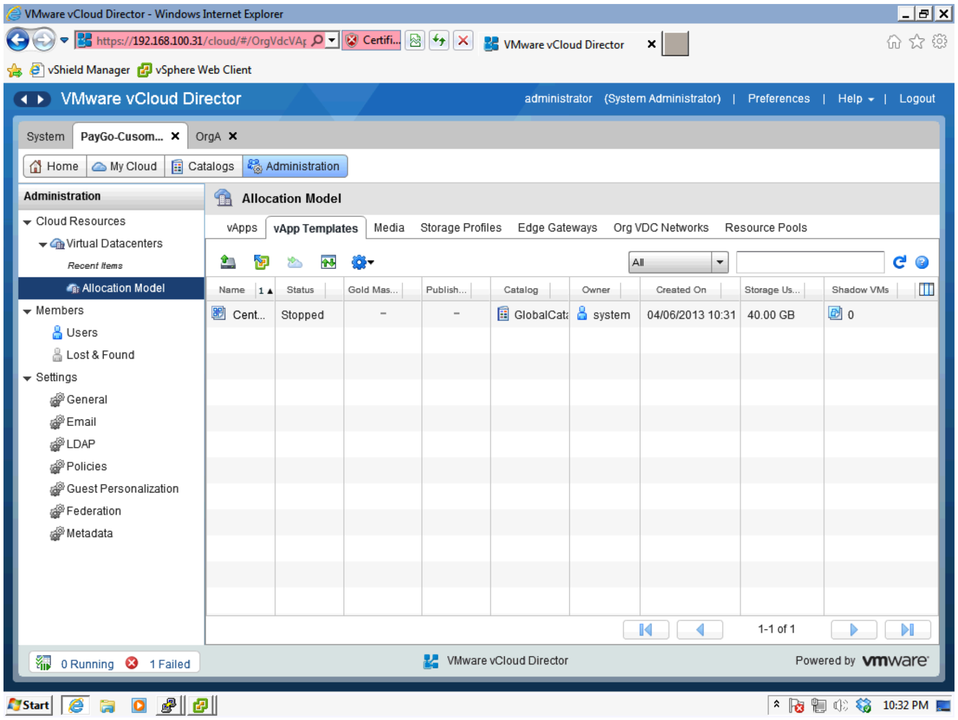

Name	Status	Gold Mas...	Publish...	Catalog	Owner	Created On	Storage Us...	Shadow VMs
Cent...	Stopped	–	–	GlobalCat:	system	04/06/2013 10:31	40.00 GB	0

Figure 8-37 Org vDC vApp Templates

5. Verify that all media (ISO and FLP) files are removed from the organization virtual datacenter

6. Verify that the organization networks have been removed. These must be removed before removing the edge gateway that provides gateway or firewall services for the organization network. Figure 8-38 shows all org networks have been removed.

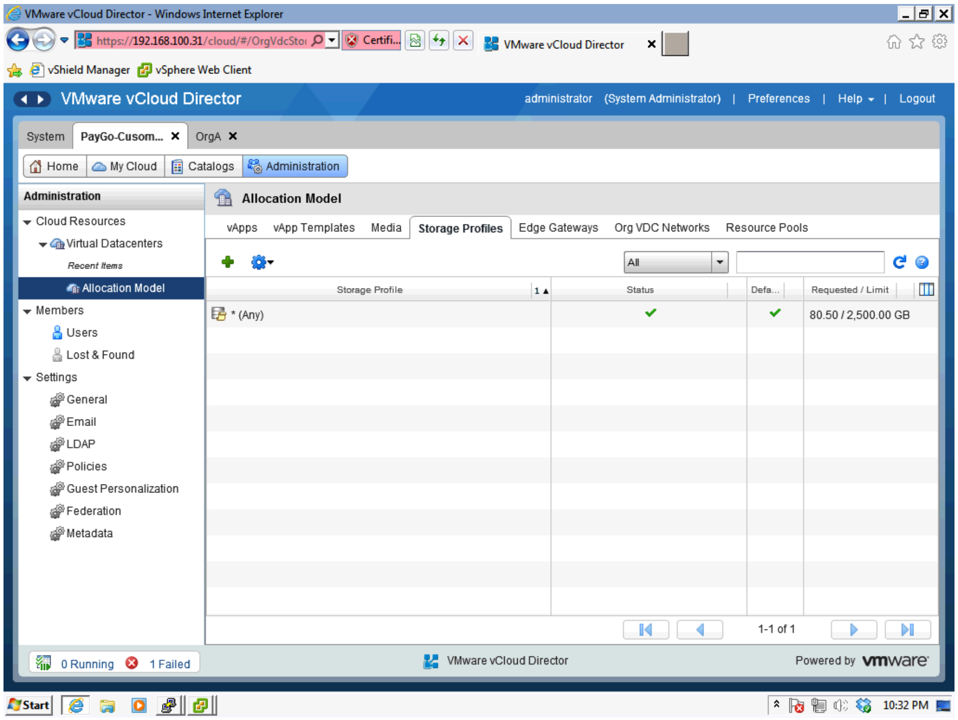

Storage Profile	Status	Defa...	Requested / Limit
* (Any)	✓	✓	80.50 / 2,500.00 GB

Figure 8-38 Org vDC Networks

7. Verify that all edge gateways have been removed from the organization virtual datacenter as shown in Figure 8-39

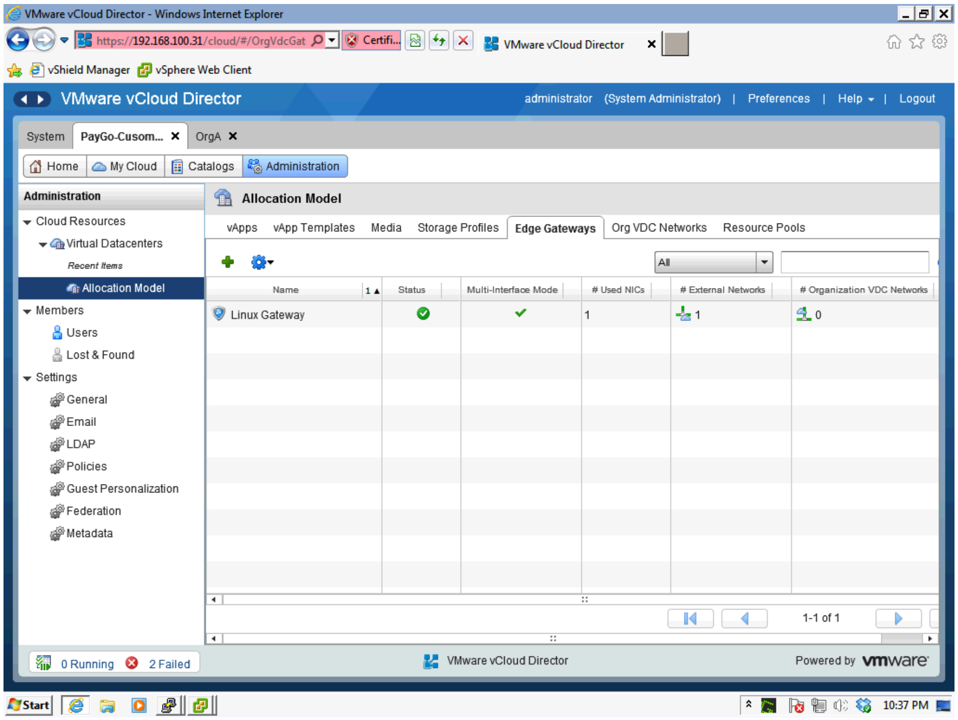

Name	Status	Multi-Interface Mode	# Used NICs	# External Networks	# Organization VDC Networks
Linux Gateway	✓	✓	1	1	0

Figure 8-39 Org vDC Edge Gateways

8. Verify that all storage profiles have been removed from the organization virtual datacenter. vApps, vApp Templates and media need to be removed from the storage profile before it can be removed from the organization virtual datacenter. Figure 8-340 shows that all storage profiles have been removed from the Org vDC.

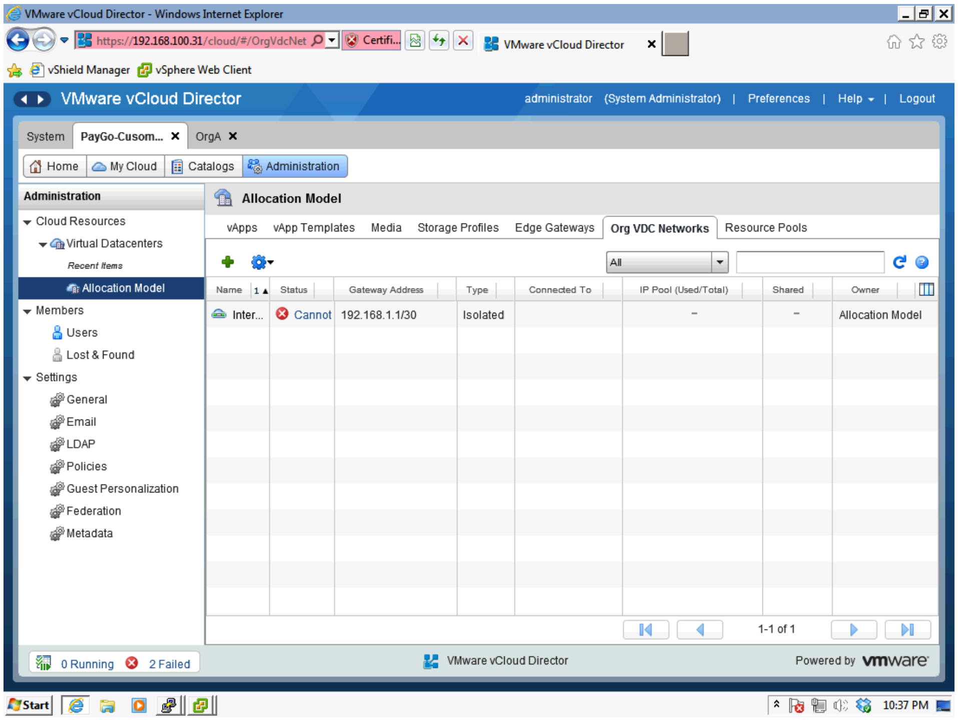

Name	Status	Gateway Address	Type	Connected To	IP Pool (Used/Total)	Shared	Owner
Inter...	Cannot	192.168.1.1/30	Isolated		–	–	Allocation Model

Figure 8-40 Org vDC Storage Profiles

9. Disable the organization virtual datacenter as shown in Figure 8-41

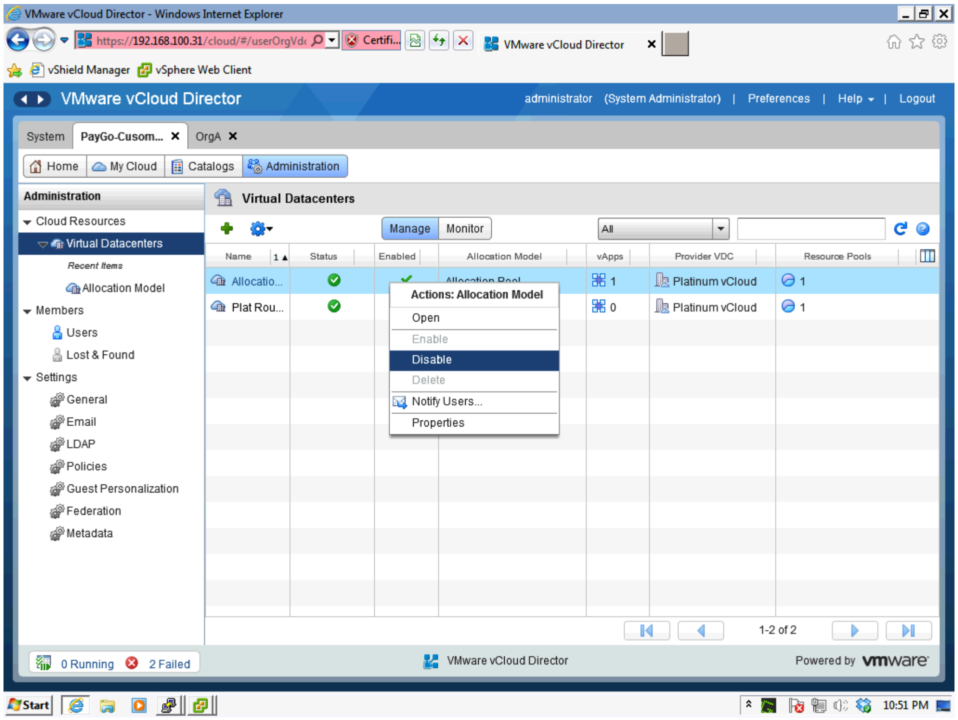

Name	Status	Enabled	Allocation Model	vApps	Provider VDC	Resource Pools
Allocatio...	✓	✓	Allocation Pool	1	Platinum vCloud	1
Plat Rou...	✓			0	Platinum vCloud	1

Figure 8-41 Disabling an Org vDC

NOTE: Disabling the organization virtual datacenter will not power off vApps or delete any items from the datacenter. Disabling the organization virtual datacenter will stop new items from being created and will prevent items from being powered on.

You now should be able to select the delete option and remove the organization virtual datacenter from the organization. If there are any resources remaining, you will be alerted with a message similar to the one shown in Figure 8-42

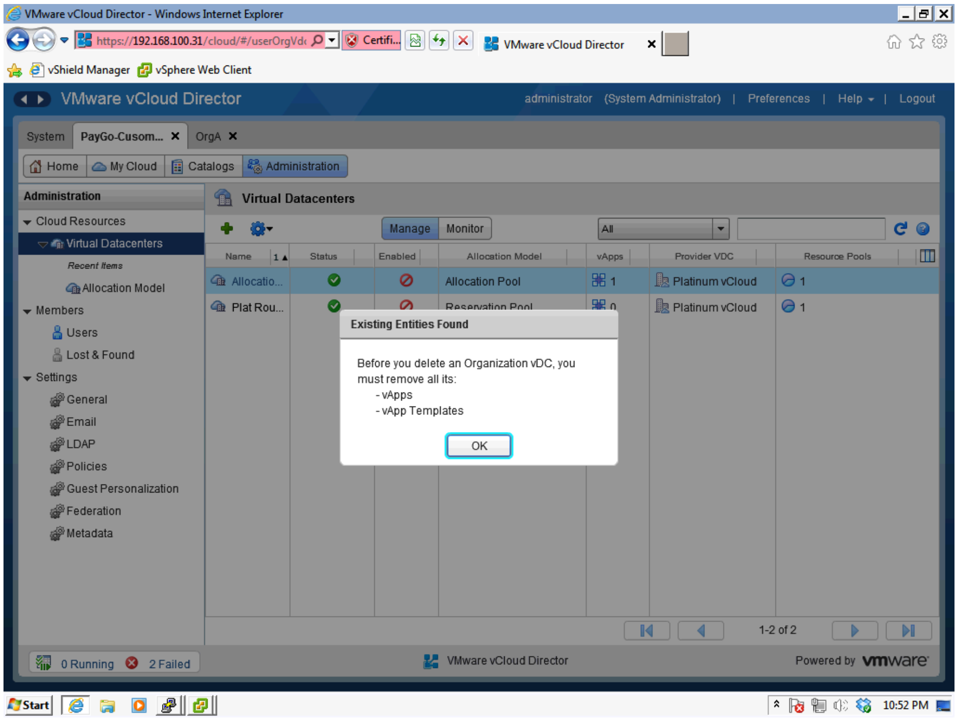

Existing Entities Found

Before you delete an Organization vDC, you must remove all its:
- vApps
- vApp Templates

OK

Figure 8-42 Org vDC Deletion Error

Catalog Management

Catalogs are a part of vCloud Director that make it easy for tenants to create virtual machines. Catalogs can be published and shared across organizations, allowing service providers to provide a base template for tenants or for enterprises to publish pre-built templates of operating systems.

Creating a catalog in vCloud Director is straightforward. The details are in the configuration of the catalog and include publishing the catalog, sharing the catalog, and populating the catalog with vApp Templates and Media. We will cover these steps in the following sections.

##key topic

When placing a vApp into the catalog you are able to define the entry as a ‘Gold Master’. This is used to distinguish a vApp Catalog entry from other items in the catalog. A good example of the use for this is in an environment where Windows servers are in the catalog. As each patch cycle comes along, the most recent version could have the ‘Gold Master’ stamp on it. Another example would be to flag the version of a vApp that is ready for publishing. This flag does not alter the vApp in any way, it is just a visual indicator to assist in highlighting a particular vApp.

##key topic

ACTIVITY 8-9 CREATE A CATALOG

1. Sign in to the organization where you want to create the catalog
2. Select the Catalog tab as shown in Figure 8-43

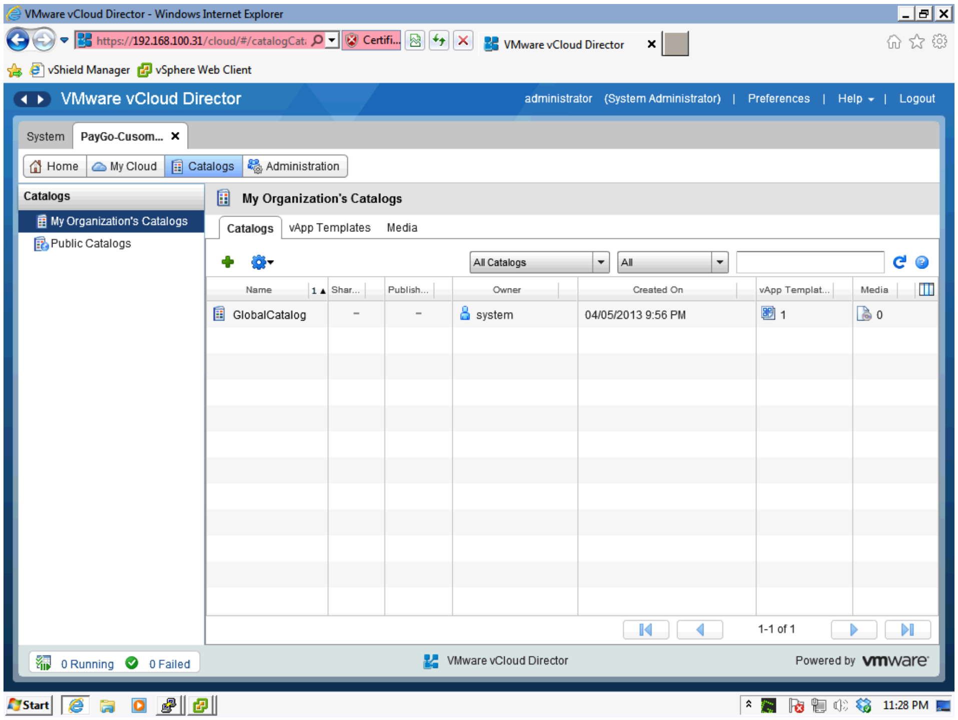

Name	Shar...	Publish...	Owner	Created On	vApp Templat...	Media
GlobalCatalog	–	–	system	04/05/2013 9:56 PM	1	0

Figure 8-43 Organization Catalogs

3. Click the ‘Green Plus Sign’ to create a new catalog shown in Figure 8-44

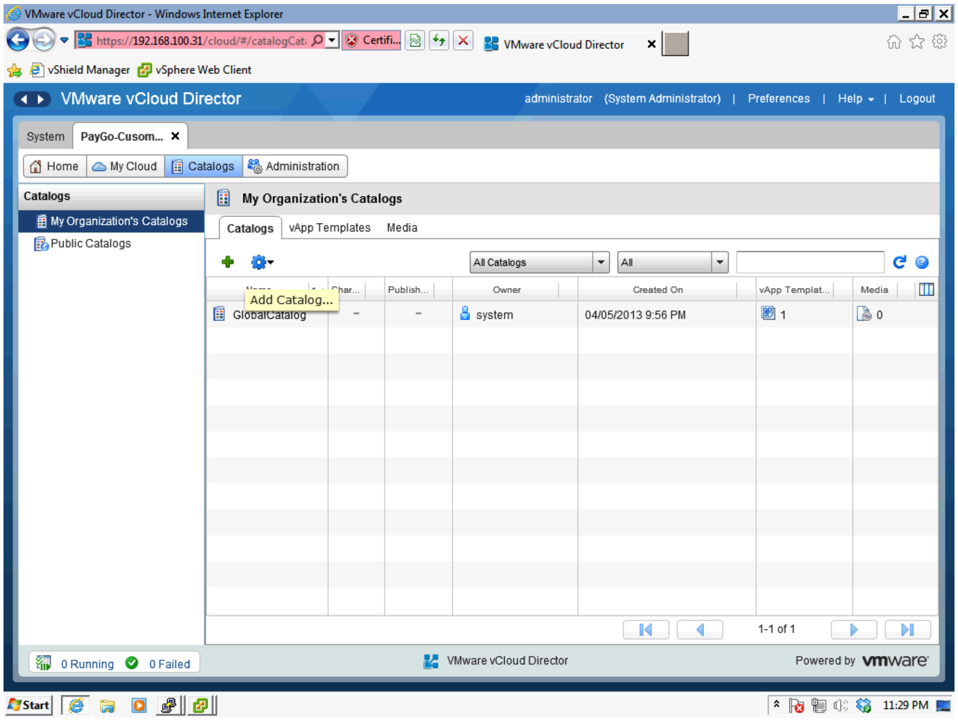

Name	Char...	Publish...	Owner	Created On	vApp Templat...	Media
GlobalCatalog	–	–	system	04/05/2013 9:56 PM	1	0

Figure 8-44 Add a Catalog Selection

4. Type in a name and description for the catalog to be created similar to the name and description shown in 8-45

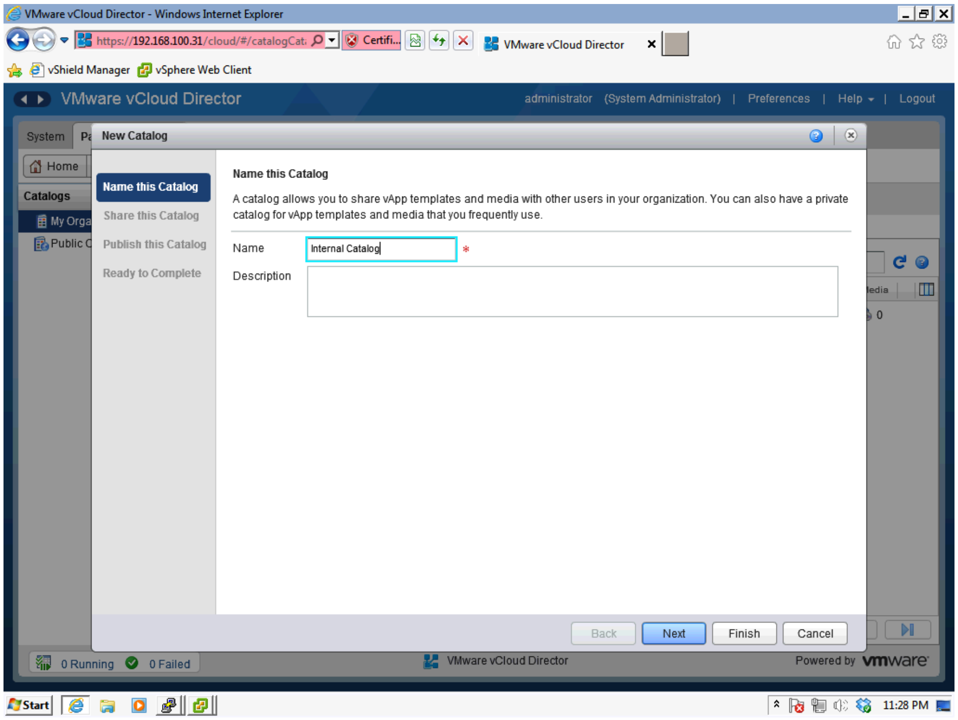

Name this Catalog

A catalog allows you to share vApp templates and media with other users in your organization. You can also have a private catalog for vApp templates and media that you frequently use.

Name: Internal Catalog *

Description:

Figure 8-45 Catalog Name and Description

5. If items in this catalog need to be accessed by more than just the original creator of the catalog, they need to be shared to those users. For simplicity an item can be shared to all users in an organization by selecting the ‘everyone’ option as shown in Figure 8-46 or you can select individual users or groups that have been defined in the organization

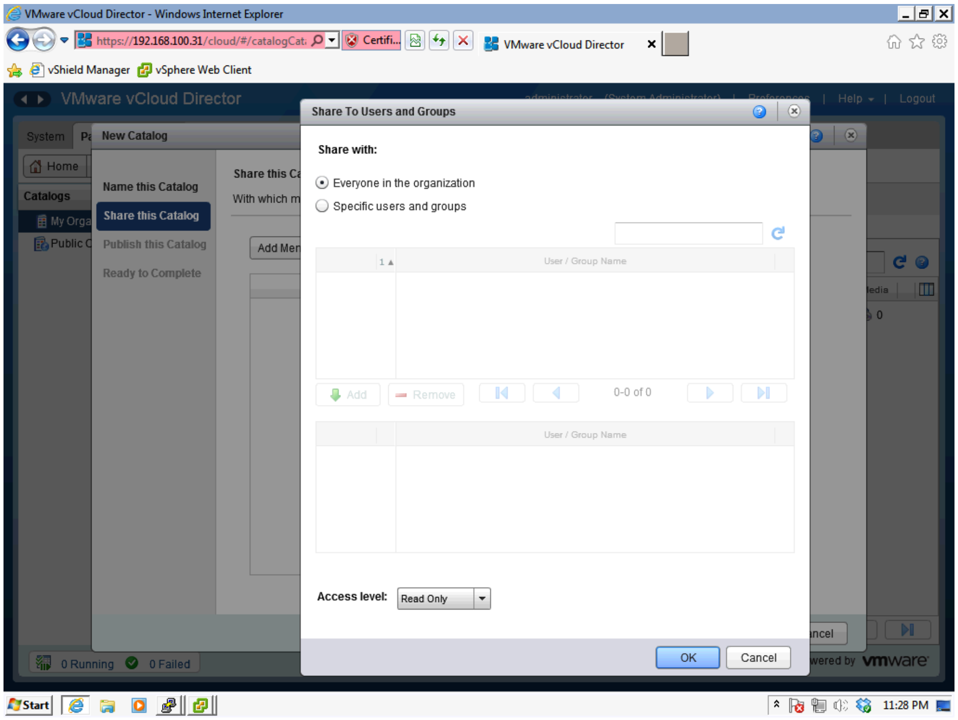

Share with:

● Everyone in the organization

○ Specific users and groups

Access level: Read Only

Figure 8-46 Catalog Sharing

NOTE: It is not possible to share across organizations. If you require the ability for other organizations to access this catalog, it must be published. Once a catalog is published, all organizations will be able to read the catalog.

6. Publishing the Catalog – This screen will appear if the organization where the catalog is being created has permission to publish a catalog. Publishing a catalog allows for all organizations in the vCloud Director instance to access it. The access is read only, meaning other organizations can deploy items from the catalog but they will not be able to add or remove items from the catalog. Figure 8-47 shows a catalog being created that will be shared to other organizations

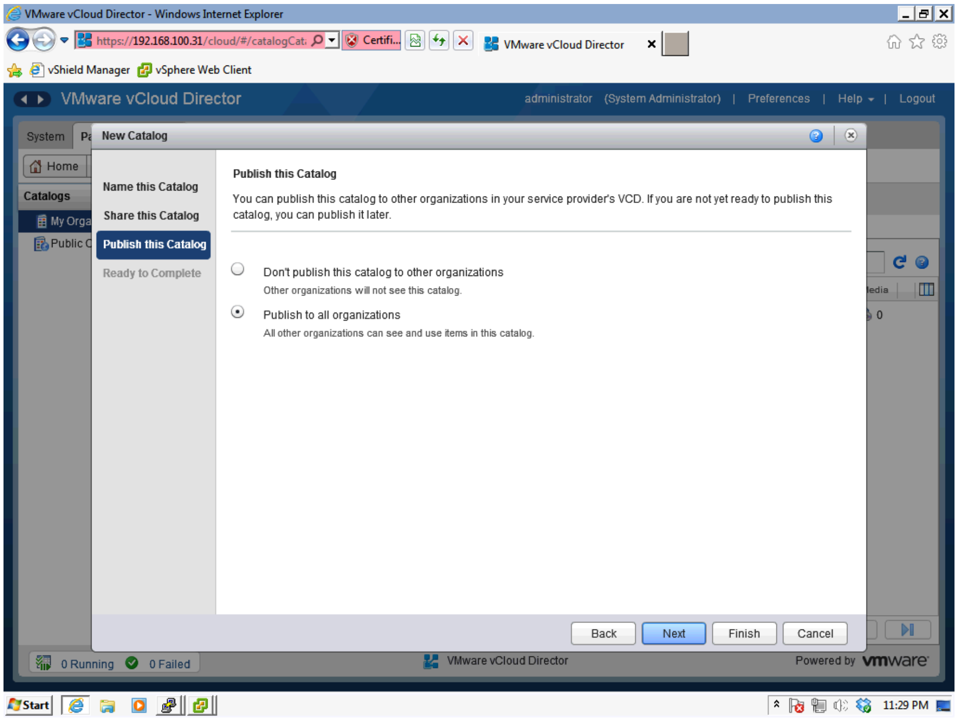

Publish this Catalog

You can publish this catalog to other organizations in your service provider's VCD. If you are not yet ready to publish this catalog, you can publish it later.

○ Don't publish this catalog to other organizations
Other organizations will not see this catalog.

● Publish to all organizations
All other organizations can see and use items in this catalog.

Figure 8-47 Publishing a Catalog

NOTE: It is possible to restrict an individual user from being able to access published catalogs. This is defined in the Roles section of the System administration section. How to create the roles and use them was covered in Chapter 03 - *Administer vCloud Users, Roles and Privileges in a vCloud Director*

Now that you have created a catalog it is time to populate (or add items to) the catalog. These steps are covered in the next section.

Populating Catalogs

vCloud Director provides two ways to populate vApp Templates to the catalog. The first is to import an existing virtual machine from a vCenter Server that is attached to vCloud Director. The other option is to upload a virtual machine(s) in OVF format.

Importing an existing virtual machine from vCenter Server requires the system administrator role. Uploading an OVF can be done by a catalog author. The uploading of an OVF can be time consuming if there are bandwidth constraints, and there is no resume feature for failed uploads.

Media can also be uploaded much in the same way that vApps can be uploaded. If importing media from vCenter Server, the datastore name and the path to the media must be known. There is no browse feature or capability in vCloud Director.

NOTE: Media cannot be shared or published across organizations in vCloud Director 5.1

Chapter Summary

- In this chapter, we discussed how vCloud Director uses a provider virtual datacenter to create a pool of resources for organizations to consume. This pool can be static in size (singe resource pool) or elastic (multiple resource pools).
- We also discussed the models that vCloud Director allows for Provider vDCs to be provisioned to organizations, as well as the removal of these resources from organizations.
- Finally we discussed the use of catalogs and the sharing of catalog entries between users of the same organization as well as sharing a catalog between organizations.

Exam Preparation Tasks

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 8-2 lists a reference of these key topics and the page numbers on which each is found.

Table 8-2 Key Topics

Key Topic Element	Description	Page
Paragraph	What a Provider vDC is	
Activity 8-1	Create a Provider vDC	
Activity 8-2	Preparing a Host	
Activity 8-3	Storage Profile Creation	
Activity 8-4	Disabling a Provider vDC	
Paragraph	Pay Go Allocation Model	
Paragraph	Allocation Model	
Paragraph	Reservation Model	
Activity 8-9	Creating a Catalog	

Definitions of Key Terms

Define the following key terms from this chapter, and check your answers in the glossary.

Provider virtual data center (Provider vDC), Organization virtual data center (Org vDC), Catalog, Storage Profile, OVF, Pay as You Go, Allocation Pool, Reservation Pool

Provider virtual data center (Provider vDC) – Contains the physical resources presented from vCenter Server. It is recommended that the Provider vDC be defined at a top-level resource pool in vCenter Server such as a cluster.

Organization virtual data center (Org vDC) – Contains an organization's defined resources. These can be reserved resources and first come first server resources. The definition of the resources and their allocation is based on the allocation model chosen.

Catalog – An organization's collection of vApp Templates and Media, a catalog is shareable between other organizations by publishing the catalog. Sharing the catalog only applies to inter-organization access to the catalog.

Storage Profile – A storage profile is a collection of one or more vCenter Server datastores with like capabilities that are presented to vCloud Director Org vDC's for consumption.

OVF – The Open Virtualization Format is an open standard for packaging and distributing virtual machines. This can be a single virtual machine or multiple virtual machines. The OVF file itself is an XML file containing the settings and configurations of the VMs associated.

Pay as You Go – The PAYG allocation model allows for individual virtual machine limits and reservations of resources.

Allocation – The Allocation Pool model defines resource settings at the Org vDC or the Resource Pool. This allows for all virtual machines in the Org vDC to consume the pool of resources and not be limited by their individual settings.

Reservation - A Reservation Pool model guarantees a set amount of resources for an Org vDC. These resources are not shared with other tenants.

Review Questions

The answers to these review questions are in Appendix A.

1. The vCPU speed in an allocation model is the limit of the vCPUs in the Org vDC.

a. True

b. False

2. The CPU quota in a Pay as You Go allocation model is elastic and can be exceeded for bursts.

a. True

b. False

3. Which Org vDC model would best be suited for an organization that requires guaranteed performance?

a. Allocation

b. PaYG

c. Reservation

4. ISOs and FLPs can be shared across organizations.

a. True

b. False

5. Which allocation models can be serviced by an elastic Provider vDC?

a. Reservation

b. Allocation

c. PaYG

6. Before removing an Org vDC, which two items must be removed or moved to another Org vDC (Choose two.)?

a. Catalogs

b. vApps

c. vApp Templates

d. Users

7. When creating a Provider vDC it is possible to assign a vCenter Server resource pool to multiple Provider vDCs.

a. True

b. False

8. With an allocation pool Org vDC it is possible to exceed the memory allocation by using ballooning and other vSphere memory management technologies.

   a. True

   b. False

9. Where in an allocation pool model is the memory resource guarantee placed?

   a. Virtual machine

   b. Resource pool

   c. Cluster

   d. None of the above

10. Where in a PaYG model is the CPU limitation placed when virtual machines are deployed?

    a. Virtual machine

    b. Resource pool

    c. Cluster

    d. None of the above

11. When does a Reservation Pool Org vDC reserve the capacity for the Organization?

    a. At creation

    b. When a virtual is deployed

    c. Increased as each VM is deployed

    d. When a virtual machine is powered on

vmware® PRESS

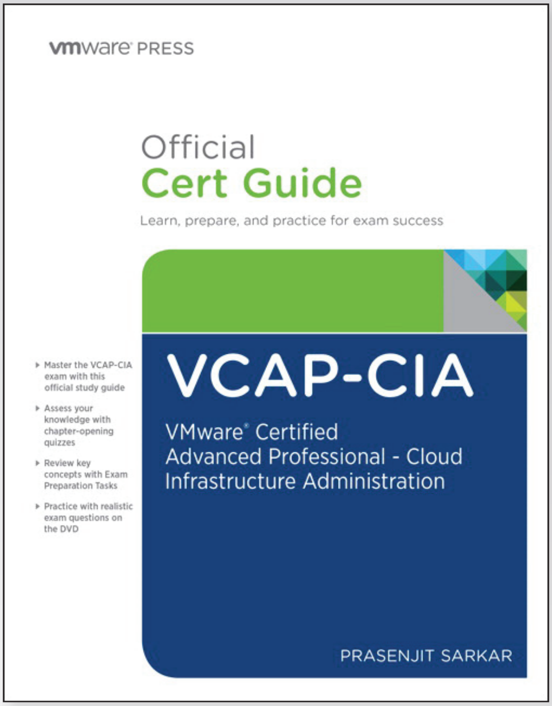

vmware® PRESS

Official

Cert Guide

Learn, prepare, and practice for exam success

- Master the VCAP-CIA exam with this official study guide
- Assess your knowledge with chapter-opening quizzes
- Review key concepts with Exam Preparation Tasks
- Practice with realistic exam questions on the DVD

VCAP-CIA

VMware® Certified Advanced Professional - Cloud Infrastructure Administration

PRASENJIT SARKAR

CHAPTER 2

Manage vSphere Resources

AVAILABLE – FEBRUARY/MARCH 2014

vmwarepress.com

ALWAYS LEARNING

PEARSON

This chapter covers the following subjects:

- Add vSphere compute resources to vCloud Director
- Manage vSphere storage resources
- Manage vSphere network resources

CHAPTER 2

Manage vSphere Resources

VMware vSphere is the foundation layer for VMware vCloud Director (vCD) and provides the compute, storage, and networking resources required for the cloud. Knowing how to manage these vSphere resources from vCD is critical. In this module, you will learn how to manage VMware vSphere resources from the vCD Web console.

Once you have added VMware vCenter servers, you can now take the resources that the vCenter exposes and create cloud constructs using them. VMware vCloud Director (VCD) treats vCenter and VMware vSphere resources as a giant pool of resources. Provider vDCs, organization vDCs, external networks, organization networks, and network pools are all considered cloud resources. After you add cloud resources to vCloud Director, you can modify them and view information about their relationships with each other.

After you add VMware vSphere resources to the VMware vCloud Director system, you can perform some management functions from vCloud Director. You can also use the VMware vSphere Client to manage these resources. vSphere resources include VMware vCenter servers, resource pools, VMware vSphere ESXi hosts, datastores, and network switches and ports.

Effective management of VMware vCloud resources (providers and networks) ensures that customers always have the resources they need while using corporate IT assets. Effective management of vCloud resources also ensures the highest efficiency and cost effectiveness in their use.

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter or simply jump to the “Exam Preparation Tasks” section for review. If you are in doubt, read the entire chapter. Table 1-1 outlines the major headings in this chapter and the corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Chapter Review Questions."

Table 2-1 "Do I Know This Already?" Section-to-Question Mapping

Foundations Topics Section	Questions Covered in This Section
Add vSphere Compute Resources to vCloud Director	1-5
Manage vSphere Storage Resources	6-7
Manage vSphere Network Resources	8-9

1. Can you add vCenter Server to the vCloud Director without vShield Manager?

a. Yes

b. No

c. Maybe

2. Which port number is used for vCenter addition to vCloud Director?

a. 443

b. 902

c. 80

d. 903

3. When you unprepare an ESXi host, which VIB gets removed from the ESXi?

a. VSLAD

b. VTEP

c. vcloud-agent

d. vcloud

4. When you unprepare an ESXi host, will it automatically vMotion those VMs which are running there?

a. Yes

b. No

5. If you have different passwords on ESXi hosts, can you prepare them all at once?

a. Yes

b. No

6. Storage Profiles are based on what type of capabilities?

   a. User-Defined Capability

   b. System-Defined Capability

   c. All of the above

   d. None of the above

7. Storage Profile is introduced in vCloud Director Version ____

   a. 1.0

   b. 1.1

   c. 5.0

   d. 5.1

8. External Networks can be derived from:

   a. vSphere Standard Switch.

   b. vSphere Distributed switch.

   c. Cisco Nexus 1000v Distributed switch.

   d. All of the above.

9. What is the minimum MTU required for VXLAN?

   a. 1500.

   b. 1524.

   c. 1575.

   d. 1600.

Foundation Topics

Add vSphere Compute Resources to vCloud Director

VMware vSphere is the foundation layer for VMware vCloud Director (vCD) and provides the compute, storage, and networking resources required for the cloud. Knowing how to manage these vSphere resources from vCD is critical. In this module, you will learn how to manage VMware vSphere resources from the vCD Web console.

vCloud Director depends on VMware vCenter Server for providing vSphere resources to the Tenants. It also depends on vShield Manager to provide network services to the cloud. Install and configure vShield Manager before you begin installing vCloud Director. You must associate each vCenter Server that you add to vCloud Director with a unique instance of vShield Manager.

Once you register vCloud Director with the vCenter servers that it uses, vCloud Director will now appear as an extension in the vSphere Client Solutions Manager tab. Also the vSphere Client sets the Managed By property for vCloud Director managed virtual machines, which protects those virtual machines added from vCloud Director from being modified using the vSphere Client. If the connection information for vCenter Server changes, or if you want to change how its name or description appears in vCloud Director, you can modify its settings by selecting Properties from the vCenters page.

Here you can also change the connection settings for the vShield Manager, or if you want to use a different VMware vShield Manager. If vCloud Director loses it connection to a vCenter server, or if you change the connection settings, you can try to reconnect.

You can achieve the above task from the Manage & Monitor tab -> vCenters -> by right-clicking the vCenter server name in the vCenters page, which will give you options to manage the registered vCenter server.

Please insert a Key Topic icon here

ACTIVITY 2-1: ADD NEW VCENTER SERVER TO VCLOUD DIRECTOR

Before you start the activity, you should have the vCenter IP Address and admin credential. You should also have the IP Address of vShield Manager IP Address and admin credentials.

1. Open the vCloud Director URL in a supported browser.
2. Login to the Cloud as Administrator. This should have been done as part of initial configuration.
3. You will be presented with the screen shown in Figure 2-1 and from here you can do the initial setup of the vCloud Director.

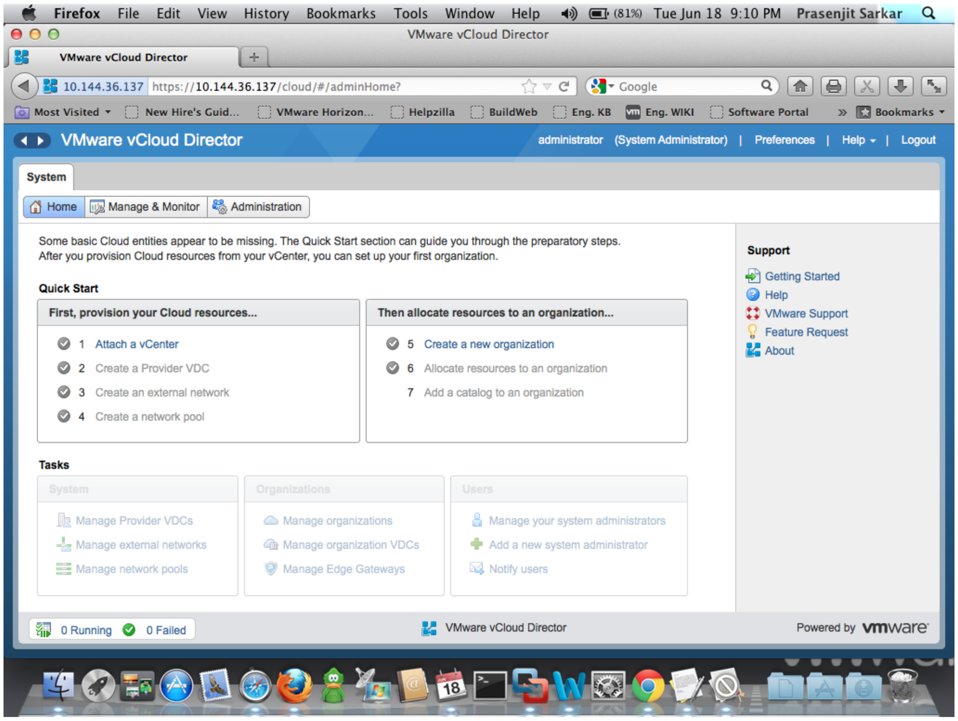
Firefox File Edit View History Bookmarks Tools Window Help
(81%) Tue Jun 18 9:10 PM Prasenjit Sarkar
VMware vCloud Director
10.144.36.137 https://10.144.36.137/cloud/#/adminHome?
Google
Most Visited New Hire's Guid... VMware Horizon... Helpzilla BuildWeb Eng. KB Eng. WIKI Software Portal Bookmarks
VMware vCloud Director
administrator (System Administrator) | Preferences | Help | Logout
System
Home Manage & Monitor Administration
Support
Getting Started
Help
VMware Support
Feature Request
About
0 Running 0 Failed
VMware vCloud Director
Powered by vmware

Some basic Cloud entities appear to be missing. The Quick Start section can guide you through the preparatory steps. After you provision Cloud resources from your vCenter, you can set up your first organization.

Quick Start

First, provision your Cloud resources...

1 Attach a vCenter
2 Create a Provider VDC
3 Create an external network
4 Create a network pool

Then allocate resources to an organization...

5 Create a new organization
6 Allocate resources to an organization
7 Add a catalog to an organization

Tasks

System	Organizations	Users
Manage Provider VDCs	Manage organizations	Manage your system administrators
Manage external networks	Manage organization VDCs	Add a new system administrator
Manage network pools	Manage Edge Gateways	Notify users

Figure 2-1 Initial Setup Screen in VMware vCloud Director

4. As you can see, first step is to attach a **vCenter Server** and the **vShield Manager**.
5. Click on **Attach a vCenter**.
6. You will be presented with the screen shown in Figure 2-2, and where you have to put the vCenter Server information.

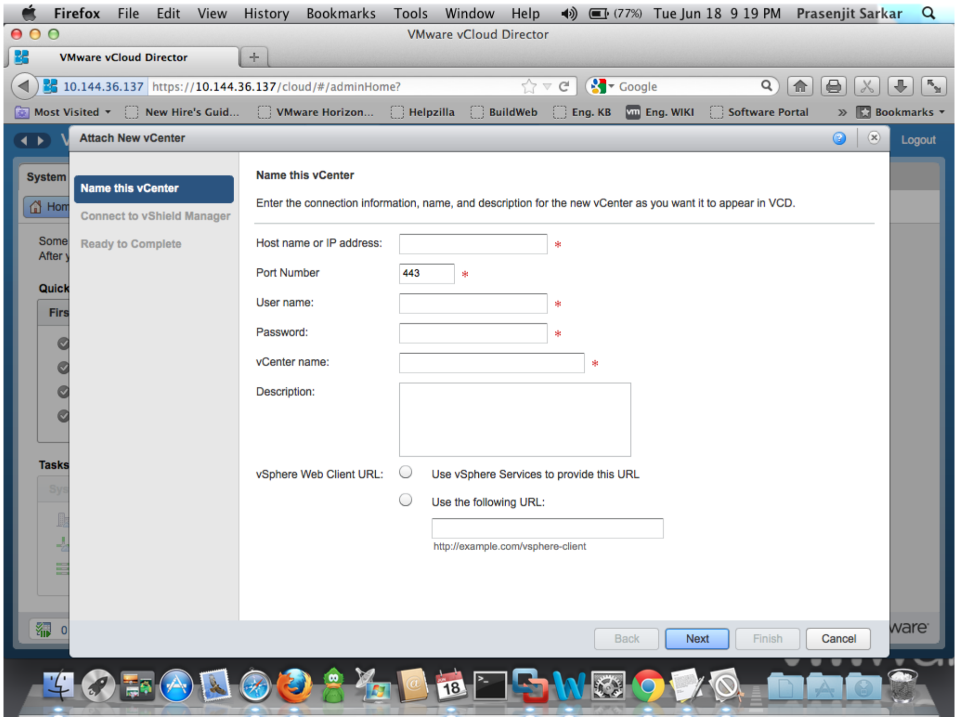
Firefox File Edit View History Bookmarks Tools Window Help
(77%) Tue Jun 18 9 19 PM Prasenjit Sarkar
VMware vCloud Director
https://10.144.36.137/cloud/#/adminHome?
Google
Most Visited New Hire's Guid... VMware Horizon... Helpzilla BuildWeb Eng. KB Eng. WIKI Software Portal Bookmarks
Logout
Name this vCenter
Connect to vShield Manager
Ready to Complete
Back
Next
Finish
Cancel

Attach New vCenter

Name this vCenter

Enter the connection information, name, and description for the new vCenter as you want it to appear in VCD.

Host name or IP address: *

Port Number 443 *

User name: *

Password: *

vCenter name: *

Description:

vSphere Web Client URL: ○ Use vSphere Services to provide this URL

○ Use the following URL:

http://example.com/vsphere-client

Figure 2-2 vCenter Connection Information in vCloud Director

7. Specify the vCenter connection information and click on **Next**.
8. You will be presented with a screen like that shown in Figure 2-3, and where you have to put the vShield Manager information.

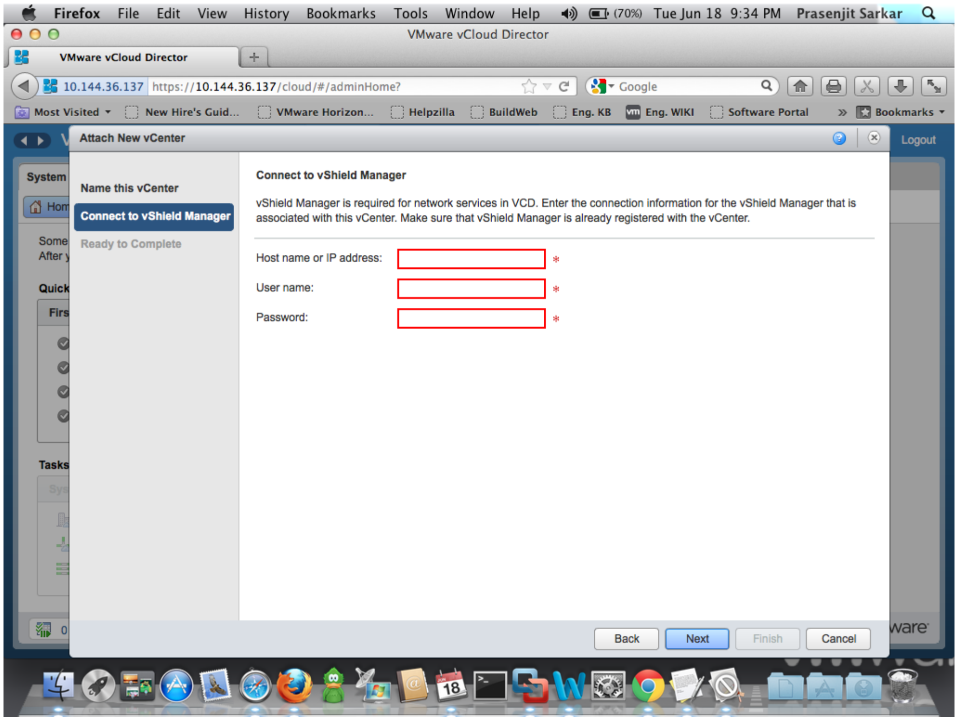

Attach New vCenter

Connect to vShield Manager

vShield Manager is required for network services in VCD. Enter the connection information for the vShield Manager that is associated with this vCenter. Make sure that vShield Manager is already registered with the vCenter.

Host name or IP address: *

User name: *

Password: *

Figure 2-3 vShield Manager Connection Information in vCloud Director

9. Specify the vShield Manager Server connection information and click on **Next**.
10. On the final screen, click on **Finish**.
11. Once you add the vCenter Server, you can see it under **Manage and Monitor** tab.
12. Go to the **Manage and Monitor** tab and under **vSphere Resources** section click on vCenters. You will see a similar screen to that shown in Figure 2-4:

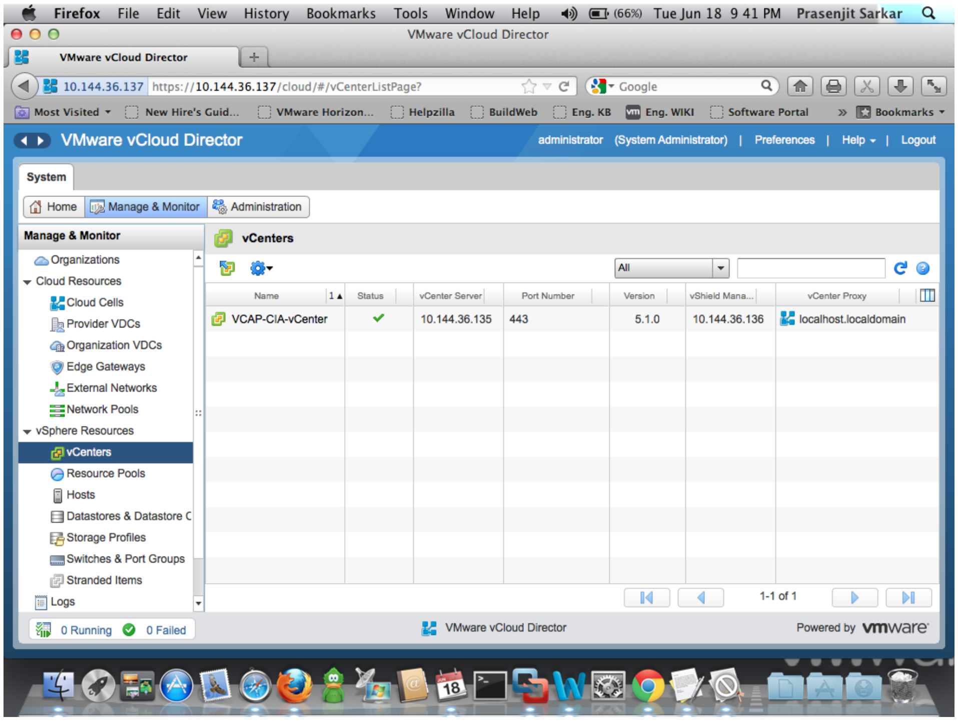

Name	Status	vCenter Server	Port Number	Version	vShield Mana...	vCenter Proxy
VCAP-CIA-vCenter	✓	10.144.36.135	443	5.1.0	10.144.36.136	localhost.localdomain

Figure 2-4 vCenter Information in vCloud Director

NOTE As a prerequisite, vCenter Server has to be registered with your vShield Manager. If not you will see an error "vShield Manager is not registered with the VC <VC Name>. Perform VC registration in vShield Manager and retry." If you get this error then follow the below steps to register your vCenter Server with vShield Manager.

1. Open the vShield Manager URL in a supported browser.
2. Login to the Cloud as **Administrator**. This should have been done as part of initial configuration.
3. In the main **Settings and Reports** section, find the **vCenter Server** section and you will see there is no vCenter Server registered with vShield Manager.
4. Click on the **Edit** button.
5. Specify the vCenter Server information and credential.
6. Click on **OK**.
7. Click **Yes** on the security warning.
8. vCenter Server should now be configured.

Immediate next task after you add vCenter Server to the vCloud Director is creating a Provider VDC. We will talk about this later in the other chapter. However in this chapter we will add an ESXi Server in the vCenter and then prepare that ESXi hosts.

When you add an ESX/ESXi host to a vSphere cluster that vCloud Director uses, you must prepare the host before a provider vDC can use its resources. However, you cannot prepare a host that is in lockdown mode. After you prepare a host, you can enable lockdown mode.

After this we will also see how to disable and unprepared ESXi hosts and Use Cases of this.

You can disable a host to prevent vApps from starting up on the host. Virtual machines that are already running on the host are not affected. Also disable the host to perform maintenance. Note that vCloud Director enables or disables the host for all provider vDCs that use its resources.

After disabling an ESXi host, you can unprepare that ESXi host.

insert key topic icon

ACTIVITY 2-2: UNPREPARE ESXI HOSTS IN VCLOUD DIRECTOR

To unprepare ESXi hosts in vCloud Director, follow the steps:

1. Start the Internet Explorer browser. Go to the URL of the vCD server

   An example would be https://serverFQDN/cloud
2. Log in to vCD by typing an administrator user ID and password
3. Click the **Manage and Monitor** tab.
4. Click **Hosts** in the left panel.
5. Right Click on the ESXi host and select **Disable Host** as shown in Figure 2-5.

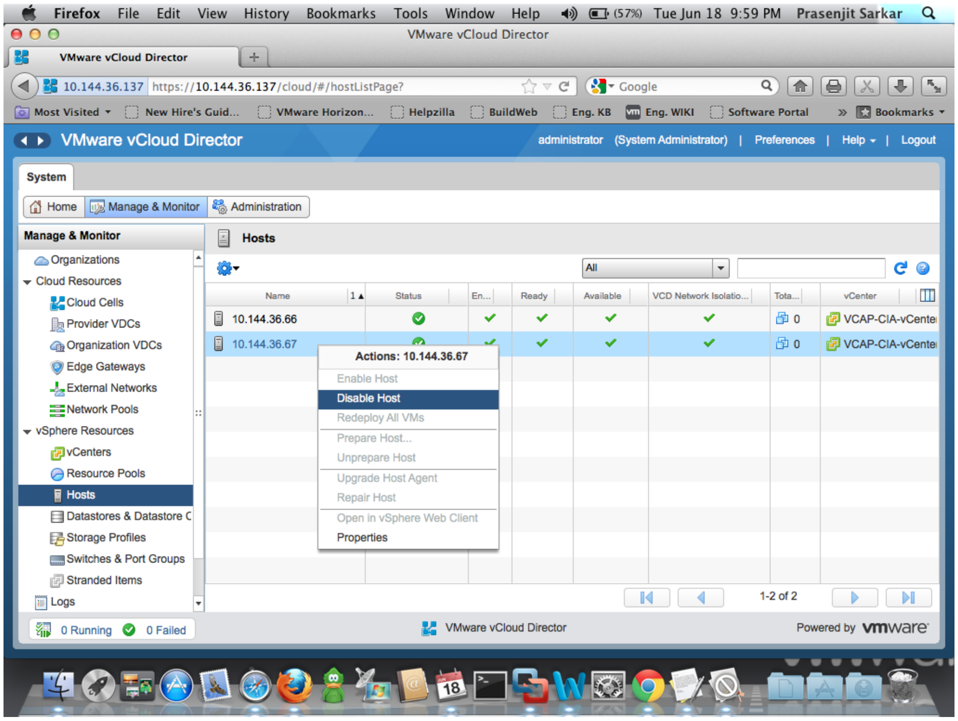

Name	Status	En...	Ready	Available	VCD Network Isolatio...	Tota...	vCenter
10.144.36.66	✓	✓	✓	✓	✓	0	VCAP-CIA-vCente
10.144.36.67	✓	✓	✓	✓	✓	0	VCAP-CIA-vCente

Figure 2-5 Disabling ESXi Host in vCloud Director

6. After this ESXi host is disabled then right click on this host and select **Unprepare Host** as shown is Figure 2-6.

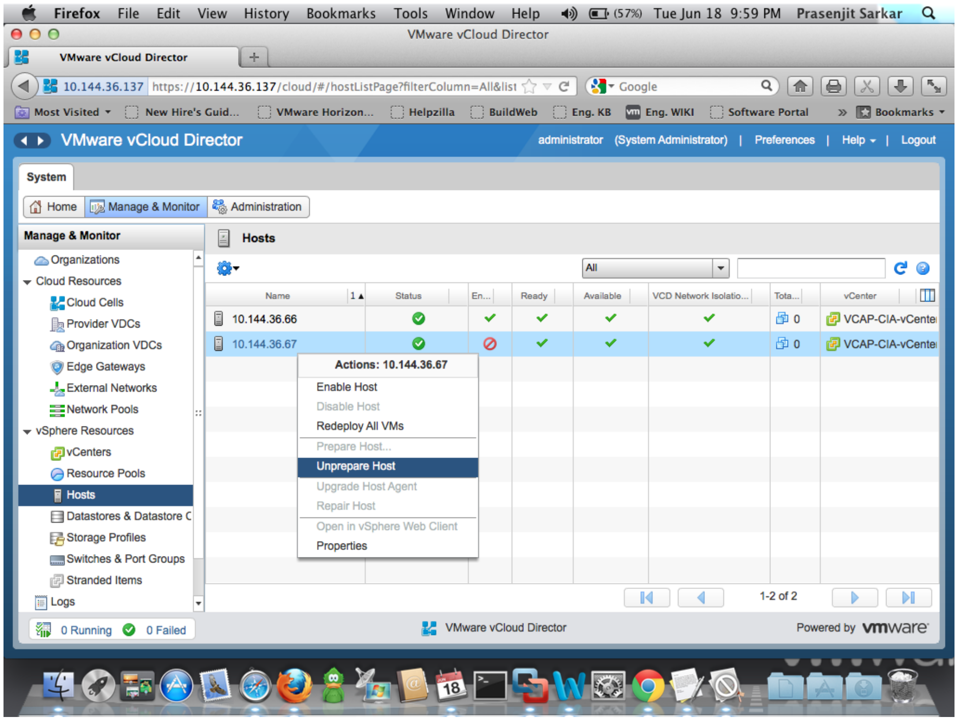

Name	Status	En...	Ready	Available	VCD Network Isolatio...	Tota...	vCenter
10.144.36.66	✓	✓	✓	✓	✓	0	VCAP-CIA-vCente
10.144.36.67	✓	⊘	✓	✓	✓	0	VCAP-CIA-vCente

Figure 2-6 Unprepare ESXi Host in vCloud Director

7. You will get a warning; select **Yes** as shown is Figure 2-7.

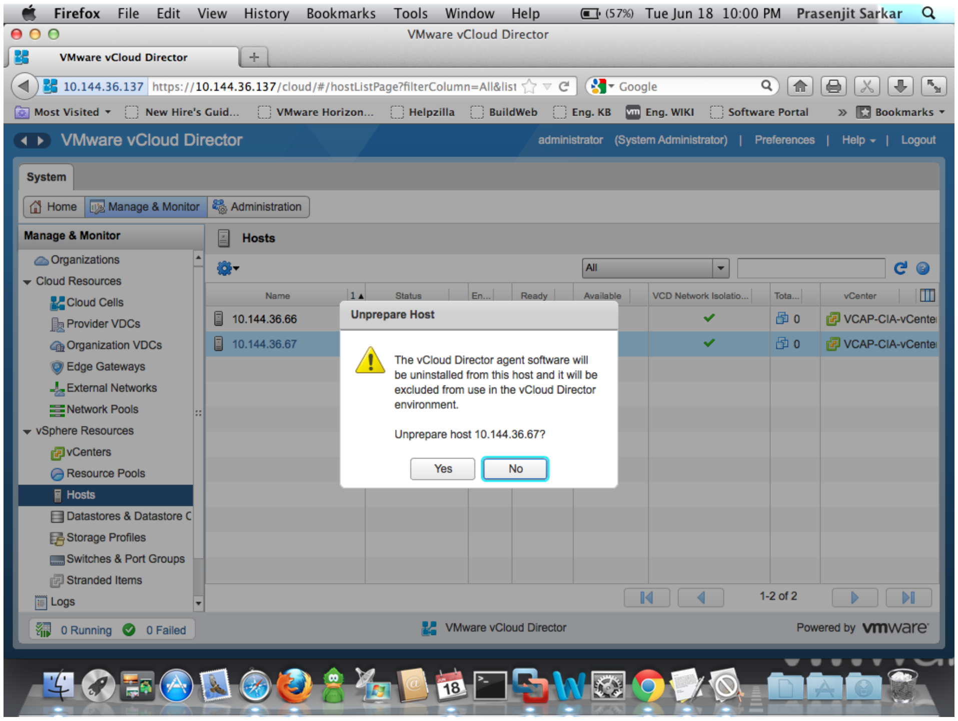

Unprepare Host

The vCloud Director agent software will be uninstalled from this host and it will be excluded from use in the vCloud Director environment.

Unprepare host 10.144.36.67?

Yes No

Figure 2-7 Warning while unpreparing ESXi Host in vCloud Director

Unprepare task will put the ESXi host in Maintenance Mode, then the vCloud Agent will be uninstalled from the ESXi host and then ESXi host will exit from the maintenance mode.

8. You can also use command to manually unprepare an ESXi host. Command is as follows:

```
~ # esxcli software vib remove –n vcloud-agent
```

If you want to increase CPU and memory resources for your pVDC, the easy way to do is to add ESXi hosts into the cluster which is backing this pVDC. You must prepare your ESXi host in vCloud Director after you add it to the vCenter and then only you can use its resources.

You should also keep in mind that you cannot prepare a host which is in Lockdown mode. However after you prepare the host; you can enable the host lockdown.

You may wonder that what is the essence of preparing a host, when you prepare an ESXi host, essentially it will install an agent on the ESXi host for vCloud Director so that host can work with vCloud.

insert key topic icon

ACTIVITY 2-3: ADD ESXI HOSTS TO VCENTER AND PREPARE ESXI HOST IN VCLOUD DIRECTOR

To add ESXi host and prepare it in vCloud Director, follow the steps.

1. Open up the vSphere Client and login to vCenter Server.
2. On the Home screen select **Hosts and Clusters**.
3. Right click on the ESXi Hosts Cluster and select **Add Host**.
4. Specify the ESXi host connection information and add the ESXi host.
5. Once ESXi host is added then start the Internet Explorer browser. Go to the URL of the vCD server.

   An example would be https://serverFQDN/cloud
6. Log in to vCD by typing an administrator user ID and password.
7. Click the **Manage and Monitor** tab.
8. Click **Hosts** in the left panel. You will see a similar screen as in Figure 2-8.

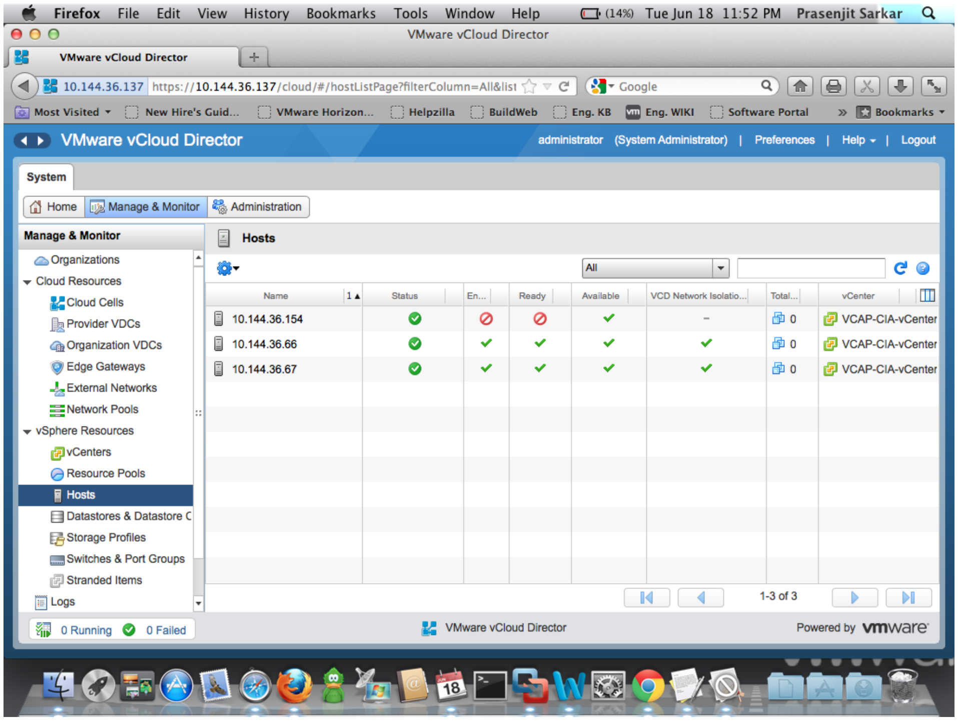

Name	Status	En...	Ready	Available	VCD Network Isolatio...	Total...	vCenter
10.144.36.154	✓	⊘	⊘	✓	–	0	VCAP-CIA-vCenter
10.144.36.66	✓	✓	✓	✓	✓	0	VCAP-CIA-vCenter
10.144.36.67	✓	✓	✓	✓	✓	0	VCAP-CIA-vCenter

Figure 2-8 ESXi Hosts in vCloud Director

9. Right Click on the newly added ESXi host and select **Prepare Host** as shown in Figure 2-9.

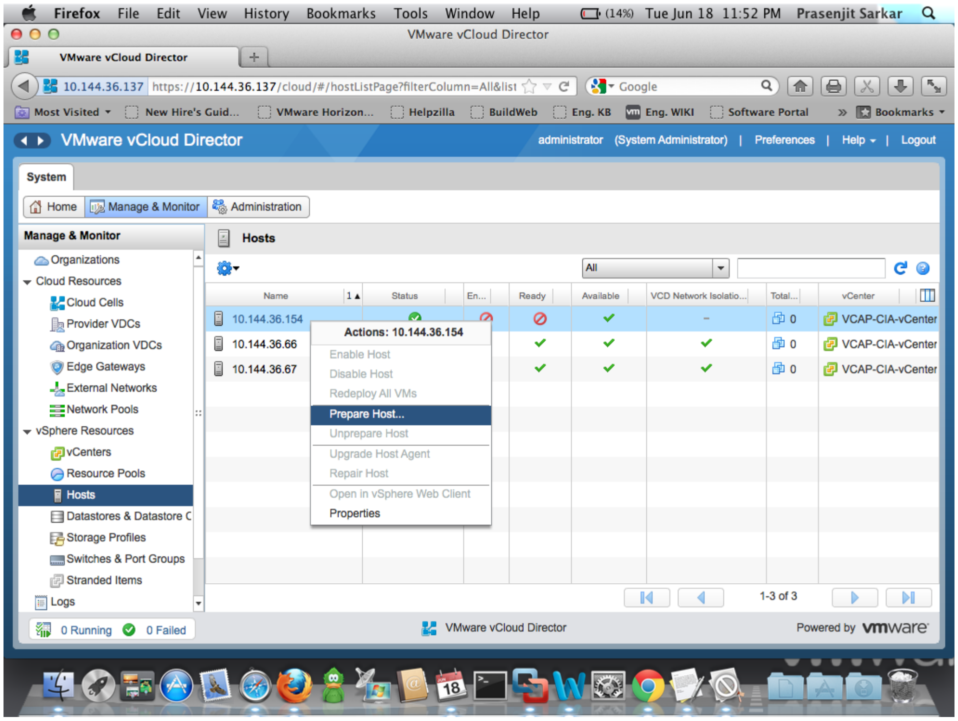

Name	Status	En...	Ready	Available	VCD Network Isolatio...	Total...	vCenter
10.144.36.154	✓	⊘	⊘	✓	-	0	VCAP-CIA-vCenter
10.144.36.66			✓	✓	✓	0	VCAP-CIA-vCenter
10.144.36.67			✓	✓	✓	0	VCAP-CIA-vCenter

Figure 2-9 Prepare ESXi Host in vCloud Director

10. It will ask for ESXi host credential. If you have global credential for all of your ESXi host then specify a single credential for all hosts. Otherwise specify credential for each host as shown in Figure 2-10.

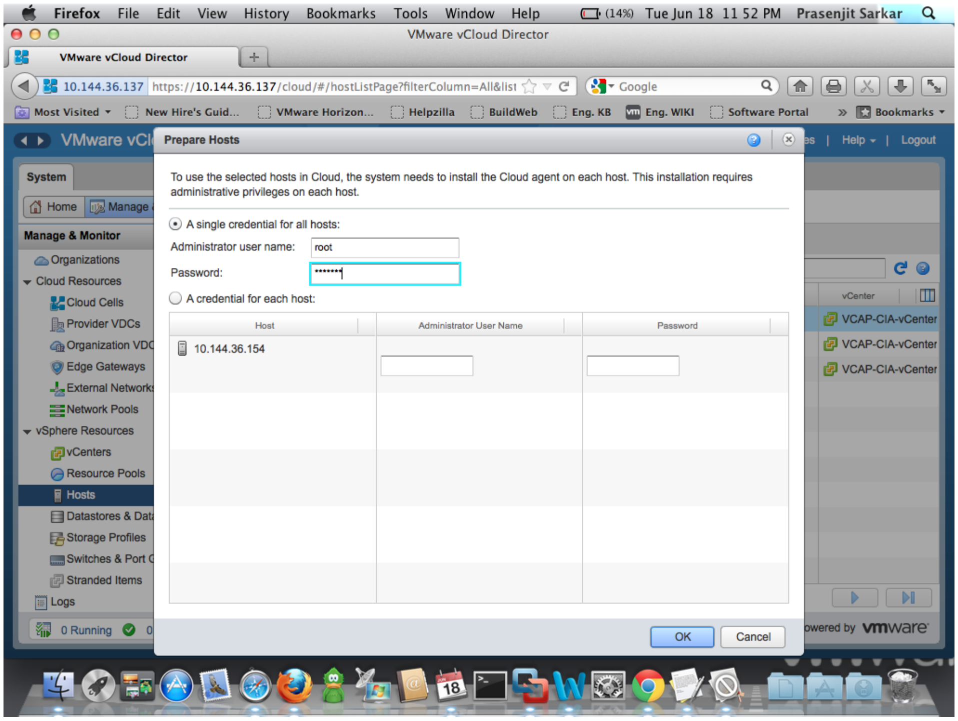

Prepare Hosts

To use the selected hosts in Cloud, the system needs to install the Cloud agent on each host. This installation requires administrative privileges on each host.

● A single credential for all hosts:

Administrator user name: root

Password: •••••••

○ A credential for each host:

Host	Administrator User Name	Password
10.144.36.154		

Figure 2-10 User Credential while preparing ESXi Host in vCloud Director

11. Once it prepares the ESXi hosts, you will see it **Enabled** and **Ready**.

After you add VMware vSphere resources to the VMware vCloud Director (vCD) system, you can perform some management functions from vCloud Director. You can also use the VMware vSphere Client to manage these resources. vSphere resources include VMware vCenter servers, resource pools, VMware ESX/VMware vSphere ESXi hosts, datastores, and network switches and ports.

You must create and configure resource pools/vSphere clusters in vSphere before you can add them to the provider vDC as every provider vDC in a vCloud Director installation requires a unique resource pool in vSphere to provide its compute and memory resources. You have to create and configure resource pools in vSphere before you can add them to a provider vDC. However you have the ability to view information about the resource pools that pVDC uses from vCloud Director.

You can view information about the used and total CPU and memory reservations for a resource pool. You can also view information about the datastores that are available to the resource pool.

Best Practice is to dedicate an entire cluster to a provider vDC. However, it is also possible to have multiple resource pools on a single cluster, with each resource pool being assigned to a provider vDC.

The type of settings used on the resource pool (reservations and limits) should be consistent with the allocation model that will be used in the organization vDC that leverages this resource pool.

insert key topic icon

ACTIVITY 2-4: MANAGE AND MONITOR VSPHERE RESOURCE POOLS

To manage and monitor VMware vSphere resource pool in vCloud Director, follow the steps:

1. Start the Internet Explorer browser. Go to the URL of the vCD server. An example would be https://serverFQDN/cloud.
2. Log in to vCD by typing an administrator user ID and password.
3. Click on the **Manage and Monitor** tab.
4. Click on the **Resource Pools** in the left panel.
5. Now you will see the resource pool there as shown in Figure 2-11.

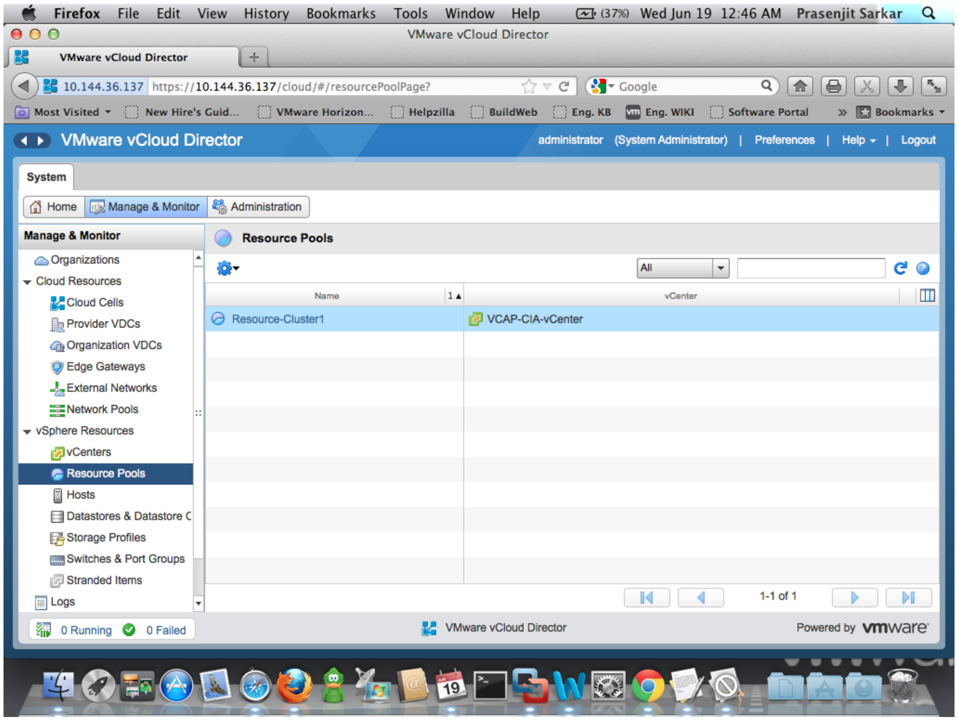

Name	vCenter
Resource-Cluster1	VCAP-CIA-vCenter

Figure 2-11 Resource Pool in vCloud Director

6. Right click on the Resource Pool there and select **Properties**. In this screen, you will see the properties of resource pool where you can see the CPU and Memory Reservations and Datastore name and it utilization as shown in Figure 2-12.

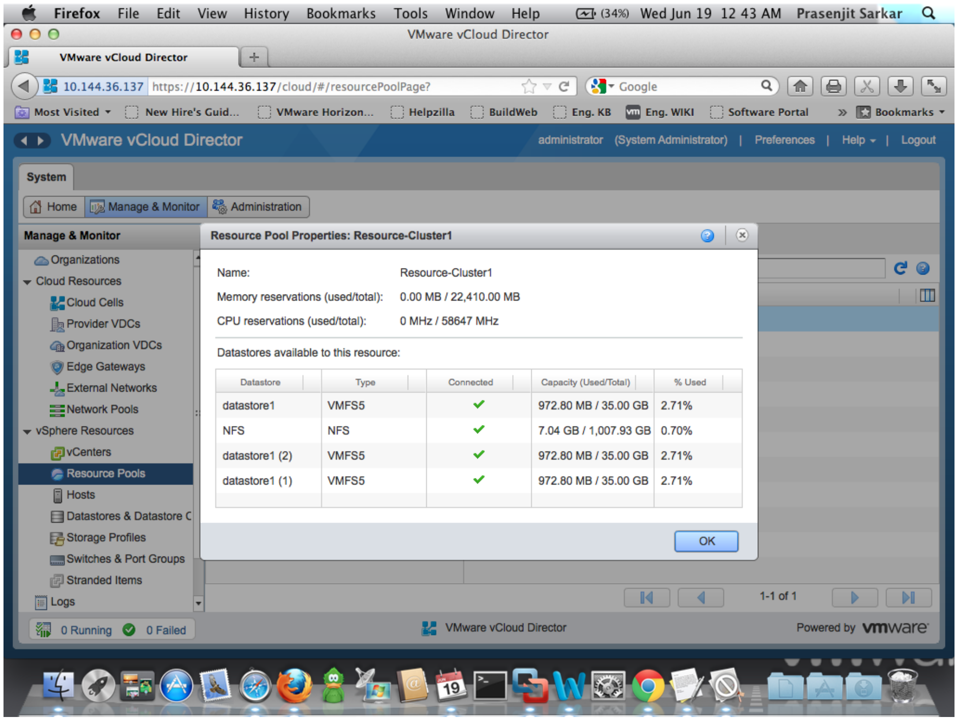

Resource Pool Properties: Resource-Cluster1

Name: Resource-Cluster1

Memory reservations (used/total): 0.00 MB / 22,410.00 MB

CPU reservations (used/total): 0 MHz / 58647 MHz

Datastores available to this resource:

Datastore	Type	Connected	Capacity (Used/Total)	% Used
datastore1	VMFS5	✓	972.80 MB / 35.00 GB	2.71%
NFS	NFS	✓	7.04 GB / 1,007.93 GB	0.70%
datastore1 (2)	VMFS5	✓	972.80 MB / 35.00 GB	2.71%
datastore1 (1)	VMFS5	✓	972.80 MB / 35.00 GB	2.71%

Figure 2-12 Resource Pool properties in vCloud Director

insert key topic icon

ACTIVITY 2-5: MANAGE VMWARE ESXI HOSTS IN VCLOUD DIRECTOR

To manage VMware ESXi hosts in vCloud Director, follow the steps:

1. Start the Internet Explorer browser. Go to the URL of the vCD server.

   An example would be https://serverFQDN/cloud.
2. Log in to vCD by typing an administrator user ID and password.
3. Click on the **Manage and Monitor** tab.
4. Click on the **Hosts**.
5. Select an ESXi host and right click and select **Disable Host**.
6. Again right click on the same host and select **Redeploy VMs** and click on **Yes**.

   The VMs which were registered with this ESXi host will be migrated to the other ESXi hosts in the cluster through vMotion.
7. Once this task is done, that is there is no VMs on this ESXi host, right click on this ESXi host and click on **Enable Host**.

Manage vSphere Storage Resources

With the release of vCloud Director 5.1, we have seen lots of good features and capabilities being added to it. One of such a thing is Storage Profiles.

vCloud Director now leverages vSphere storage profiles capability for providing profile/class driven storage to vCloud customers. Without even creating any storage profile, VMware provides a generic default storage profile and that is *(Any) storage profile. This profile includes all of the datastores from all of your ESXi hosts in your vSphere cluster. That means your local datastores will also be added to it.

This is not something you want to have it right? Yes I know; you never want to keep your Cloud VMs on an ESXi local datastore. So essentially what you want to do is to create a storage profile which will access only the specified datastores (Shared) and using that create your pVDC and assign class driven storage profiles (e.g. Gold, Silver, Bronze).

vSphere Storage Profile are based on either VASA capabilities or user-defined storage capabilities. When you create a provider vDC you have to assign at least one vSphere Storage Profile to the provider vDC. You can also assign storage from multiple vSphere Storage Profile instances to a single provider vDC and that may depend on different class of storage (SSD, SAS, and SATA). Organization vDCs gets their storage from a single provider vDC. If the provider vDC has access to storage from multiple vSphere Storage Profile instances, storage from those same multiple vSphere Storage Profile instances is also available to the organization vDC.

vSphere Storage Profile instances must still be defined on the resource cluster. That means you still need to enable your resource cluster to use Storage Profile capability.

NOTE If you are upgrading your vCloud from 1.5 to 5.1 then there are some major caveats around storage profiles and that has to be rectified almost immediately before you start provisioning your cloud VMs again. For more information on this follow this article:

http://stretch-cloud.info/2013/01/upgrading-your-vcloud-from-1-5-to-5-1-watch-out-for-the-any-storage-profile-caveat/.

insert a Key Topic icon

ACTIVITY 2-6: CONFIGURE STORAGE PROFILES IN VCENTER

Let’s first define some Storage Capabilities for our Datastores that we will use in Storage Profiles.

1. Login to the vSphere Web Client as root

2. You will be landed on the **Home** screen as shown in Figure 2-13

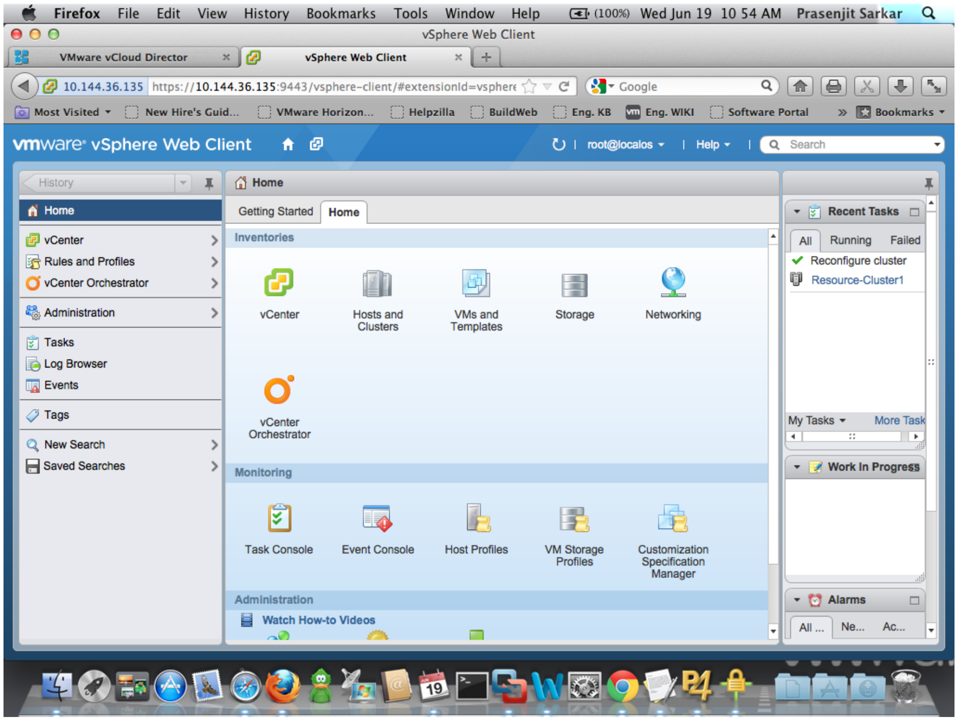

Figure 2-13 vCenter Home Screen

3. Select the **Storage** link from the Home screen.

4. Expand the Datastores on the left side pane of the screen.

5. Select the shared datastore and click on **Assign Storage Capability** on the right hand side of the screen as shown in Figure 2-14.

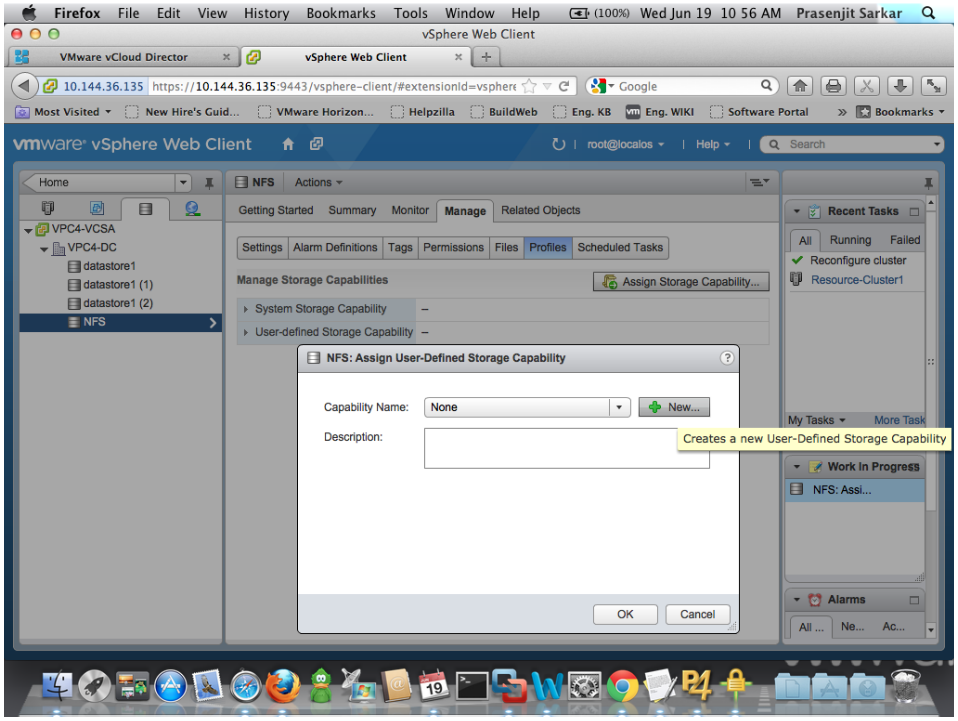
Firefox File Edit View History Bookmarks Tools Window Help
Wed Jun 19 10 56 AM Prasenjit Sarkar
vSphere Web Client
VMware vCloud Director
vSphere Web Client
https://10.144.36.135:9443/vsphere-client/#extensionId=vsphere
Most Visited
New Hire's Guid...
VMware Horizon...
Helpzilla
BuildWeb
Eng. KB
Eng. WIKI
Software Portal
Bookmarks
vmware vSphere Web Client
root@localos
Help
Search
Home
VPC4-VCSA
VPC4-DC
datastore1
datastore1 (1)
datastore1 (2)
NFS
NFS Actions
Getting Started Summary Monitor Manage Related Objects
Settings Alarm Definitions Tags Permissions Files Profiles Scheduled Tasks
Manage Storage Capabilities
Assign Storage Capability...
System Storage Capability –
User-defined Storage Capability –
Recent Tasks
All Running Failed
Reconfigure cluster
Resource-Cluster1
My Tasks
More Tasks
Work In Progress
NFS: Assi...
Alarms
All ... Ne... Ac...
NFS: Assign User-Defined Storage Capability
Capability Name:
New...
Description:
Creates a new User-Defined Storage Capability
OK
Cancel

Capability Name: None

Description:

Figure 2-14 Assign Storage Capability on Datastore

6. Click on the **New** button and specify a capability name.
7. Click on **OK** and **OK** again. Now we will create Storage Profile that will use this storage capability.
8. Click on the **Home** page in the vSphere Client.
9. Click on the **Rules and Profiles** section.
10. Click on the **VM Storage Profiles**.
11. Click on the **Create new Storage Profile** icon as shown in Figure 2-15.

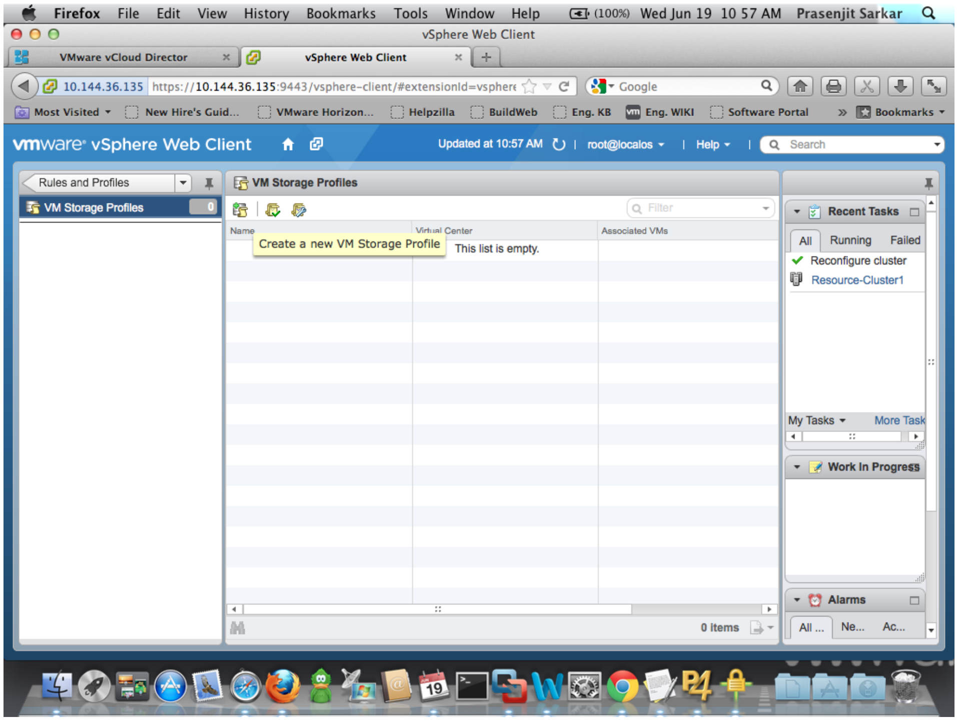

Figure 2-15 Storage Profile creation in vSphere Web Client

12. Specify a name in the Storage Profile and select the **Storage Capability** just created as shown in Figure 2-16.

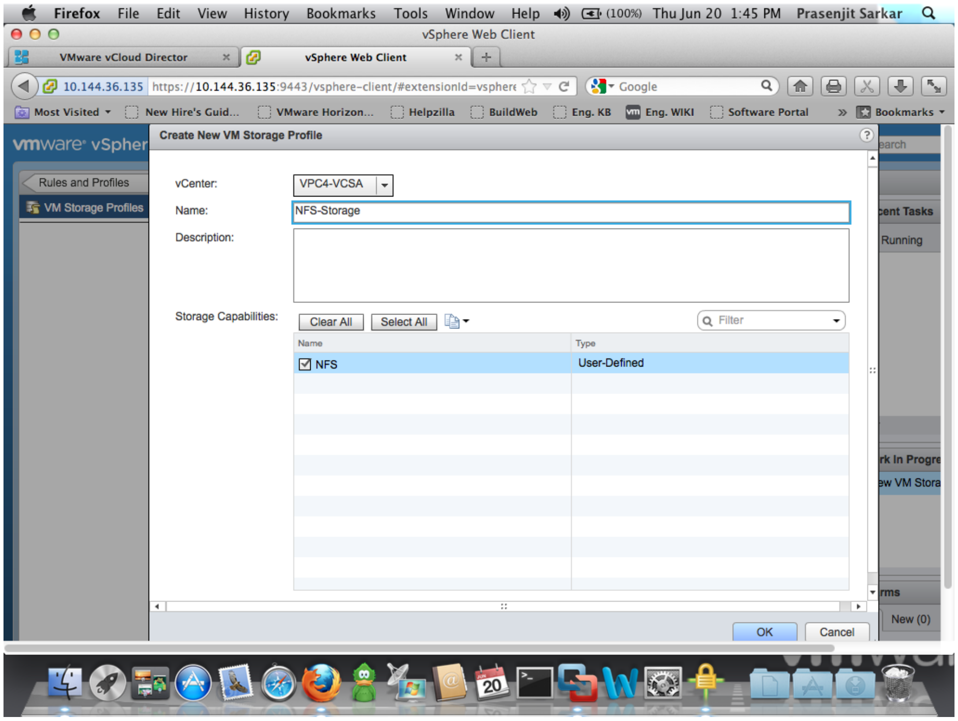

Create New VM Storage Profile

vCenter: VPC4-VCSA

Name: NFS-Storage

Description:

Storage Capabilities:

Name	Type
☑ NFS	User-Defined

Figure 1-16 Storage Profile creation in vSphere Web Client using Storage Capability

13. Click **OK**.
14. Now you need to enable VM storage profiles per cluster (Compute Resource). Click the **Enable VM Storage Profiles per Compute Resource** button as shown in Figure 2-17.

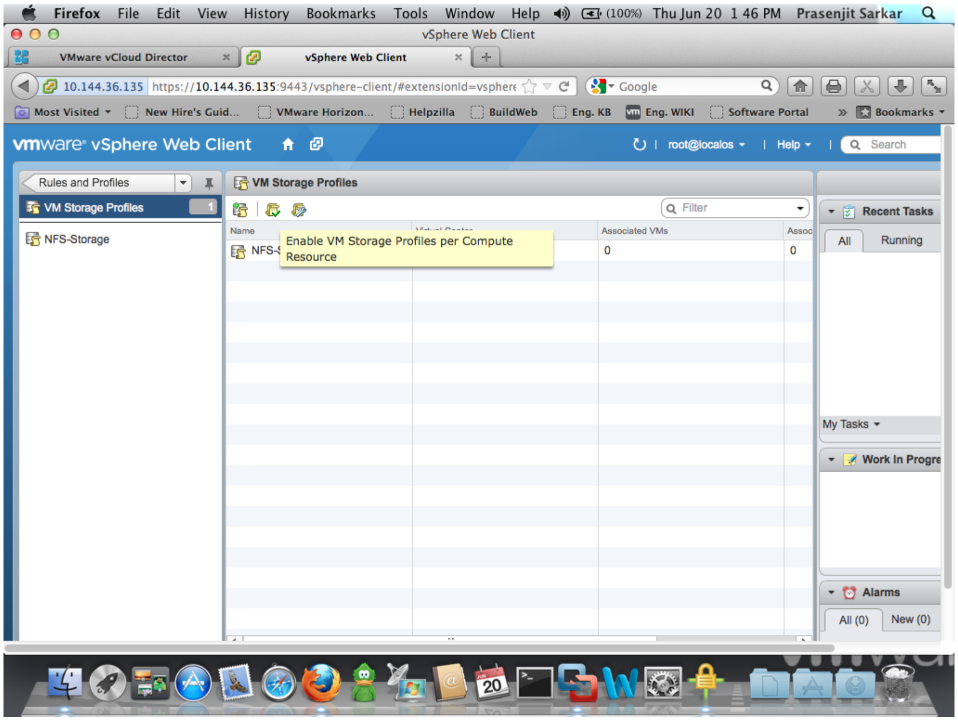
Firefox File Edit View History Bookmarks Tools Window Help
Thu Jun 20 1 46 PM Prasenjit Sarkar
vSphere Web Client
VMware vCloud Director
vSphere Web Client
https://10.144.36.135:9443/vsphere-client/#extensionId=vsphere
Most Visited New Hire's Guid... VMware Horizon... Helpzilla BuildWeb Eng. KB Eng. WIKI Software Portal Bookmarks
vmware vSphere Web Client
root@localos Help Search
Rules and Profiles
VM Storage Profiles
NFS-Storage
VM Storage Profiles
Filter
Name
Associated VMs
Assoc
NFS-S
Enable VM Storage Profiles per Compute Resource
0
0
Recent Tasks
All
Running
My Tasks
Work In Progre
Alarms
All (0)
New (0)

Figure 2-17 Enabling Storage Profile per Compute Resource - 1

15. Select the appropriate cluster and click on **Enable** button as shown in Figure 2-18.

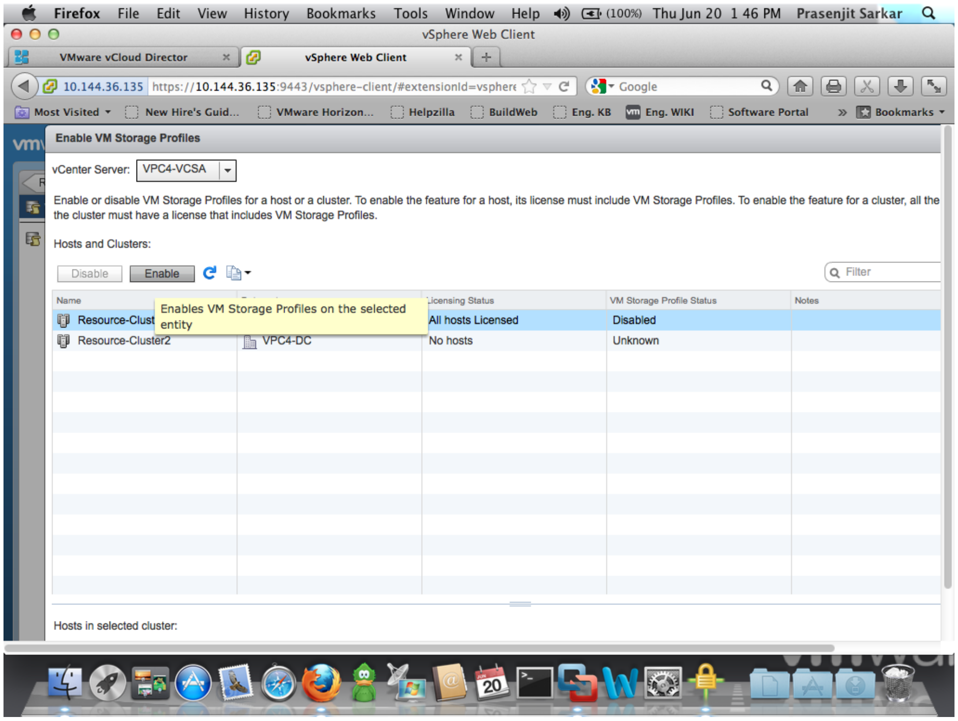

Enable VM Storage Profiles

vCenter Server: VPC4-VCSA

Enable or disable VM Storage Profiles for a host or a cluster. To enable the feature for a host, its license must include VM Storage Profiles. To enable the feature for a cluster, all the the cluster must have a license that includes VM Storage Profiles.

Hosts and Clusters:

Enables VM Storage Profiles on the selected entity

Name		Licensing Status	VM Storage Profile Status	Notes
Resource-Clust		All hosts Licensed	Disabled	
Resource-Cluster2	VPC4-DC	No hosts	Unknown	

Hosts in selected cluster:

Figure 2-18 Enabling Storage Profile per Compute Resource - 2

16. Click on the **Close** button.
17. Start the Internet Explorer browser. Go to the URL of the vCD server.

    An example would be https://serverFQDN/cloud.
18. Log in to vCD by typing an administrator user ID and password.
19. Click on the **Manage and Monitor** tab.
20. Select the vCenter section.
21. Right click on the vCenter Name and select **Refresh Storage Profiles** as shown in Figure 2-19. Now this Storage Profile will appear once you create a Provider vDC.

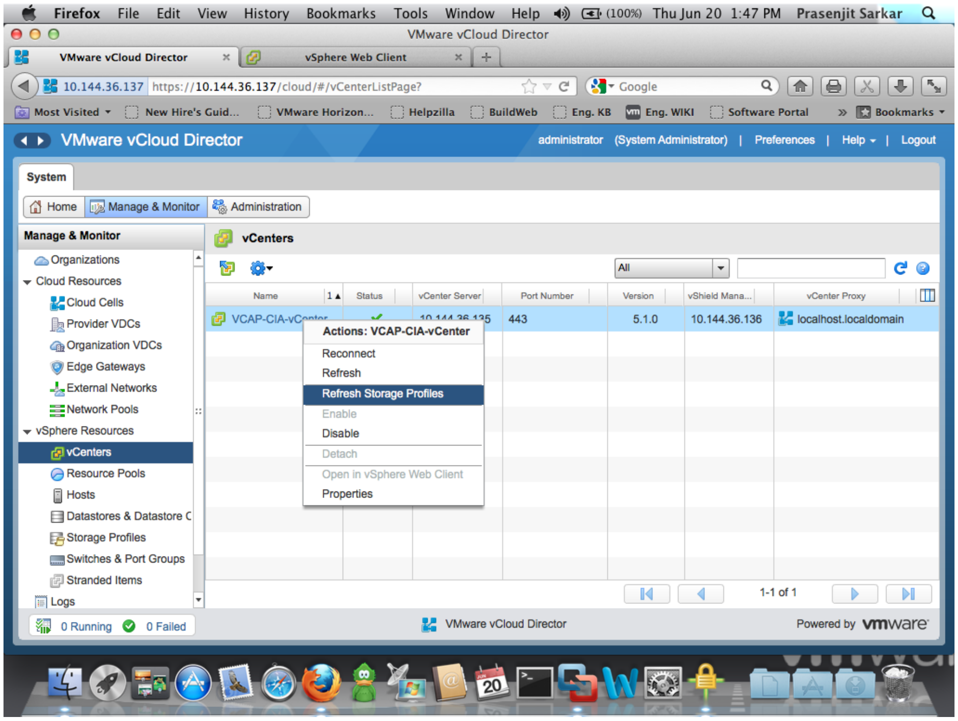

Name	Status	vCenter Server	Port Number	Version	vShield Mana...	vCenter Proxy
VCAP-CIA-vCenter	✓	10.144.36.135	443	5.1.0	10.144.36.136	localhost.localdomain

Figure 2-19 Refreshing Storage Profiles of vCenter in vCloud Director

You can determine exactly which vSphere Storage Profiles are attached to a specific datastore. You can see the number of the vDCs using each VMware vSphere Storage Profile (both provider and organization), the number of datastores in each vSphere Storage Profile, and how much space has been used, provisioned, and requested. You can also see how much storage space on the selected datastore is being actively used by the vSphere Storage Profile.

insert a Key Topic icon here

ACTIVITY 2-7: MANAGE & MONITOR VSPHERE STORAGE PROFILES IN VCLOUD DIRECTOR

To manage and monitor our vSphere Storage Profiles in vCloud Director, follow the steps:

1. Start the Internet Explorer browser. Go to the URL of the vCD server.

   An example would be https://serverFQDN/cloud.
2. Log in to vCD by typing an administrator user ID and password.
3. Click on the Manage & Monitor tab.
4. In the left pane, click **Storage Profiles** and you will see the Storage Profiles listed there as shown in Figure 2-20.

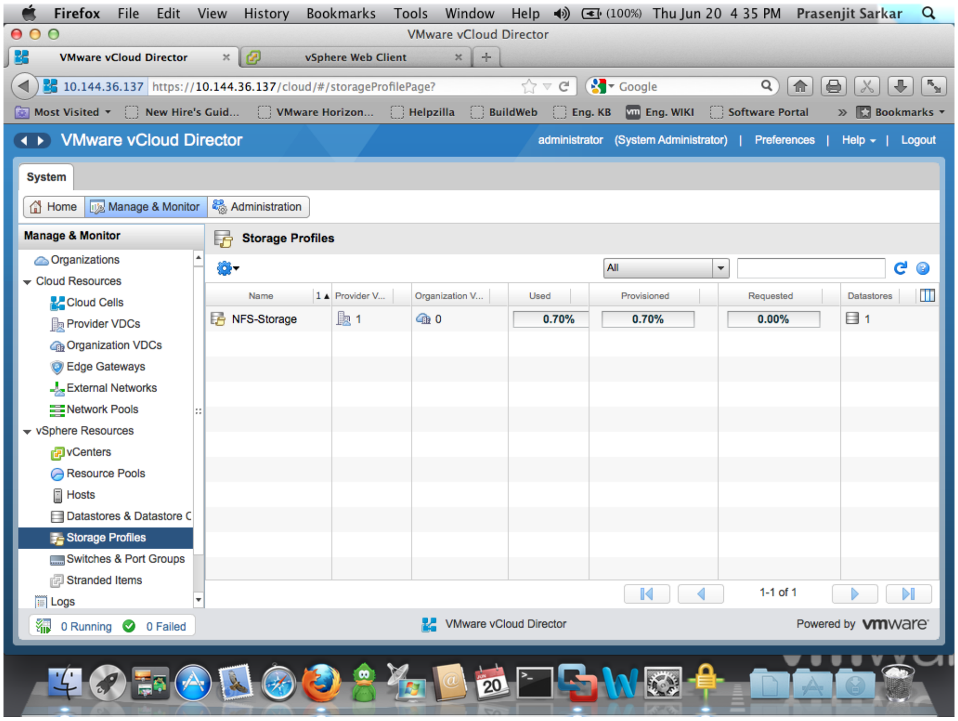

Name	Provider V...	Organization V...	Used	Provisioned	Requested	Datastores
NFS-Storage	1	0	0.70%	0.70%	0.00%	1

Figure 2-20 Storage Profiles in vCloud Director

5. Right click on the Storage Profile and click on **Properties**.

6. You will now see which Datastore and Datastore Clusters are in the selected Storage Profile, what percentage of storage has been used in the storage profile in each datastore as shown in Figure 2-21.

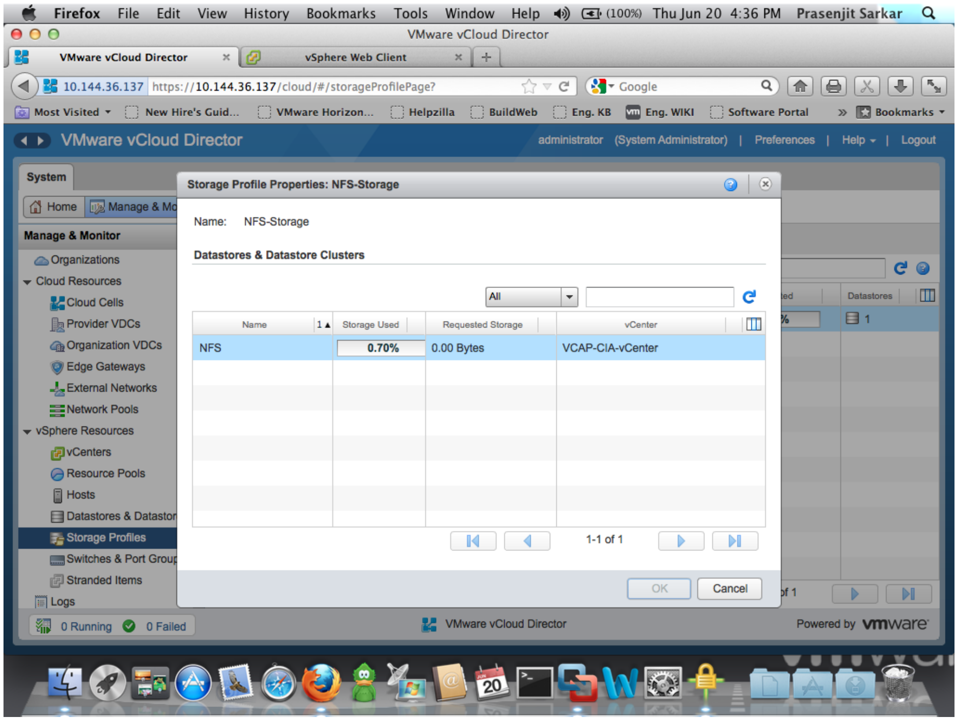

Storage Profile Properties: NFS-Storage

Name: NFS-Storage

Datastores & Datastore Clusters

All

Name	Storage Used	Requested Storage	vCenter
NFS	0.70%	0.00 Bytes	VCAP-CIA-vCenter

1-1 of 1

OK Cancel

Figure 2-21 Properties of Storage Profiles in vCloud Director

Manage vSphere Network Resources

In vCloud Director, you have different types of networks such as External Network and different types of network pools. External Networks are logical and based on vSphere port groups. These port groups can be derived from standard virtual switch, distributed virtual switch or even from third party distributed virtual switch which is Cisco Nexus 1000v. Each external network is backed by a single port group. As a best practice, you should use a vSphere Distributed Switch as a single distributed virtual switch can have several port groups in it and each one backing a different external network. If you want to create multiple external networks then it must be separated by VLANs. You should have these port groups already created in vSphere and then can use it in vCloud Director.

Sometimes you may consider this external network as internet facing. Though it is true that most of the times it is intended to be hooked up with internet, however, this is not mandatory. It is just an external network for your vCloud organizations.

It is utmost important for one to understand the how external networks at the provider side are built from vSphere networks. An example of the external network and org network connecting to it is shown in Figure 2-22.

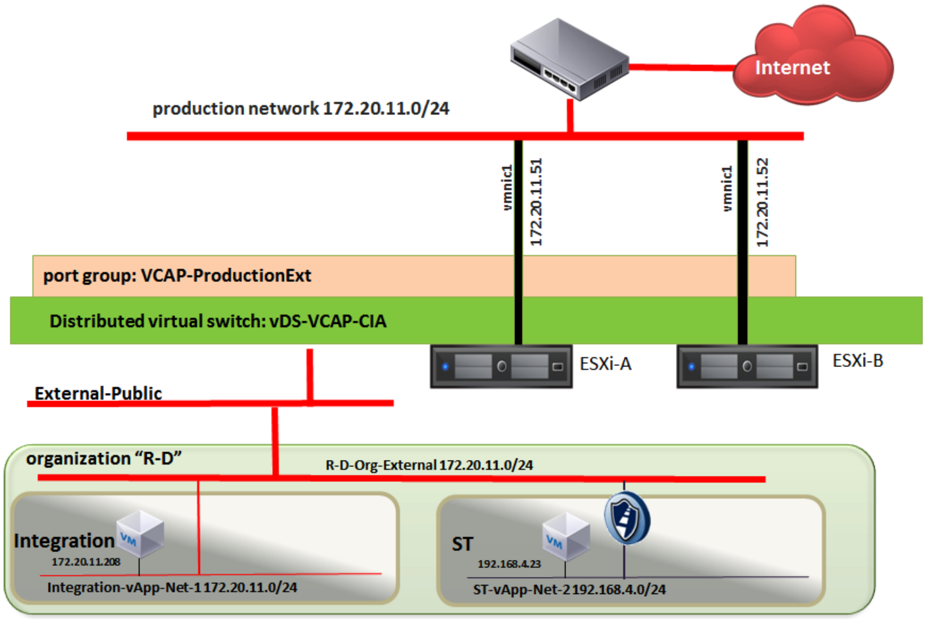
Internet
production network 172.20.11.0/24
vmnic1
172.20.11.51
vmnic1
172.20.11.52
port group: VCAP-ProductionExt
Distributed virtual switch: vDS-VCAP-CIA
ESXi-A
ESXi-B
External-Public
organization "R-D"
R-D-Org-External 172.20.11.0/24
Integration
VM
172.20.11.208
Integration-vApp-Net-1 172.20.11.0/24
ST
VM
192.168.4.23
ST-vApp-Net-2 192.168.4.0/24

Figure 2-22 Logical Diagram of External Network in vCloud Director

Here you can see External-Public, a provider-level external network, is built off of the VCAP-ProductionExt port group. VCAP-ProductionExt port group is located in the vDS-VCAP-CIA distributed virtual switch. The hosts ESXi01 and ESXi02 and connected to the vDC-VCAP-CIA distributed virtual switch.

insert key topic icon here

ACTIVITY 2-8 CREATE AND MANAGE VSPHERE PORT GROUPS

To create and manage vSphere Port Groups for vCloud Director, follow the steps:

1. Login to the vSphere Web Client as root
2. You will be landed on the Home screen, click on **Networking** link there.
3. Expand the DVSwitch on the left side of the panel.
4. Right click on the DVS and click on **New Distributed Port Group** as shown in Figure 2-23.

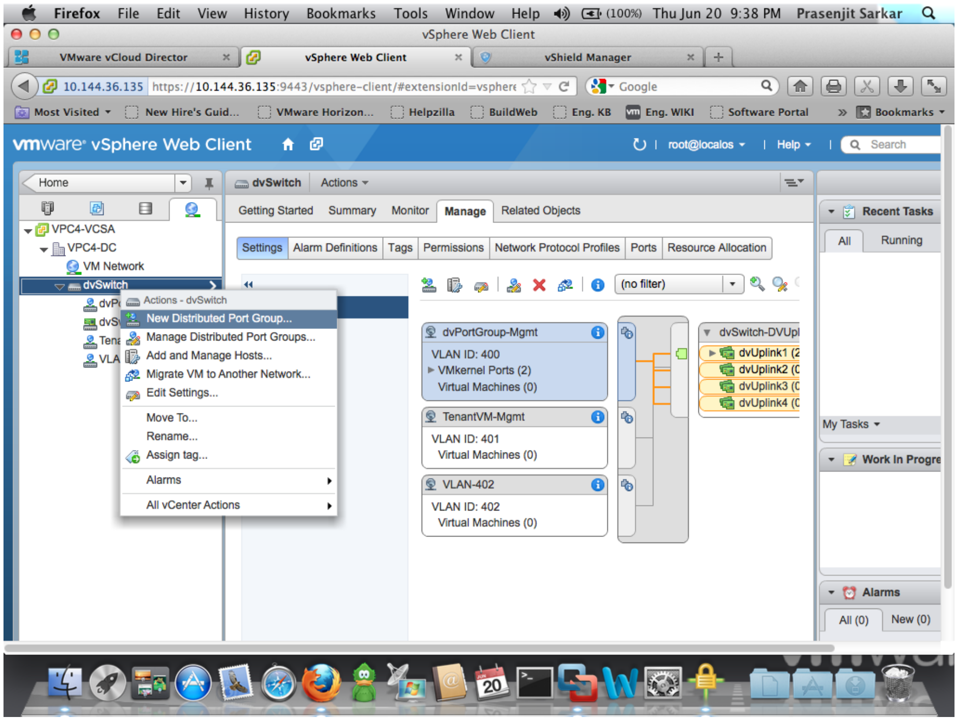

Figure 2-23 Creating Distributed Port Group in vSphere Web Client

5. Specify a name of the port group and click on **Next**.
6. Specify a VLAN number (Optional) as shown in Figure 2-24 and click on **Next**.

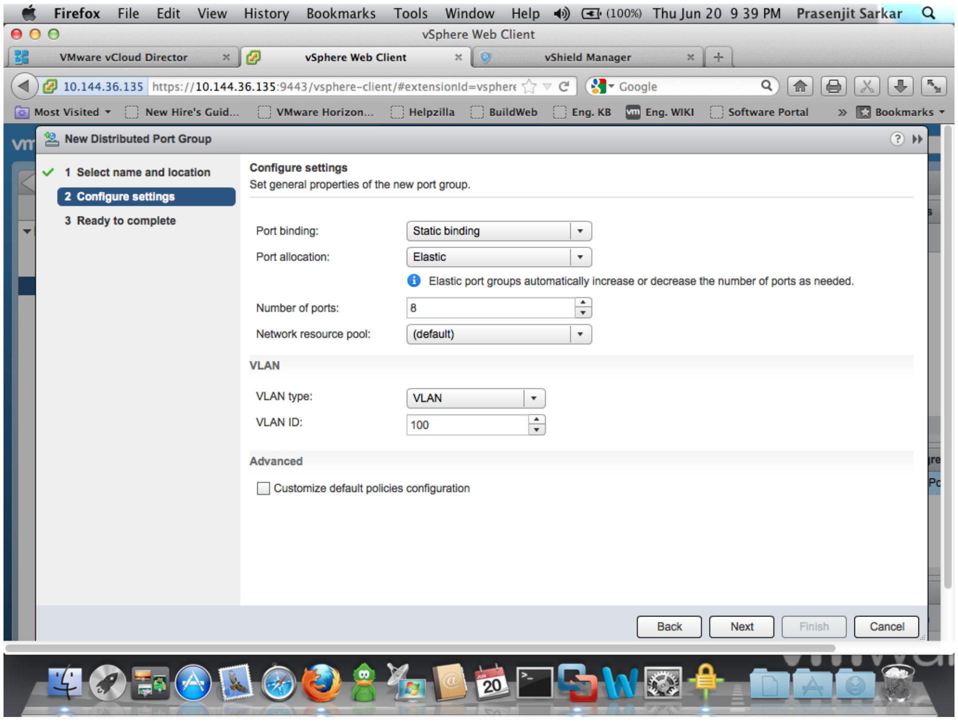

New Distributed Port Group

1 Select name and location
2 Configure settings
3 Ready to complete

Configure settings

Set general properties of the new port group.

Port binding: Static binding

Port allocation: Elastic

Elastic port groups automatically increase or decrease the number of ports as needed.

Number of ports: 8

Network resource pool: (default)

VLAN

VLAN type: VLAN

VLAN ID: 100

Advanced

☐ Customize default policies configuration

Figure 2-24 Distributed Port Group creation process in vSphere Web Client

7. Click on **Finish** to create this Port group.

You also need to change the MTU settings of the vDS so that you can use this vDS for your VXLAN deployment which will be discussed after this activity.

1. On the **Networking** page, expand the vDS there on the left side.
2. Click on the vDS and at the right hand side click on **Manage** Tab.
3. Click on the **Properties** in the left side and under the **Advanced** section, you should see **MTU** as **1500 bytes** as shown in Figure 2-25.

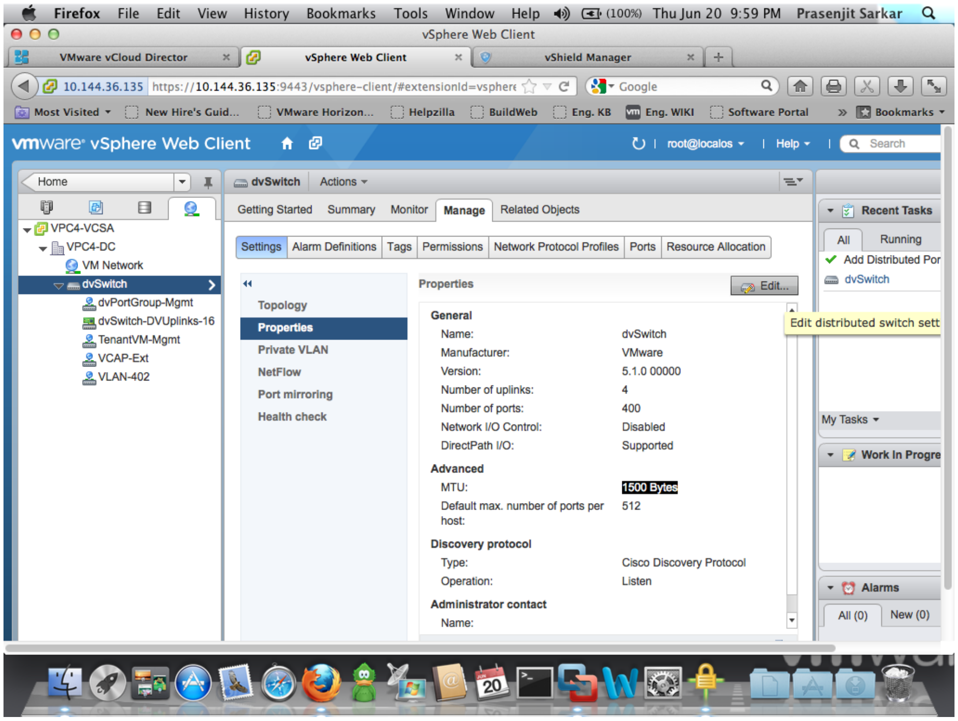
Firefox File Edit View History Bookmarks Tools Window Help
Thu Jun 20 9:59 PM Prasenjit Sarkar
vSphere Web Client
VMware vCloud Director
vSphere Web Client
vShield Manager
https://10.144.36.135:9443/vsphere-client/#extensionId=vsphere
Most Visited New Hire's Guid... VMware Horizon... Helpzilla BuildWeb Eng. KB Eng. WIKI Software Portal Bookmarks
vmware vSphere Web Client
root@localos Help Search
Home
VPC4-VCSA
VPC4-DC
VM Network
dvSwitch
dvPortGroup-Mgmt
dvSwitch-DVUplinks-16
TenantVM-Mgmt
VCAP-Ext
VLAN-402
dvSwitch Actions
Getting Started Summary Monitor Manage Related Objects
Settings Alarm Definitions Tags Permissions Network Protocol Profiles Ports Resource Allocation
Topology
Properties
Private VLAN
NetFlow
Port mirroring
Health check
Edit...
Recent Tasks
All Running
Add Distributed Por
dvSwitch
Edit distributed switch sett
My Tasks
Work In Progre
Alarms
All (0) New (0)

Properties

General

Name:	dvSwitch
Manufacturer:	VMware
Version:	5.1.0 00000
Number of uplinks:	4
Number of ports:	400
Network I/O Control:	Disabled
DirectPath I/O:	Supported

Advanced

MTU:	1500 Bytes
Default max. number of ports per host:	512

Discovery protocol

Type:	Cisco Discovery Protocol
Operation:	Listen

Administrator contact

Name:

Figure 2-25 Distributed Switch properties in vSphere Web Client

4. Click on the **Edit** button and then click on **Advanced** link on the left side.
5. Change the **MTU** from 1500 to 1600 as shown in Figure 2-26 and click on **OK**.

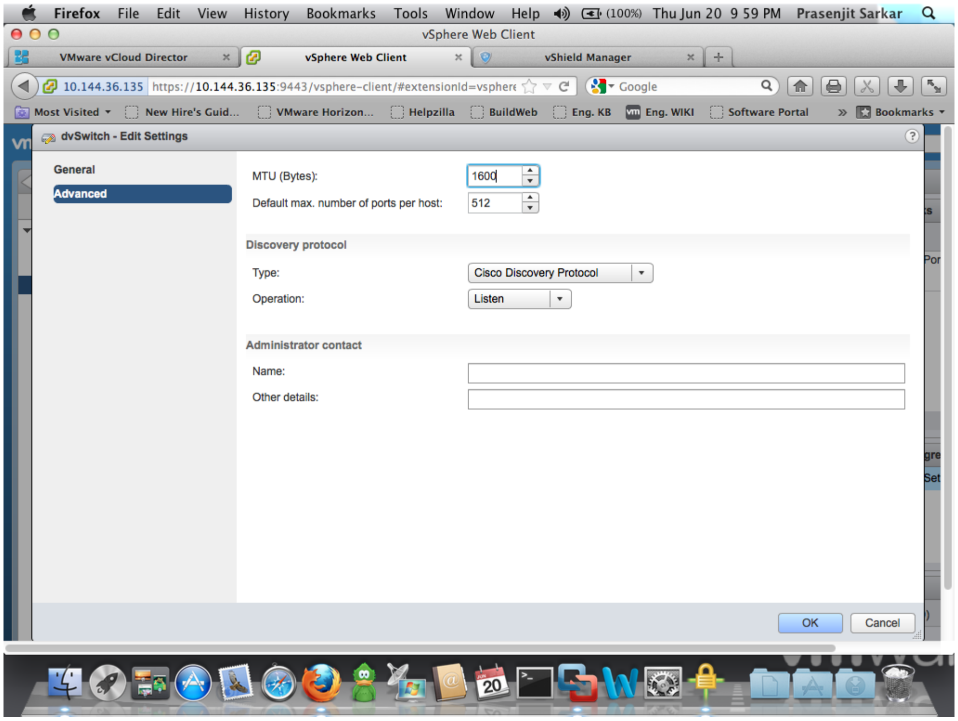
Firefox File Edit View History Bookmarks Tools Window Help
Thu Jun 20 9 59 PM Prasenjit Sarkar
vSphere Web Client
VMware vCloud Director
vSphere Web Client
vShield Manager
https://10.144.36.135:9443/vsphere-client/#extensionId=vsphere
Most Visited New Hire's Guid... VMware Horizon... Helpzilla BuildWeb Eng. KB Eng. WIKI Software Portal Bookmarks
OK
Cancel

dvSwitch - Edit Settings

General

Advanced

MTU (Bytes): 1600

Default max. number of ports per host: 512

Discovery protocol

Type: Cisco Discovery Protocol

Operation: Listen

Administrator contact

Name:

Other details:

Figure 2-26 Changing MTU of Distributed Virtual Switch in vSphere Web Client

VMware network virtualization uses the virtual extensible local area network (VXLAN) overlay networking technology to create a logical network. A VXLAN based network virtualization solution addresses several challenges that are faced in a traditional physical network. VXLAN virtual networking must be configured correctly for the layer 2, logical broadcast domain to operate properly.

Some of the benefits of VXLAN are:

- Allow noncontiguous cluster expansion
- Allow availability domains within a DC
- Leverage capacity across multiple L2 POD's
- Overcome IP addressing challenges
- VLAN sprawl
- VLAN scale

VXLAN enables multi tenancy at scale without encountering VLAN scale limitations; however, VXLAN implementation requires VMware vSphere Distributed Switch 5.1.

VXLAN frames are Ethernet frames that are carried in IP packets across the physical infrastructure. It increases the layer-2 broadcast domain space. IP and UDP encapsulation are used to build large-scale fabrics. It also permits LACP hashing and equal cost multipath (ECMP) routing. For VXLAN to work properly, MTU for the physical network and vDS must be at least 1600 bytes.

To accommodate the VXLAN encapsulation overhead, at least a 1600 byte MTU is required which will help to avoid packet fragmentation. However physical infrastructure must carry 50 bytes more than the virtual machine vNIC MTU size.

Table 2-2 MTU classification for IP classes

	IPv4 (bytes)	**IPv6 (bytes)**
Guest frame	1514	1514
Guest VLAN tag	4	4
VXLAN transport header	50	70
VXLAN transport VLAN tag	4	4
IPv6 data and control	-	8
MTU size	1572	1600

VXLAN relies on IP multicast in the physical infrastructure for broadcast, multicast, and unknown unicast traffic. VXLAN Network ID map to multicast groups (1:1 or m:n mapping). Now you may wonder how it works.

In VXLAN, Virtual machine-to-virtual machine traffic is tunneled over a layer 3 network by a VXLAN module in each ESXi host. In this process node learning is done through multicast, not broadcast.

Some of the mandatory components for VXLAN are as follows:

- VMware vSphere Enterprise Plus Edition and VMware vCloud Networking and Security licenses are required.
- VMware vShield Manager
- vSphere Distributed Switch
- Virtual Tunnel End Point (VTEP)
- VMware vShield Edge

A logical layer 2 across layer 2 (VLAN) network with a single vCenter Server looks as Figure 2-27.

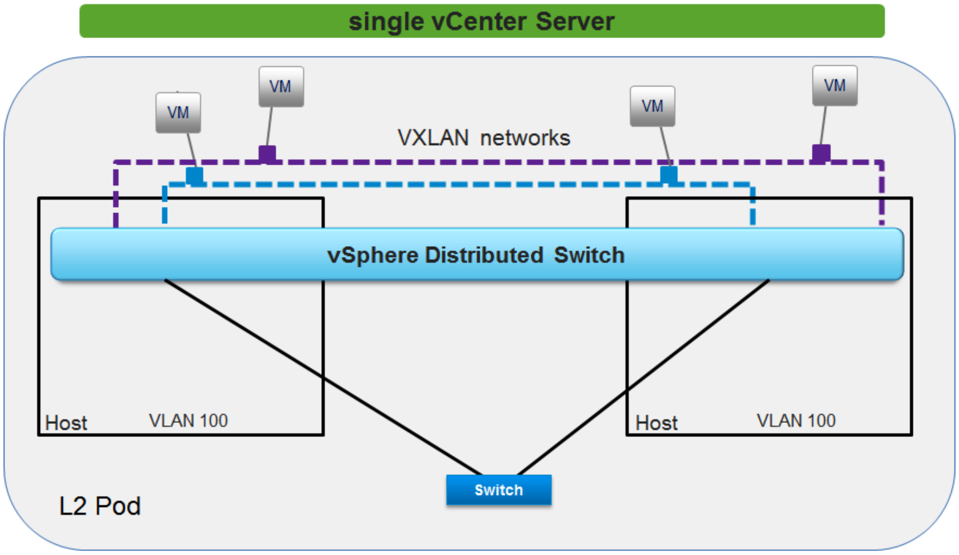

Figure 2-27 Logical Layer 2 across Layer 2

Now Figure 2-28 shows logical layer 2 networks across layer 3 over multiple vDS.

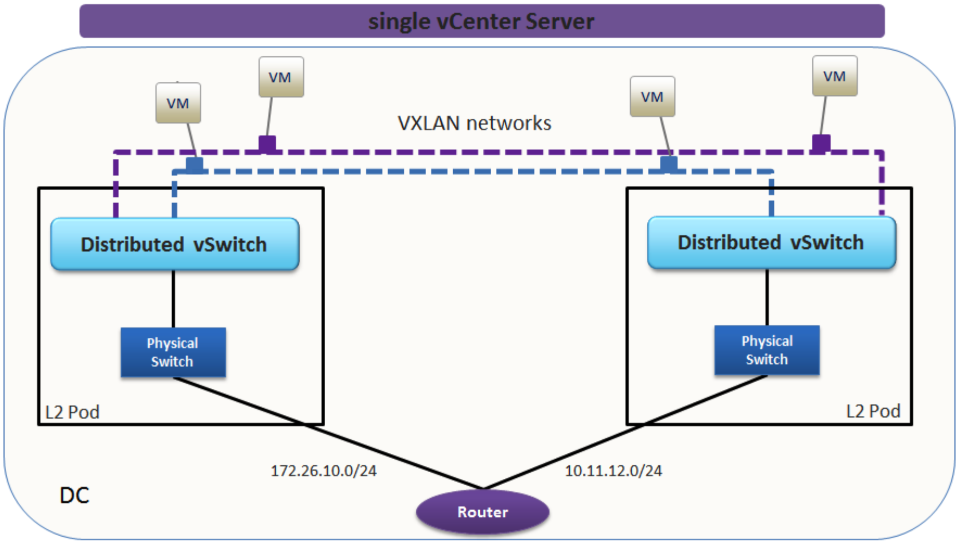

Figure 2-27 Logical Layer 2 across Layer 3 with multiple vDS

Figure 2-29 shows logical layer 2 networks across layer 3 over a single vDS.

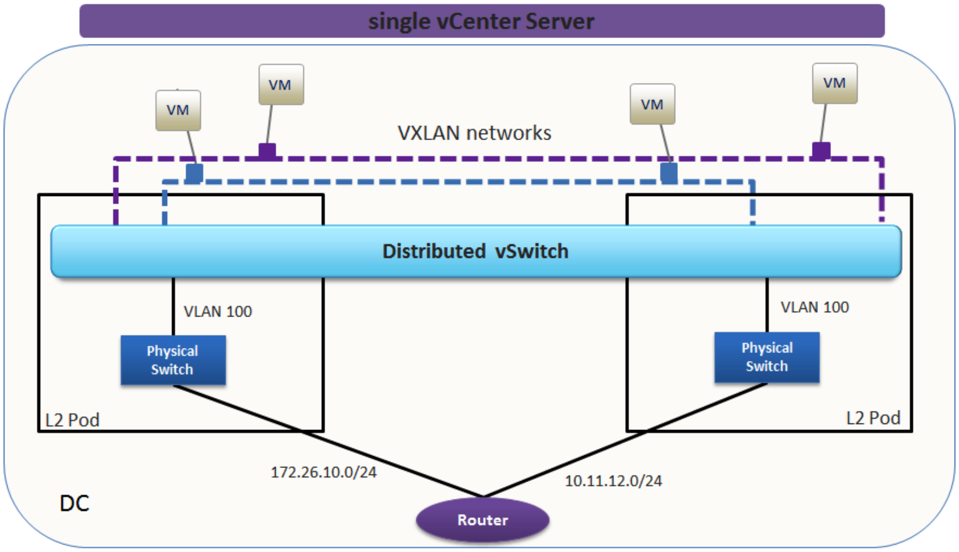

Figure 2-29 Logical Layer 2 across Layer 3 with single vDS

VXLAN networks use multicast for broadcast, multicast, and unknown unicast traffic and for those two protocols are used:

Internet Group Management Protocol (IGMP)

- Used by L2 switches to learn which hosts are subscribed to which multicast groups
- IGMP snooping – Controls which ports receive specific multicast traffic
- IGMP snooping querier – Issues periodic IGMP queries that trigger IGMP report messages from hosts

Protocol-Independent Multicast (PIM)

- Used by routers to exchange multicast membership information
- Use either the Sparse or Bidirectional Multicast mode
- PIM Sparse Mode (PIM-SM) is suitable for groups where a very low percentage of the nodes will subscribe to the multicast session.
- Bidirectional PIM scales better than Sparse mode.

If you are wondering that how a logical representation of two VMs in a same VXLAN looks like then Figure 2-30 shows the same.

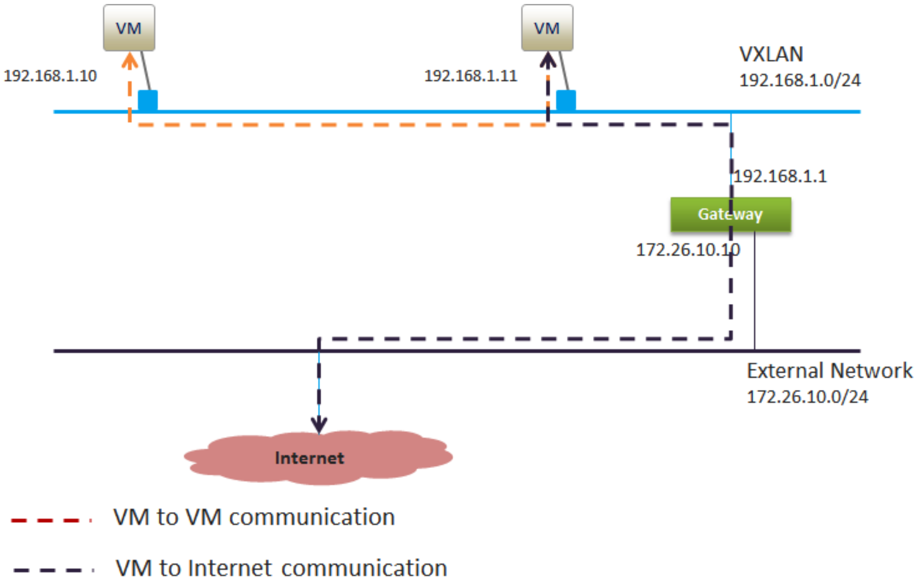

Figure 2-30 Two VMs in same VXLAN

Similarly Figure 2-31 shows a logical representation of two VMs in two different VXLAN.

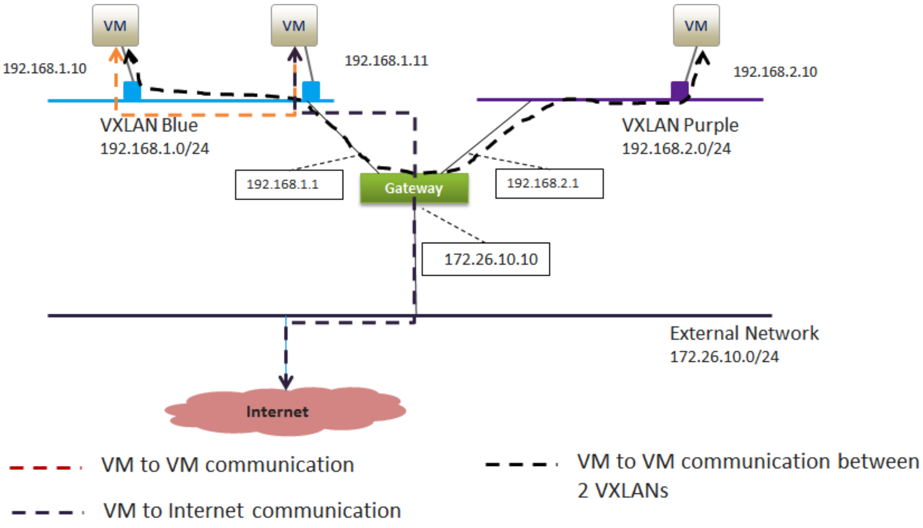

Figure 2-31 Two VMs in two different VXLAN

insert key topic icon here

ACTIVITY 2-9: PREPARE VSPHERE CLUSTER FOR VXLAN

To prepare your vSphere cluster for VXLAN, follow the steps:

1. Open up Internet Explorer and login to the vShield Manager.
2. In the left pane, expand the **Datacenters** container and select your datacenter item.
3. In the right pane, click the **Network Virtualization** tab.
4. Click the **Preparation** link.
5. Click on the **Segment ID** button.
6. Click on the **Edit** button located on the far right of the vShield Manager interface as shown in Figure 2-32.

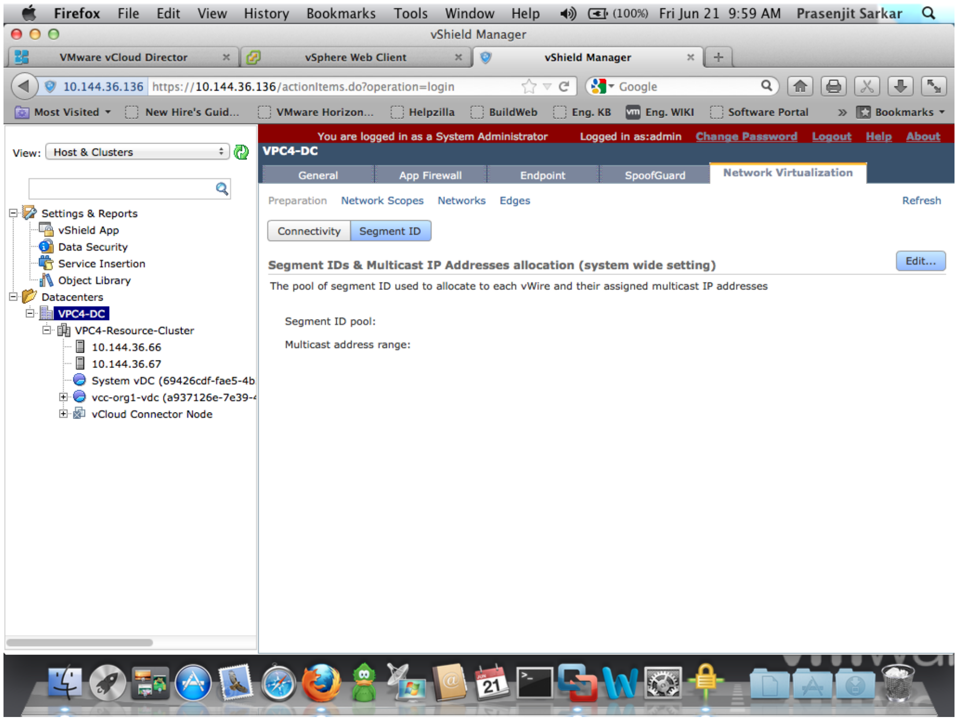

You are logged in as a System Administrator

Logged in as:admin

Segment IDs & Multicast IP Addresses allocation (system wide setting)

The pool of segment ID used to allocate to each vWire and their assigned multicast IP addresses

Segment ID pool:

Multicast address range:

Figure 2-32 VXLAN Preparation in vShield Manager

7. Specify a segment ID pool and multicast address range to help distribute VXLAN traffic across the physical infrastructure as shown in Figure 2-33.

   The lowest segment ID allowed is 5000. Multicast addresses are allocated by your network administrator.

   VXLAN Segment ID is a 24-bit value used to designate the individual VXLAN overlay network on which the communicating virtual machines are connected.

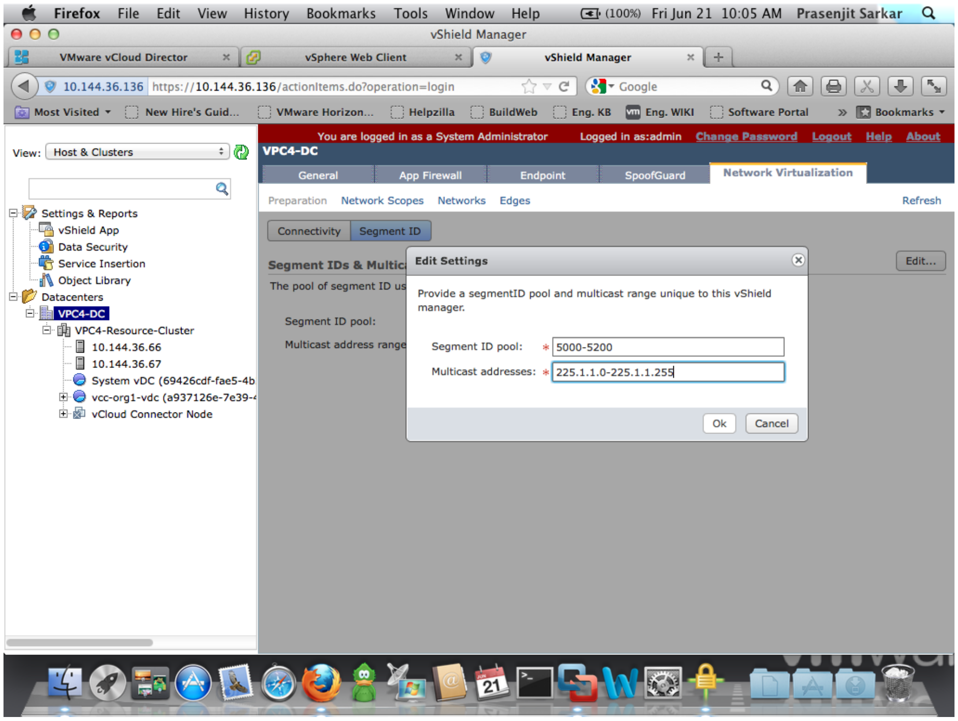

Provide a segmentID pool and multicast range unique to this vShield manager.

Segment ID pool: * 5000-5200

Multicast addresses: * 225.1.1.0-225.1.1.255

Figure 2-33 VXLAN Segment ID Configuration

8. Click on **OK**.
9. Now we need to configure the cluster connectivity. Click on the **Connectivity** button.
10. Click the **Edit** button located on the far right of the vShield Manager interface.
11. Select the vSphere Cluster, its associated Distributed vSwitch and specify VLAN as shown in Figure 2-34 and click on **Next**.

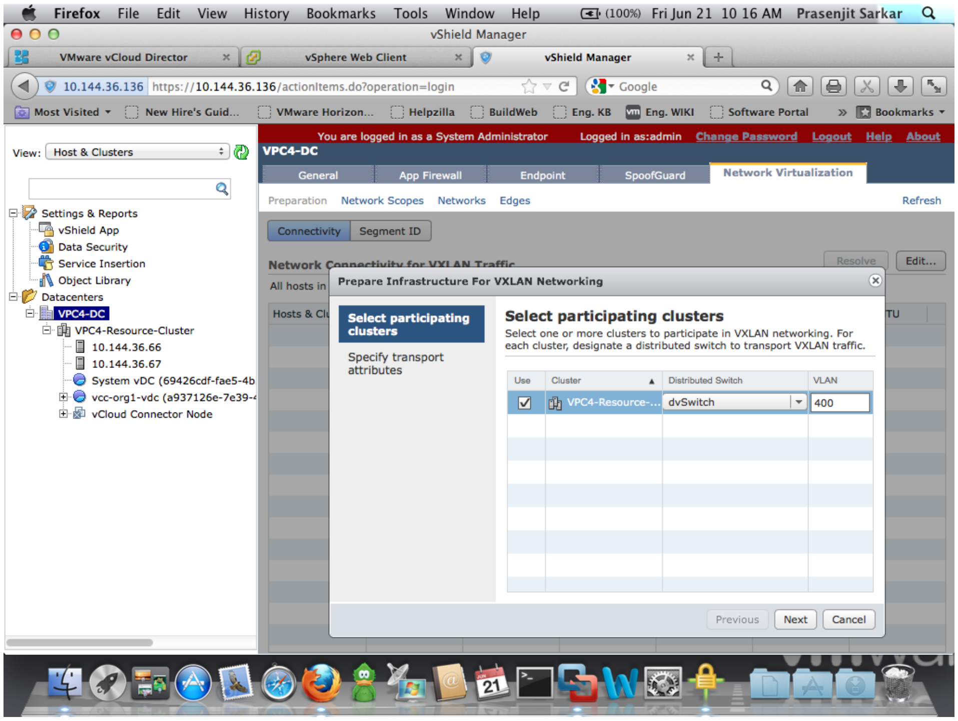

Prepare Infrastructure For VXLAN Networking

Select participating clusters

Select one or more clusters to participate in VXLAN networking. For each cluster, designate a distributed switch to transport VXLAN traffic.

Use	Cluster	Distributed Switch	VLAN
☑	VPC4-Resource-...	dvSwitch	400

Figure 2-34 VXLAN Cluster Configuration - 1

12. Under Specify transport attributes, select **Fail Over** from the Teaming Policy drop-down menu, for your distributed switch.

13. Leave the MTU (bytes) value at 1600 for your distributed switch as shown in Figure 2-35.

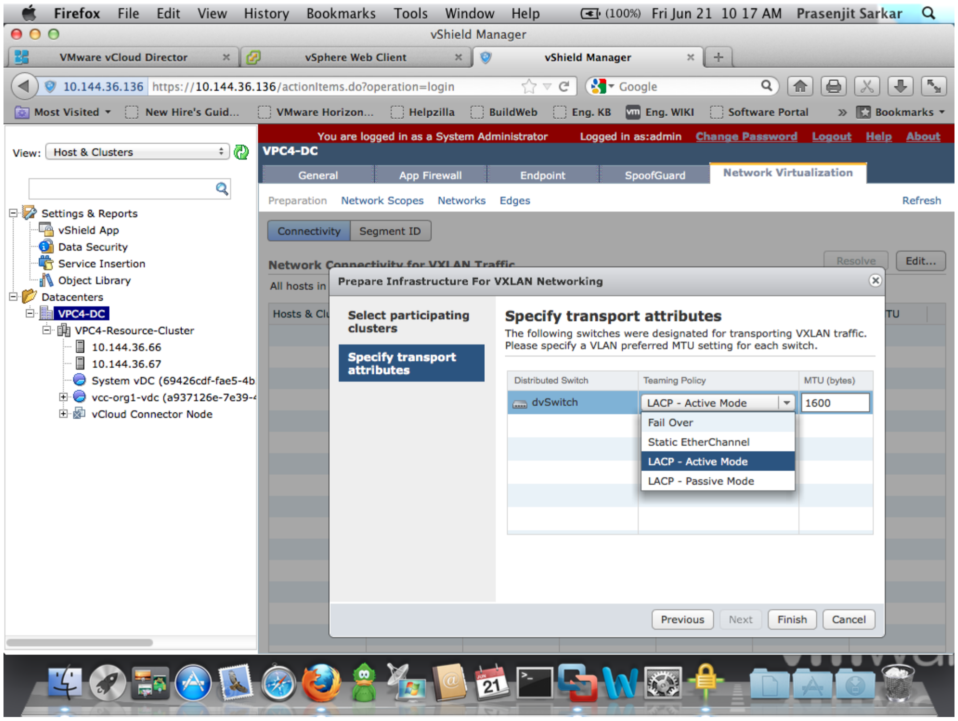

Prepare Infrastructure For VXLAN Networking

Specify transport attributes

The following switches were designated for transporting VXLAN traffic. Please specify a VLAN preferred MTU setting for each switch.

Distributed Switch	Teaming Policy	MTU (bytes)
dvSwitch	LACP - Active Mode	1600

Figure 2-35 VXLAN Cluster Configuration - 2

14. Click **Finish**.

15. Verify cluster and host readiness.

When the status of your cluster is ***Ready***, expand the cluster item and verify the following:

- The status of each ESXi host is ***Ready***.
- Each vmnic has acquired an appropriate DHCP assigned address. (You can use static address as well).

It maps each cluster that is to access the VXLAN network to a distributed switch. vmkernel modules (the VTEP) are pushed and enabled on all of the hosts in the cluster and all hosts in the cluster are automatically enabled for VXLAN networking.

As a process new dvPort groups and vmknic interfaces are created on the distributed switch. The vmknic which will be created will have DHCP enabled. However, you can configure a static address. You can identify the new dvPort groups by the unique naming convention, vxw-dvs-xxx-virtualwire-xxxx.

Exam Preparation Tasks

Review All the Key Topics

Review the most important topics from inside the chapter, noted with the Key Topic icon in the outer margin of the page. Table 2-3 lists these key topics and the page numbers where each is found.

Table 2-3 Key Topics for Chapter 2

Key Topic Element	Description	Page Number
Activity 2-1	Add New vCenter Server to vCloud Director	
Activity 2-2	Unprepare ESXi Hosts in vCloud Director	
Activity 2-3	Add ESXi Hosts to vCenter and Prepare ESXi Host in vCloud Director	
Activity 2-4	Manage and Monitor vSphere Resource Pools	
Activity 2-5	Manage VMware ESXi Hosts in vCloud Director	
Activity 2-6	Configure Storage Profiles in vCenter	
Activity 2-7	Manage & Monitor vSphere Storage Profiles in vCloud Director	
Activity 2-8	Create and Manage vSphere Port Groups	
Table 2-2	MTU classification for IP classes	
Activity 2-9	Prepare vSphere Cluster for VXLAN	

Review Questions

The answers to these review questions are in Appendix A.

1. What are the fields you can monitor for a Resource Pool in vCloud Director?

- **a.** Memory Reservations
- **b.** CPU Reservations
- **c.** Datastore Utilization
- **d.** All of the above

2. What is the operation a Redeploy all VMs does on a disabled ESXi host in vCloud Director?

    a. vMotion the VMs

    b. Shutdown the VMs

    c. Restart the VMs

    d. Recreate the VMs

3. What does Refresh Storage Profiles do in vCloud Director?

    a. Create new storage profile.

    b. Delete storage profiles.

    c. Update storage profiles.

    d. Refresh the changes made in vSphere for storage profile.

4. Storage Profile properties give us the ability to monitor:

    a. Datastores and Datastore Clusters.

    b. %Storage has been used.

    c. Datastores in each Storage Profile.

    d. All of the above.

5. Benefits of using VXLAN include?

    a. Allow noncontiguous cluster expansion.

    b. Allow availability domains within a DC.

    c. Leverage capacity across multiple L2 POD's.

    d. Overcome IP addressing challenges.

    e. VLAN sprawl

    f. VLAN scale

    g. All of the above

6. What are the mandatory components required for VXLAN?

    a. VMware vSphere Enterprise Plus Edition and VMware vCloud Networking and Security licenses.

    b. VMware vShield Manager.

    c. vSphere Distributed Switch.

    d. Virtual Tunnel End Point (VTEP).

    e. VMware vShield Edge
    f. All of the above

7. VXLAN Segment ID should start at?

    a. 2000.
    b. 4000.
    c. 5000.
    d. 5200.

8. VXLAN Segment ID is a ___ bit value?

    a. 10.
    b. 20.
    c. 24.
    d. 50.

vmware® PRESS

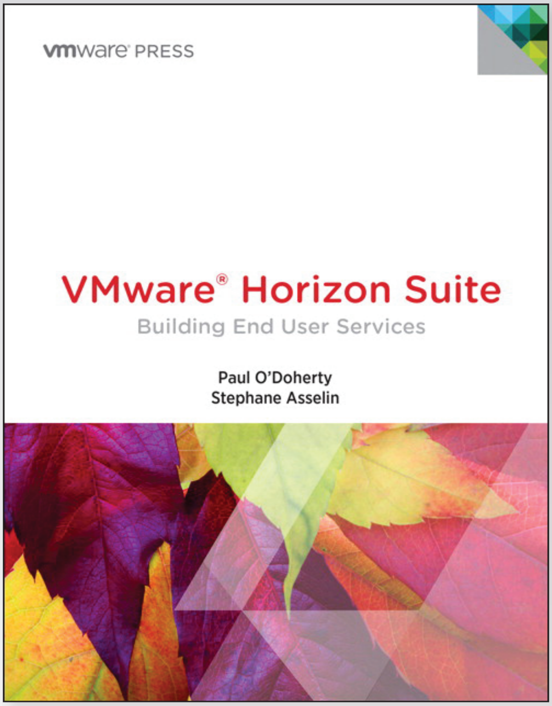

vmware® PRESS

VMware® Horizon Suite

Building End User Services

Paul O'Doherty
Stephane Asselin

CHAPTER 1

The New End User Model

AVAILABLE – MARCH/APRIL 2014

vmwarepress.com

Chapter 1

The New End User Model

The Evolution of the End User

VMware's product line has rapidly evolved to become a more complete end user suite. VMware acquired Wanova and their flagship product Mirage. The founders of Wanova had extensive experience in Wide Area Networking (WAN) and had already developed and sold a WAN services company, Actona that was acquired by CISCO. The experience in WAN optimization was the basis of the company name Wanova.

Mirage however was designed to centralized desktop images within the datacenter. Mirage used layering to separate the management and delivery of OS images. The end point image is cloned so into the datacenter however a local cached copy is still available for local execution.

View runs on a vSphere environment and supports an offline mode for checking out a View desktop to a Windows based endpoint running the local mode client. The local mode is dependent on the underlying operating system so is considered a layer 2 hypervisor. Mirage does not have this dependency on an underlying OS and is designed to make use of the resources on the end point vs. the datacenter. VMware announced that View would be known a VMware Horizon View in order to align it closer with VMware Horizon Mirage and Workspace.

VMware Horizon Workspace is targeted at delivering services in a way that is more tablet, mobile and Cloud friendly. It became clear that if we wanted to provide a current reference we would have to address more than just Horizon View. As the technology has evolved so too has our original subject matter. Anyone considering deploying virtual desktops needs to consider a post PC era. Even as we write this book the trending data is that tablets are quickly eclipsing sales of laptops and desktops. What is even more interesting is the speed in which this is occurring. Although it seems as if tablets have been around for a while in, they have only been available in the market place for three.

In this book we will use several terms interchangeably. When describing a virtual desktop we may also use the term virtual instance, View desktop and the aforementioned virtual desktop. When describing the larger Horizon View Environment we will use View infrastructure and Virtual Desktop Infrastructure (VDI) as well as the abbreviation View. When describing VMware Horizon Mirage and Workspace, both Mirage and Workspace will be used to reference the products. As we have much more to cover, we will focus on the core products and there related architecture. It is assumed that the reader has an understanding of vSphere and its related components so these items will not be covered in this book.

This book will cover Horizon View 6, Horizon Mirage and Horizon Workspace. For each we will review the architecture, planning considerations and how to properly install and configure the environment.

An End User Service Catalog

The understanding of Cloud technology will become as fundamental as the understanding of virtualization is in IT today. One of the key concepts behind Cloud is the idea of a service catalog. A service catalog is not a Virtual Machine (VM) or a collection of VMs but rather a complete solution that is designed to be consumed by the end user. In making the transition from supplying virtual desktops to delivering services to end users this is an important concept to understand.

How can you deploy virtual desktops, provide end user driven storage repositories, integrate Software-as-a-Service (SaaS) applications and traditional window applications in a service? This is the value message of Horizon Workspace which integrates all these services.

You will see the idea of a service catalog used heavily in Horizon Workspace where a View desktop is one of the services you can entitle a user to. Entitlement is the term used across the end user product line to enable a service for a user. In Horizon View you entitle desktops to Active directory users or groups. In Horizon Workspace you entitle users to access different services to build a service catalog as shown in Figure 1.1.

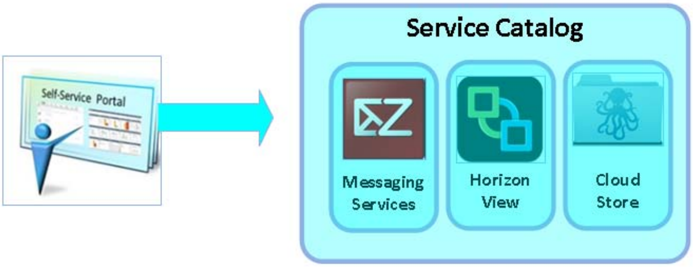

Figure 1.1 Service Catalog

To make the transition to the post PC era you must consider entitling users to service catalogs of which a View desktop is an important component. In this way you can ensure that the planning you are doing now will allow you to service any type of end user device without the reengineering of the entire environment. Desktops will be around for some time however as while there is allot of development of Cloud based apps, a large percentage of our day-to-day applications are still using traditional enterprise based client-server applications. The goal of this book is to allow you to deploy a single framework that leverages Horizon View, Horizon Workspace and Horizon Mirage to deal with any multitude of end user service requirements.

How Do View, Mirage and Horizon WorkspaceDeliver a Service Catalog

Delivering a service catalog involves understanding what is delivered by each of the products within the Horizon Suite. When it was View it was pretty straightforward. With Horizon Suite you have View, Mirage and Horizon Workspace.

VMware Horizon View

Each component of the Horizon Suite delivers different business value and meets different end user requirements. Horizon View delivers virtual desktops running in the datacenter from a centralized vSphere environment. It provides several technologies designed to manage key components of the end user experience; Persona for end user data management, Composer for deployment and image management and ThinApp to enable delivery without interoperability problems. These technology components are centrally managed through the View Connection server to deliver a consistent and robust end user experience enabling the desktop to be delivered as a service.

VMware Horizon Mirage

Horizon Mirage provides centralized image management so is similar to View however it allows the image to run decentralized on the endpoint. It supports an offline mode so can provide services to laptops without the prerequisite of running an underlying OS. This allows desktops services to be delivered to physical offline devices with the same level of consistency and control as a View instance running in the datacenter. Mirage allows you to extend your service catalog to deliver centralized and decentralized desktop management.

VMware Horizon Workspace

Workspace provides the Software as a Service interface to enable you to present Desktops as a Service (DaaS). In addition to presentation it integrates other Cloud like services such as unifying web and cloud and traditional windows applications through a single portal. It also provides users an easy way to exchange files and documents by integrating data services. The integration of Workspace enables you to provide a complete suite of services with the flexibility to deliver to any device form factor. If the ideal end user experience is a View Desktop you can entitle users through workspace. If the ideal end user experience is direct access to web or cloud applications you can entitle these services.

Considerations for Deploying View, Mirage and Horizon

Each component of the Suite is deployed very differently. Horizon Workspace comes as a vApp made up of 5 virtual appliances while View and Mirage are more traditional installations. It is important to understand some of the considerations when deploying each

VMware Horizon View

When deploying View there are many considerations, however one of the first things to understand before getting into the design is what is the business strategy for the virtual desktops as shown in Figure 1.2? Is it to replace physical desktops? Is it part of a Bring Your Own Device initiative? Understanding the intended use allows you to properly plan out the integration of virtual desktops in your environment. Once the use case is understood it is important to understand what services are required by the end users as this will greatly influence your design.

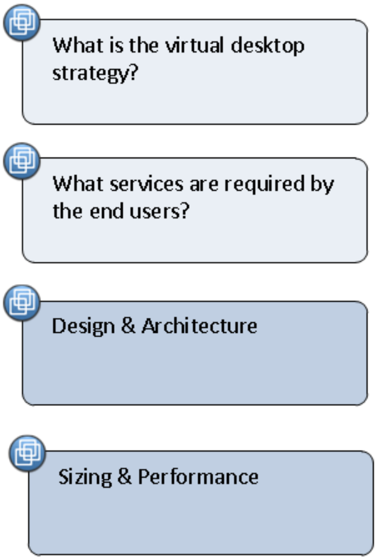
What is the virtual desktop strategy?
What services are required by the end users?
Design & Architecture
Sizing & Performance

Figure 1.2 Considerations

How the use case can influence design is perhaps best illustrated with an example. In the first example company XYZ runs several manufacturing plants in which computers are deployed throughout the shop floor and are used by various personnel throughout the day in 24 hour shifts. The primary application is used for text input in order to control various aspects of the manufacturing process. No information is stored on the computer.

The company has decided to use tablets to remove the need for fixed computer stations and to allow the operators to inspect the adjustments they are making within the computer program against what happening on the line. A VMware View environment will be used to centralized the desktop and enable access through the computer tablets.

In the second example company ABC is an engineering design company that designs large industrial flow and control pump replacement parts. The company would like to use VMware View to provide more flexibility to their engineers to design irrespective of where they are located. In addition they want to enable access to the engineer’s desktops while they are consulting and reviewing aspects of the design with ABC’s customers.

You can see clearly in these examples why knowing the use case upfront can have a large impact on the architecture and the design. For company XYZ the end user would likely be classified as a light user so resources might not be my primary concern but View compatibility with the tablet device likely would be. For company ABC, 3D rendering can often pose problems if not considered carefully in the design. In addition engineers are likely to require a significant amount of CPU and memory allocated to their View desktop.

VMware has made significant strides in improving graphics capabilities in the platform and now offers three distinct types of video support you plan for requirements like company ABC.

[lb] **Soft 3D and SVGA** – Software 3D Renderer and Super Video Graphics Array are provided through the VMware display driver. The Soft 3D and SVGA support is installed when you install the VMware Tools within the virtual machine. Because it has no hardware dependencies, services such as VMotion and Distributed Resource Services (DRS) or the automation of VMotion are fully supported.

[lb] **vDGA** – Virtual Direct Graphics Acceleration is a PCI pass-through to an underlying graphics card. It allows you to pass-through the graphic processing of a virtual machine to an underlying physical graphics card and dedicates a Graphic Processing Unit to the VM. The relationship between the VM and GPU is 1 to 1. As it is a 1-to-1 relationship VMotions and DRS are not supported. It is a property of the VM however so to VMotion a VM you can disable vDGA temporatily.

[lb] **vSGA** – virtual Shared Graphics Acceleration allows you to pass-through the graphic processing of a virtual machine to an underlying physical graphics card that has multiple Graphic Processing Units or CPUs shared amongst VMs. The amount of video memory that can be assigned is restricted to 512 MB however it does allow you to address the engineering requirements of ABC from within a virtual desktop. The configuration of vSGA is very flexible and can be set to Software, Hardware or Automatic. The Software setting only uses vSphere software 3D Rendering even if physical GPUs are available in the host. VMotion and DRS are supported using this setting. Hardware is the opposite and forces the requirement of a GPU in the vSphere host. If one is not present then a VMotion will be restricted and the VM may not be able to be powered on. Automatic is the option in the middle which allows you to switch from vSphere software 3D Rendering to physical GPUs based on availability.

Once you know the business strategy it is easier to deduce which features are required for the end users. As I alluded to in the last paragraph the business case drives the use case and the use case determines which features of View will be most critical to the users. An understanding of the features required also helps influence the design. For example in company XYZ stateless desktops (the assignment of a user to a desktop and any customizations are not preserved between reboots) are likely ideal which means View Composer will likely be a key component of the architecture. View Composer allows a single image to be represented to many different users without requiring a full clone of the parent image for each virtual desktop. A single 20 GB desktop image can be shared to multiple virtual machines (VMs) while appearing to be an independent desktop OS to each user through the use of a linked clone tree.

In addition it is likely that View Blast would meet the requirement. View Blast is the integration of HTML5 support that was added in View 5.2 SP1. It allows the desktop to be delivered over a web browser that supports HTML5 without the requiring the installation of a client. VMware recommends it for users that do not spend a significant amount of time interacting with the View desktop. For users that do spend hours working on their desktop then PCoIP will provide the most robust high fidelity experience.

For company ABC if each engineer and designer has a specific set of tools then a state full desktop (the relationship between the end user and assigned View desktop is preserved along with any user changes) that is associated to each individual user is more appropriate. Having a large number of state full desktops also heavily influences the design and architecture.

Design and architecture is influenced by the business case, end user requirements and also the scale of the environment. If you are designing an environment that will scale to deliver 1000’s of View instances then how you structure the View environment is very important. If I am scaling a View environment from 1000 to 5000 to 10,000 should I just keep adding capacity as shown in Figure 1.3?

Deleted: f

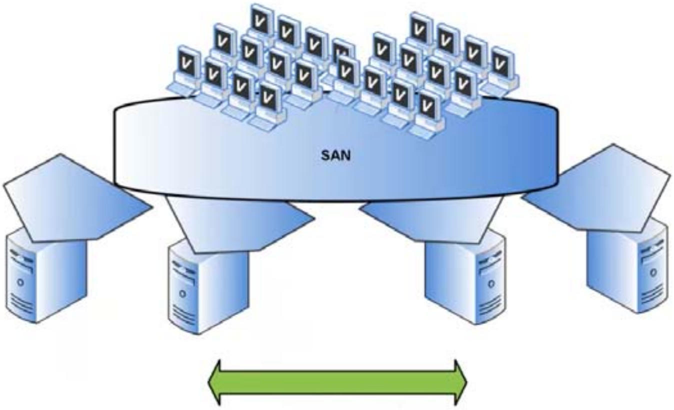

Figure 1.3 Scale out

VMware does not recommend this approach but instead recommends deploying a cluster of View Connection servers in a modular fashion or a 'View Pod' as shown in Figure 1.4. A Pod is a cluster of View Connection servers that replicate metadata through the Active Directory Lightweight Directory Service (AD LDS). Each View Connection server must be connected by a LAN and a maximum of 7 Connection Servers is supported within a cluster or Pod. Each View Connection Server supports 2000 concurrent connections for a theoretical maximum of 14,000 however it is recommended that you do not exceed 10,000 concurrent connections per Pod. A Pod would be deployed on a separate VMware vSphere Cluster that includes not just the Connection servers but also other View servers such as the Composer, Security and Transfer Servers (We will discuss these later in the book in greater detail), This cluster is generally referred to as a Management Block.

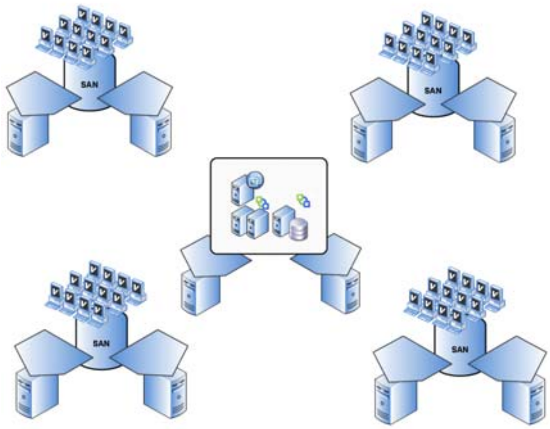

Figure 1.4 A Management Block controlling View Blocks

A Pod would control View desktops deployed in a vSphere environment managed by a separate vCenter server more commonly known as a View Block. When scaling to 10,000 View desktops multiple vCenter View Blocks are recommended. Using multiple View Blocks allows you to redistribute operations that can be resource intensive like redeploying View desktops in order to bring the OS back to a clean state (Operationally this is known as a View Refresh for additional details please see Chapter REFERENCE TO BE ADDED). View Blocks are resource pools controlled by management blocks. This modular approach allows you to scale in a predictive fashion. Each management block is designed to control 10,000 concurrent connections, so if each View Block is designed to support 2500 users then you know that you will need one Management Block for every four View Blocks.

Sizing and performance goes hand-in-hand with your design and architecture. There are several aspects of sizing and performance from the CPU and Memory required per View Desktop to the calculation of disk and network IO. Assuming users are migrating from physical desktops, there is no better way to calculate resource requirements for a virtual desktop than running a proper performance and capacity analyzer tool against the current physical desktops. This is critical if you are designing an environment that will scale to thousands of View desktops.

There are some general guidelines which you can apply in sizing View desktops. Several common considerations are how many vCPUs should be assigned, how much memory, and should I use an x86 or 64-bit OS? Multiple vCPU's only benefit multithreaded applications. The reality is however that most modern Windows desktop operating systems and applications support multithreading. For example multithreading is supported in Microsoft Office 2010. As multiple vCPUs add overhead you generally should not add more than one unless utilization is over 60% of a single vCPU.

The question of whether to deploy an x86 vs. 64-bit desktop OS is tied to the memory allocation. If your View desktop does not require more than 4 GB of memory an x86 desktop is likely to meet your requirement. If it requires more than 4 GB then a 64-bit desktop OS is required. For example, according to Microsoft the Windows 7 64-bit Operating System has a memory limitation of 192 GB.

Network requirements vary considerably depending on what services are enabled within the View desktop. It is therefore important to understand the user requirements so you can estimate your bandwidth requirements. As an example the basic PCoIP (PC over IP) protocol requirements are approximately 250 Kbps per session. Offering high resolution video can require an additional 4,096 Kbps per session. Table 1.1 lists a sampling of the PCoIP bandwidth requirements.

Table 1.1 Sample PCoIP Bandwidth Requirements

PCoIP base requirements	250	Kbps
Multimedia video	1,024	Kbps
3D graphics	10,240	Kbps
480p video	1,024	Kbps
1080p video	4,096	Kbps
Bidirectional audio	500	Kbps
USB peripherals	500	Kbps
Sterio audio	500	Kbps
CD quality audio	2,048	Kbps

Storage sizing is a bit of a science in virtual desktop environments. This is because virtual desktop environments have different IO characteristics depending on the operational activity or state of the desktop. For example if 100 virtual desktops are simultaneously powering on then a burst of IO activity is occurring. If the virtual desktops are powered on and users are logged in, this is considered normal or operational IO activity. To properly size the environment from a storage perspective you need to understand both these IO properties (burst and operational) and also the capacity of size of storage required. While this may seem straightforward size and IO requirements may be a stark contrast of one another. For example View Composer thinly provisions storage requirements so from a capacity perspective I require less storage. As View Composer uses a linked clone tree the Storage IO requirements may be very high.

Storage vendors treat burst IO and capacity and operational requirements distinctly from a sizing perspective as they are typically pinned to different storage tiers. Burst IO requirements are often placed on SSD or flash drives as these technologies can deliver a tremendous amount of IOPS. As they tend to be higher in cost general storage or capacity requirements are usually place on SATA or SAS disks. To determine the burst IO it is common to follow some general guidelines:

[lb] Determine the high watermark for IOPS per View Desktop

[lb] Take the total IOPS and separate them as a percentage of Reads vs. Writes

[lb] Factor the performance penalty on writes by 4

NOTE:

You can aggregate these further for example I have seen IO calculators that consider different states of the virtual desktop (powering on, idle, low utilization etc.) in order to factor a more accurate IO number.

To provide an example let’s look at Figure 1.5. We have estimated 25 IOPS per View Desktop with 60% reads and 40% writes. We have 100 View Desktops to deploy. We calculate the reads at 15 X 100 (25 X 60%) to give us 1500 IOPS. We then add the write IOPS with the penalty; (10 X 1000) X 4) to give us 4000 IOPS. Our expected burst IO is estimated to be 5500 IOPS.

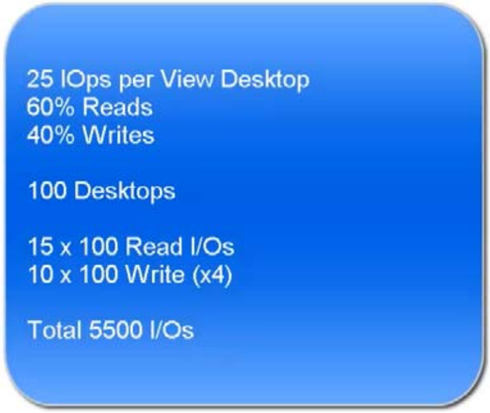

25 IOps per View Desktop
60% Reads
40% Writes

100 Desktops

15 x 100 Read I/Os
10 x 100 Write (x4)

Total 5500 I/Os

Figure 1.5 IOPS Example

VMware Horizon Mirage

Horizon Mirage is based on a Distributed Desktop Virtualization (DDV) architecture which makes use of physical endpoints in order to run a virtual desktop. It is designed to address remote and mobile workers. In a simple architecture as shown in Figure 1.6 the desktop image is stored in the datacenter on a Mirage Server. This desktop image is referred to as a Centralized Virtual Desktop (CVD). The endpoint runs the Mirage client which caches the CVD and locally and ensures any changes are synchronized.

A Mirage desktop is not a single layer but is actually made up of five layers that are treated distinctly. What makes Mirage so powerful is that it will only synchronize the changes to the datacenter made within a layer not the entire image making it extremely efficient on the network. In addition it is possible to update individual layers centrally and have those changes pushed down to the endpoint making it also a great migration tool. When you change a layer it does require a reboot however in upcoming releases, VMware is looking at which layers can be dynamically changed without rebooting. Currently Mirage supports migrations between Windows XP and Windows 7 only although Windows 8 is on the roadmap.

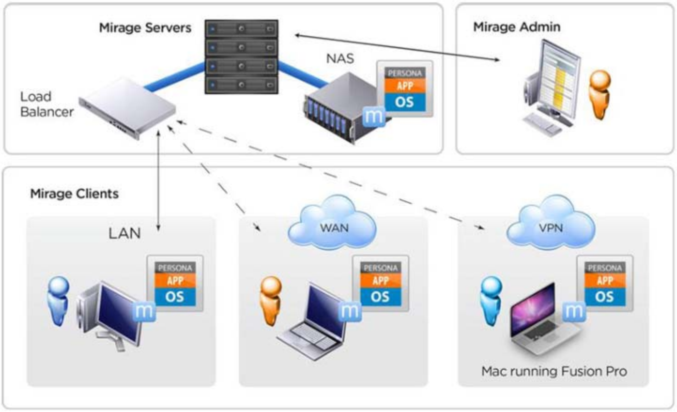

Figure 1.6 Simple Mirage Architecture

The use cases for Horizon Mirage are diverse and include

1. Centralized Image Management
2. Centralized desktop data backup
3. Migration from Windows XP to Windows 7
4. As an enhancement to disaster recovery to provide endpoints
5. System provisioning of desktops

When deploying Mirage there are several things to consider. The initial delivery of the CVD is predicated on the availability of the network. Unlike View that streams the display of the desktop using PCoIP, Mirage delivers the image over the network. Although designed to be highly efficient on the network there are some key considerations when deploying Mirage in large environments. In order to avoid each endpoint downloading the CVD you can make use of a Branch Reflector. A Branch Reflector acts as a proxy for the downloading and synchronization of CVDs and endpoints. Using a Branch Reflector can reduce the bandwidth requirements during mass deployments.

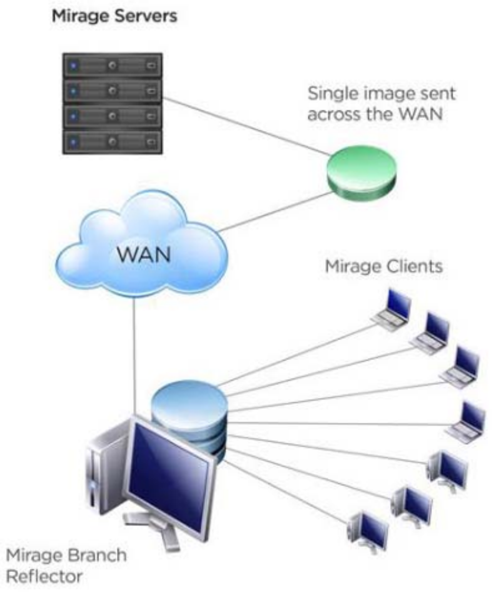

Figure 1.7 Branch Reflector

Another key consideration is the availability of your Mirage Servers. It is recommended that Mirage Servers be load balanced to ensure that one is always available to service the environment.

The deployment considerations for Mirage require a proper understanding of how much bandwidth is required for downloading and uploading between the Mirage clients and Servers, and how to mitigate these requirements by properly placing Branch Reflectors in the environment. In addition, load balancing should be incorporated to ensure the availability of the Mirage Servers.

VMware Horizon Workspace

Horizon Workspace is like a universal aggregator for a variety of end user services. The 1.0 release of Horizon Workspace was released on March 4, 2013. VMware acquired a virtualization technology for Android and iOS based phones from Trango in October 2008. Trango developed a Mobile Virtual Platform (MVP) for phones and that was eventually released by VMware as Horizon Mobile. Horizon Mobile is designed to deliver Mobile Access Management (MAM) for business applications to smart phones.

In actual fact MVP is only used for Android phones while App wrapping is used for iOS applications. App wrapping imbeds a policy within an iPhone App to enforce security boundaries. Horizon Workspace 1.5 consolidates Horizon Workspace 1.0 and Horizon Mobile into a single universal broker for delivering applications to PCs, Tablets and Smart phones. In addition the Horizon Mobile API is likely to take advantage of new Mobile Device Management features of Applie iOS 7.

While we will discuss Horizon Workspace deployment in Chapter 5 “Implementing Horizon Workspace”, one of the key considerations is what services you will aggregate initially. With Horizon Workspace you can integrate View. You can also deploy Data Service which was formerly known as project Octopus and is similar in nature to Dropbox only designed for enterprise customers. You can also provide access to any number of third party Web or Cloud based applications as well as access to Smart phones. In addition for Windows clients you can stream application virtualization packages using VMware ThinApp. All these options are available as different Modules in workspace as shown in Figure 1.8.

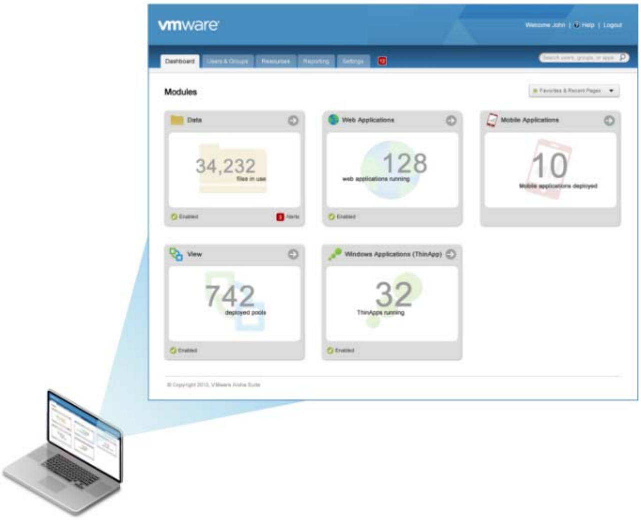

Figure 1.8 Horizon Workspace Modules

An important consideration is understanding your end-user requirements in order to deploy modules that will meet a specific business requirement. Also, you need to verify with the phone manufacturer whether the Smart phone model is VMware-enabled for the MAM component of Horizon; this is only applicable if it is based on Android. VMware has several agreements in place with Verizon and Samsung in the US, however you should verify that the make and model of the Android phone is on the Hardware Compatibility List (HCL) as part of your deployment planning.

Each module within Horizon Workspace has its own deployment and architecture considerations. For example with Data services you need to understand how much storage you will provide for each user in order to determine how much is allocated to Data services.

It is better to have a short list of modules that you plan on enabling initially and then to bring additional modules or service online as required. This actually is generally true of Horizon View as well. Many a virtual desktop project has floundered because too much emphasis was put on deploying features vs. deploying aspects of the technology that address core business requirements. The nice thing about View and Workspace is it is easy to extend the architecture as additional business requirements are identified.

An Introduction to VMware Horizon View, Mirage and Workspace

Although we have discussed each product at a high level and reviewed some deployment considerations, we have not looked at the basic building blocks of each. In this section we will have a look at each product and its base level components.

VMware Horizon View

VMware Horizon View is made up four major server roles; View Connection Server, View Replica Server, View Security Server, View Transfer Server and View Composer as shown in Figure 1.9.

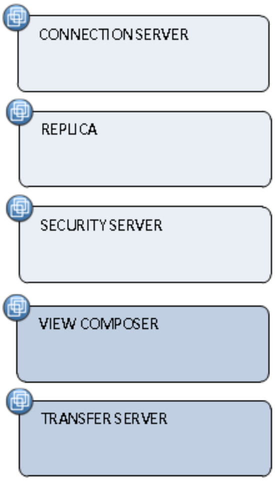

Figure 1.9 View Servers

The View Connection Server fulfills the role of the traditional connection broker which is a core piece of all virtual desktop environments. The Connection Server is to a View environment what vCenter is to a vSphere environment. In other words it is the one stop shop for management, maintenance, configuration and administration of the environment. As it plays such a key role it is a good idea to install not one but two Connections Servers. The second Connection server is referred to a Replica as it shares the metadata that is stored in the Active Directory Lightweight Directory Service (AD LDS). AD LDS provides Active Directory services for AD aware applications without the overhead of AD domains and forests. It is designed purely to replicate application information between servers. Both the primary Connection server and Replica leverage AD LDS to ensure they are in sync from a configuration perspective. What is unique with View is that neither the Connection of Replica requires a database to synchronize information between them. Database services are used for other services such as the Event and Composer database.

The View Security server provides a secure gateway service and is typically deployed within a DMZ as shown in Figure 1.10. It enables the entire View environment to be presented securely through HTTPs over port 443 reducing the number of ports that are opened in the forward facing firewall. It essentially acts as a client proxy.

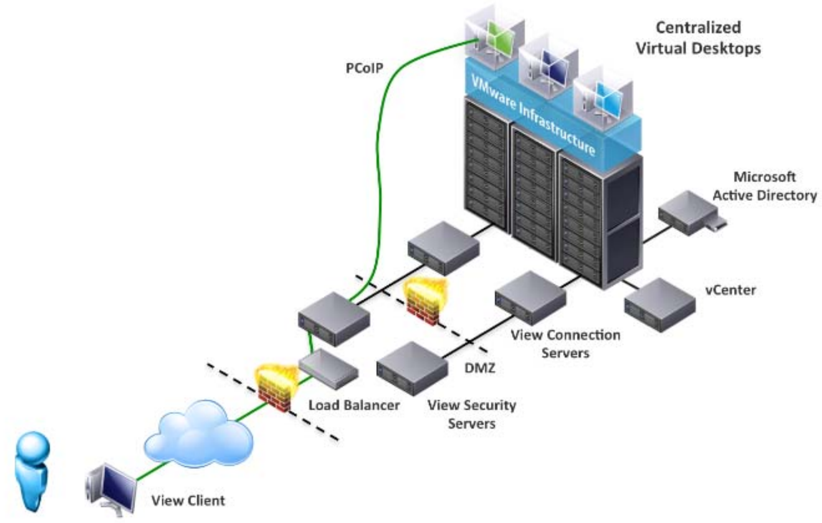

Figure 1.10 Security Server

The View Transfer server enables the Local Mode Client and offloads the checking out or copy of the View desktop and synchronization of the changes from the Connection Server to a dedicated server in the environment as shown if Figure 1.11.

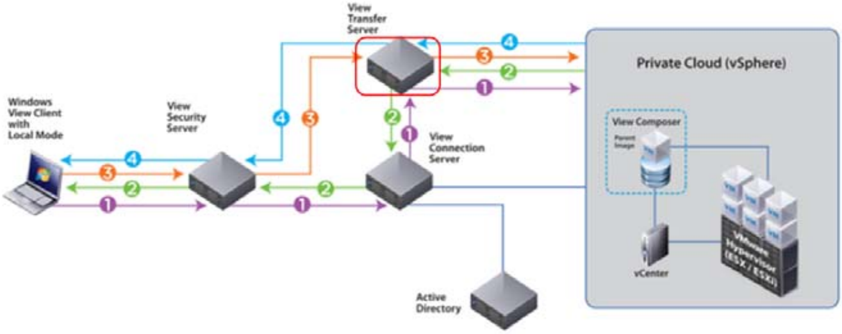

Figure 1.11 Transfer Server

View Composer enables the deployment of linked clones within a View environment. It does require a database to keep track of the connections and components of the service. As View 5.1 Composer can be a dedicated server vs. a service that runs on vCenter as shown in Figure 1.12. This enables Composer to scale much better than prior releases.

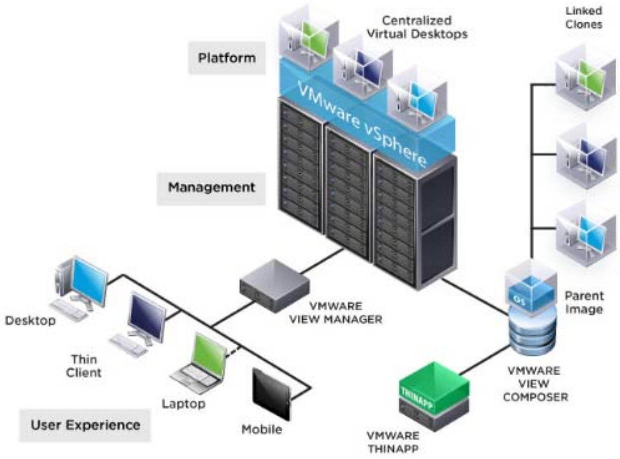

Although there are other key components of View such as the end user data management service, View Persona we have a chapter that provides enhanced details. In chapter one it is important that you get a high level overview of the key pieces of View architecture.

VMware Horizon Mirage

Horizon Mirage is designed to take advantage of resources on the endpoint and provides local execution of the virtual desktop image as shown in Figure 1.13.

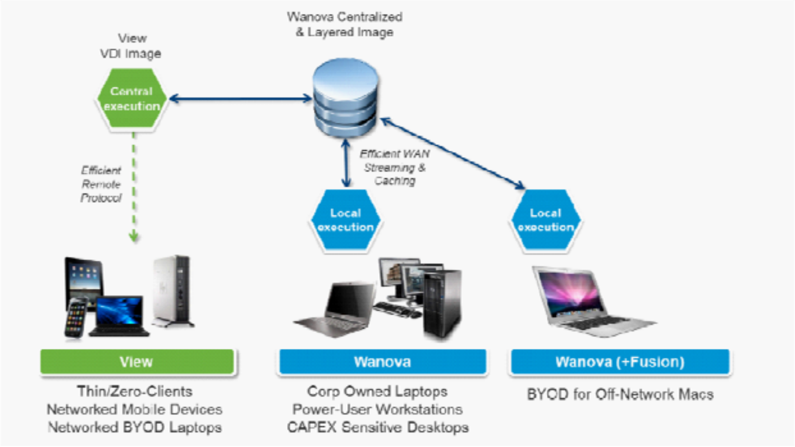

Figure 1.13 Horizon Mirage

There are several key pieces of the architecture that are required to make this work efficiently. At a high level they are the Mirage Server, Centralized Virtual Desktop of CVD, the mirage client and the Branch Reflector as shown if Figure 1.14.

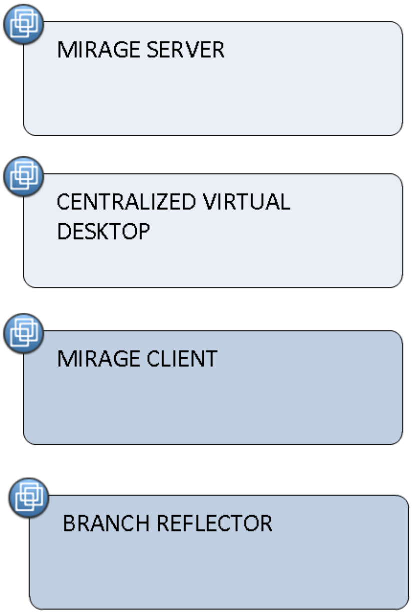

Figure 1.14 Horizon Mirage Components

The Mirage server resides in the datacenter. It provides storage and management of the CVDs. Multiple Mirage servers can be clustered together to ensure availability of the service. It is important the Mirage server is dedicated to this role and it is not recommend that it provide any other services.

The CVD is a layered desktop image that includes 5 layers; a Base Layer which generally includes the OS and core applications such as Anti-virus software. In addition there is an App Layer which is a department or user group specific applications. A Driver Profile which is a repository of drivers designed to be used with specific hardware platforms and a Customized layer which includes machine state information such as the hostname and unique identifier for the desktop. There is also a User settings and data layer which encompasses any changes made by the end user.

The Mirage Client is installed on the endpoint device and enables the execution of a CVD. In addition to running a CVD it can also convert an existing desktop image to a CVD so it used to migrate images to the datacenter as well.

The Branch Reflector optimizes the downloads of the CVD to avoid each Mirage Client downloading the CVD directly or synchronizing all the changes directly. The Branch Reflector proxies the upload and download requests to reduce the number of client connections. In a remote site rather the Mirage Clients connect to the Branch Reflector on the local LAN and the Branch Reflector connects to the Mirage Servers as shown in Figure 1.15.

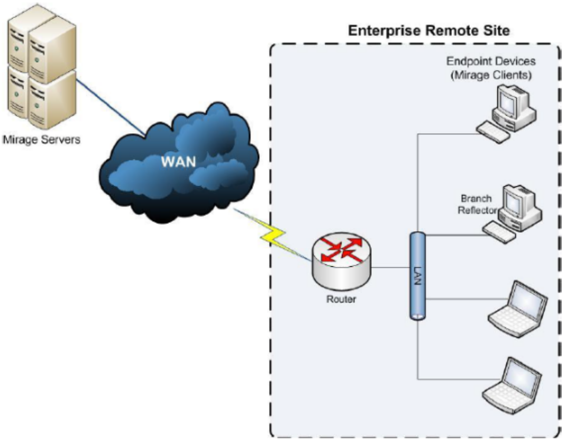

Figure 1.15 Branch Reflector

We break down these high-level components further in Chapter 9 *VMware Horizon Mirage*.

VMware Horizon Workspace

Horizon Workspace architecture is made up of a series of virtual appliances that provide various services to the Workspace environment as shown in Figure 1.16.

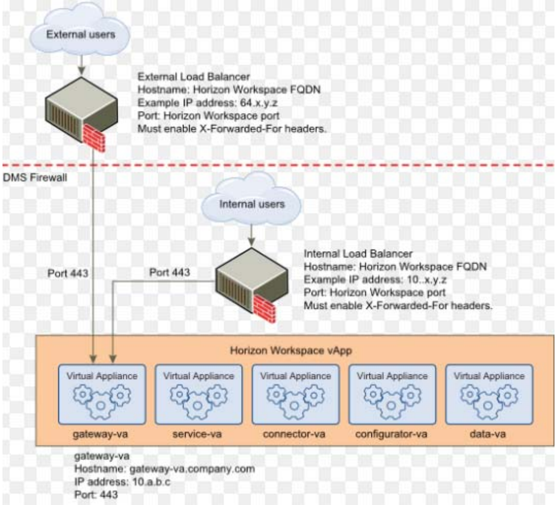

Figure 1.16 Horizon Workspace

The six virtual appliances (VA) provide the services listed in Table 1.2:

NOTE

A virtual appliance the Horizon Mobile Manager, was added when Horizon Mobile was integrated in 1.5

Table 1.2 Horizon Workspace vApp virtual appliances

Virtual appliance	Description
Configurator-va	The configurator VA is used to set and configure and push the configurations to the other vs
Service-va	The service VA provides ThinApp synchronization and user and group management through a web interface
Connector-va	The connector VA provides synchronization of directories, View Pools and ThinApp catalogs along with user authentication
Data-va	The data VA controls the file storage and sharing service

	(formerly project Octopus)
Gateway-va	Is the gateway appliance routes end user connections to the appropriate back end Workspace virtual appliance
Horizon Mobile Management	HMM controls, manages and allows you to configure corporate workspace on a user's smart phone. Rather than enforced management model like in a Mobile Device Management (MDM) Platform, access is managed allowing the user to self-select the corporate applications. This is generally referred to as Mobile Access Management (MAM).

To access Horizon Workspace you need a Horizon Web client. VMware makes a Windows, Android, iOS and Mac client. In addition to these standard clients there is also a Web Client for agentless installations.

Summary

In this chapter we introduced the core pieces of the Horizon Suite: View, Mirage and Workspace. In addition we reviewed some high level pre-deployment considerations and the base level components of each solution. In the following chapter we will drill down into the architecture and outline how each is installed and configured. In addition we will discuss operational aspects of each technology. We also discussed the importance of looking at your next View project as an opportunity to build proper services for your end users that are designed around a service catalog. A service catalog allows users to consume these services as required without the need for IT to deploy each individually. Instead IT can simply entitle the user to access the service.

As Horizon Mirage is a new acquisition to the VMware Horizon product line we outlined the high-level differences and use cases in which Mirage might be more appropriate then centralized View desktops. At the end of this book you will feel comfortable architecting, deploying and managing the solutions in the Horizon Suite in concert with each other. More importantly you will be able to provide a wide variety of business services that are easy for the users to consume and simple for you to manage.

vmware PRESS

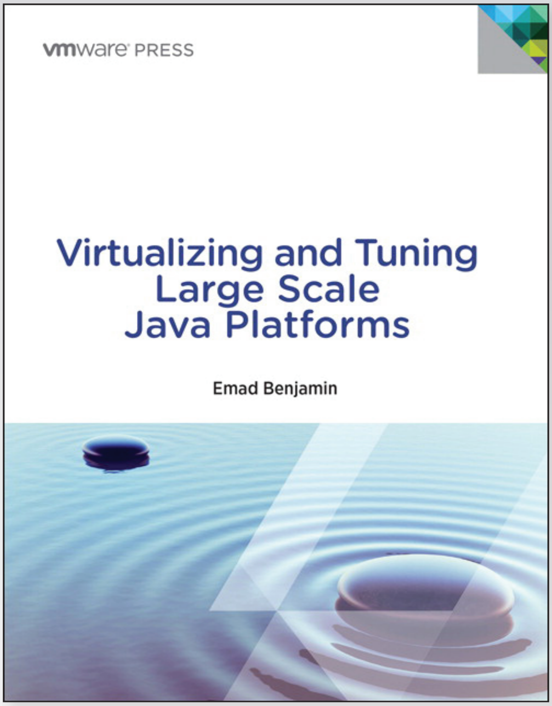

vmware PRESS

Virtualizing and Tuning Large Scale Java Platforms

Emad Benjamin

CHAPTER 1

Introduction to Large Scale Java Platforms

AVAILABLE – JANUARY/FEBRUARY 2014

vmwarepress.com

Chapter 1

Introduction to Large Scale Java Platforms

In this chapter, we will define large scale Java platforms as three different types of categories:

- **Category 1**: Large number of Java Virtual Machines (JVMs) (100s-1000s of JVMs)
- **Category 2**: Smaller number of JVMs with large heap sizes
- **Category 3**: A combination of category 1 consuming data from category 2.

In addition to this various trends and technical considerations are outlined in order to help you understand the range of technical issues associated with designing large scale Java platforms.

Large Scale Java Platform Categories

There are essentially three categories of large scale Java platforms. Although this is the definition I have come across as most suitable based on the many field interactions we have had with customers. The three categories are as follows:

- **Category 1[md]large number of JVMs category**: In this first category there are commonly 100s-1000s of JVMs deployed on the Java platform, and these are typically JVMs as part of a system that maybe servicing millions of users. I have seen some customers with as many as 15000 JVMs. Whenever you are dealing with thousands of JVM instances; you have to consider the manageability cost and if there are opportunities to consolidate the JVM instances.
- **Category 2[md]JVMs with large heap size**: In this category there are almost always fewer (1[nd]20 JVMs); however, the individual JVM heap size is quite large, within a range of 8GB[nd]256GB and potentially higher. These are typically JVMs that have an in memory databases deployed on them. In this category certainly Java Garbage Collection (GC) tuning becomes critical, and we will delve deeper into this in later chapters.
- **Category 3**: The third category is a combination of the first two categories, where perhaps 1000s of JVMs running enterprise applications that are consuming data from category-2 type of large JVMs in the backend.

Across these three categories I see four key requirement trends when virtualizing and tuning large scale Java platforms:

- Compute resource consolidation
- JVM consolidation
- Elasticity and flexibility
- Performance.

Let's look at each one of these trends in more detail.

Large Scale Java Platform Trends and Requirements

In this section key trends and requirements of large scale Java platforms are briefly discussed. Trends such as compute resource consolidation, JVM instance consolidation, elasticity and flexibility, and performance, are some of the major trends that exist within large scale Java platform migration projects.

Compute Resource Consolidation

Many VMware customers find that their middleware deployments have proliferated and are becoming an administrative challenge associated with higher costs. We see a trend across customers who look to virtualization as a way of reducing the number of server instances. At the same time, customers are taking the consolidation opportunity to rationalize the number of middleware components needed to service a particular load. Middleware components most commonly run within a JVM with an observed scale of 100 to 1000s of JVM instances and provide many opportunities for JVM instance consolidation. Hence, middleware virtualization provides an opportunity to consolidate twice[md]once to consolidate server instances, and, secondly, to consolidate JVM instances. This trend is far-reaching, because every IT shop on the planet is considering the cost savings of consolidation.

One customer in the hospitality sector went through the process of consolidating their server footprint and at the same time consolidated many smaller JVMs that were less than 1GB heap. They consolidated many of these smaller 1GB JVMs into 2 categories, those that were 4GB, and others that were 6GB. They performed the consolidation in such manner that the net total amount of RAM available to the application was equal to the original amount of RAM, but with fewer JVM instances. They did all of this while improving performance and maintaining good SLAs. They also reduced the cost of administration considerably due to the reduced number of JVM instances they had to originally manage, and refined environment that helped easily achieve SLA.

Another customer, in the insurance industry, was able to achieve the same as the above customer, but additionally was able to over-commit CPU in development and QA environments in order to save on third-party software license costs.

JVM Instance Consolidation

On the other hand, sometimes we come across customers that have a legitimate business requirement to maintain one JVM for an application, and/or one JVM per a line of business. In these cases, you cannot really consolidate the JVM instances, as that would cause intermixing of the lifecycle of one application from one line of business with another. However, while such customers don't benefit from eliminating additional JVM instances through JVM consolidation, they do benefit from more fully utilizing the available compute resource on the server hardware, that otherwise would have been underutilized in a non-virtualized environment

Elasticity and Flexibility

It is increasingly common to find applications with seasonal demands. For example, many of our customers run various marketing campaigns that drive seasonal traffic towards their application. With VMware, you can handle this kind of traffic burst, by automatically provisioning new virtual machines and middleware components when needed, and then automatically tear down these VMs when the load subsides.

In addition, the ability to change updating/patching hardware without causing outage is paramount for middleware that supports the cloud era scale and uptime. VMware VMotion gives you the ability to move VMs around without needing to stop applications and/or the VM. This flexibility alone makes virtualization of middleware worthwhile when managing large-scale middleware deployments. One customer in the financial space, handling millions of transactions per day, used VMotion quite often to schedule their hardware upgrades without any time downtime; a process that otherwise would be a costly scheduled downtime to their business.

Performance

Customers often report improved middleware platform performance when virtualizing. Performance improvements are partly due to the updated hardware that customers will typically refresh during a virtualization project. There is also some performance improvement due to the robust VMware hypervisor. The VMware hypervisor has improved considerably in the last few years and in Chapter 5, "Performance Studies," we look details of a few performance studies done to showcase some of the heavy workloads that were tested in a virtualized environment.

Large Scale Java Platform Technical Considerations

There are many technical considerations when designing large scale Java platforms. For example, a good understanding of Java Garbage Collection (GC) and JVM architecture, hardware, and hypervisor architectures are essential to building good large scale Java platforms. At a high level GC, and Non Uniform Memory Architecture (NUMA), and theoretical versus practical memory limits will be discussed. In later chapters a more detailed description will be provided, but it is imperative to start at a high level understanding of the issues surrounding large scale Java platform designs.

Theoretical and Practical Limits of Java Platforms

Figure 1-1 depicts the theoretical and practical sizing limits of Java workloads, where we see that there are critical limits that we need to be cognizant of when sizing JVM workloads.

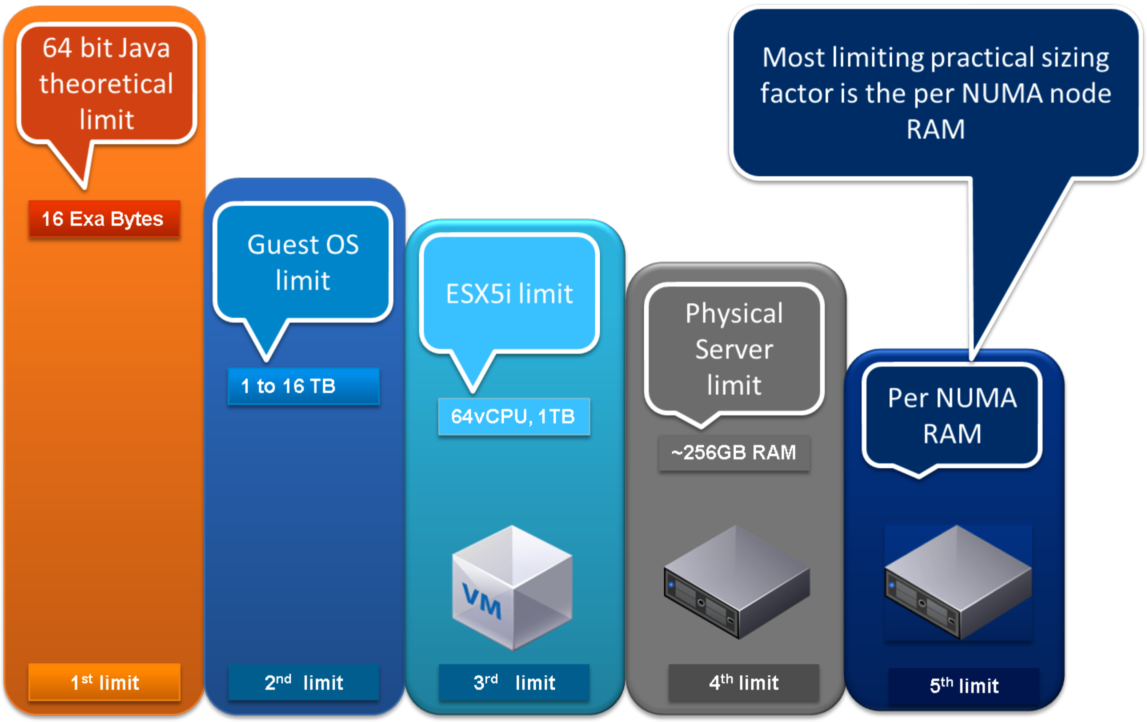

Figure 1-1 Theoretical and Practical Limits of Java Platforms

- It is important to highlight that the JVM theoretical limit is 16Exa Bytes; however, there is no practical system that can provide this amount of memory. Hence, we capture this as the first theoretical limit.
- The second limit is the amount of memory a Guest OS can support; in most practical cases, this is several TBs, and is dependent on the operating system being used.
- The third limit is the ESXi5 1TB RAM per VM, which is ample for any workload that we have encountered with our customers.
- The fourth limit (really the first practical limit) is the amount of RAM that is cost effective on typical ESX Servers. We find on average vSphere hosts have 128-144GB, and at the top end 196GB to 256GB, certainly from a feasibility standpoint the hard limit is probably around 256GB. There are of course larger RAM based vSphere hosts, such as 384GB to 1TB, however these are probably more suited for Category-2 type of in-memory database workloads, and more likely suited for traditional Relational Database management Systems (RDBMS) that would utilize such vast compute resource. The primary reason that RDBMS systems need such large vSphere hosts is because majority (with some minor exceptions such as Oracle RAC) traditional RDBMS don't scale-out and mainly scales up. In the case of Category-1 and Category-2 a scale-out approach is available and hence the ability to select a more cost effective vSphere host configuration is afforded. In Category-1 type of Java workloads you would be best served to consider vSphere hosts with more reasonable RAM range of less than 128GB.
- The fifth limit is the total amount of RAM across the server, and how this is divided into number of NUMA nodes, where each processer socket will have one NUMA node worth of

NUMA-local memory. The NUMA-local memory can be calculated as the total amount of RAM within the server divided by the number of processor sockets. We know that for optimal performance, you should always size a VM within the NUMA node memory boundaries; no doubt ESX has many NUMA optimizations that come into play, but it is always best to stay NUMA local.

If the ESX Host, for example, had 256GB of RAM across 2 processor sockets, i.e. it has 2 NUMA nodes with 128GB (256GB/2) of RAM across each NUMA node. This implies that when you are sizing a VM, it should not exceed the 128GB limit in order for it to be NUMA local.

The limits outlined in the preceding figure and list will help drive your design and sizing decision as to how practical and feasible it is to size large JVMs. However, there are other considerations that come with sizing very large JVMs, such as GC tuning complexity and knowledge needed to maintain large JVMs. In fact, most commonly sized JVMs within our customer base are in the vicinity of 4GB of RAM for the typical enterprise web application, or what has been referenced in this book as Category-1 workloads. On the other hand, larger JVMs exist, and we have customers that run large scale monitoring systems and large distributed data platforms (in memory databases) on JVMs ranging from 4GB to 128GB. This is also true for in memory databases such as vFabric GemFire and SQLFire where individual JVM members within a cluster can be as big as 128GB, and total cluster size of 1 to 3TB. With such large JVMs comes the need to have a better knowledge of GC tuning. At VMware, we have helped many of our customers with their GC tuning activities over the years, even though GC tuning on physical is no different than on virtual. The reason being is that we have uniquely integrated the vFabric Java and vSphere expertise into one spectrum, which has helped our customers run many Java workloads on vSphere optimally. When faced with the decision of whether to vertically scale the size of the JVM and VM, always first consider a horizontal scale out approach[md]as we always find our customers get better scalability with the a horizontally scaled out platform . If horizontal scalability is not feasible, then consider increasing the size of the JVM memory and hence VM memory. When opting to increase the size of the JVM by increasing the heap space/memory, the next point of consideration will be GC tuning and the in house knowledge you have to be able to handle large JVMs.

NOTE On the 3rd limit as of the writing of the book, ESXi 5.1 is the GA released official version; however, it is anticipated that by the time this book is published, it is possible that some of these maximum vSphere limits may change. Double check official VMware product documentation for the latest maximums.. It is worthwhile noting that at these VM limits there is not cost effective hardware that would need such a large number of vCPU, but yet assuring for those that may need it.

Having discussed earlier in the chapter about the three categories within large scale Java platforms that exist in the enterprise today, Figure 1-2 shows the various workload types and relative scale. A common trend is that as the size of the JVM increases so too does the required JVM GC tuning knowledge.

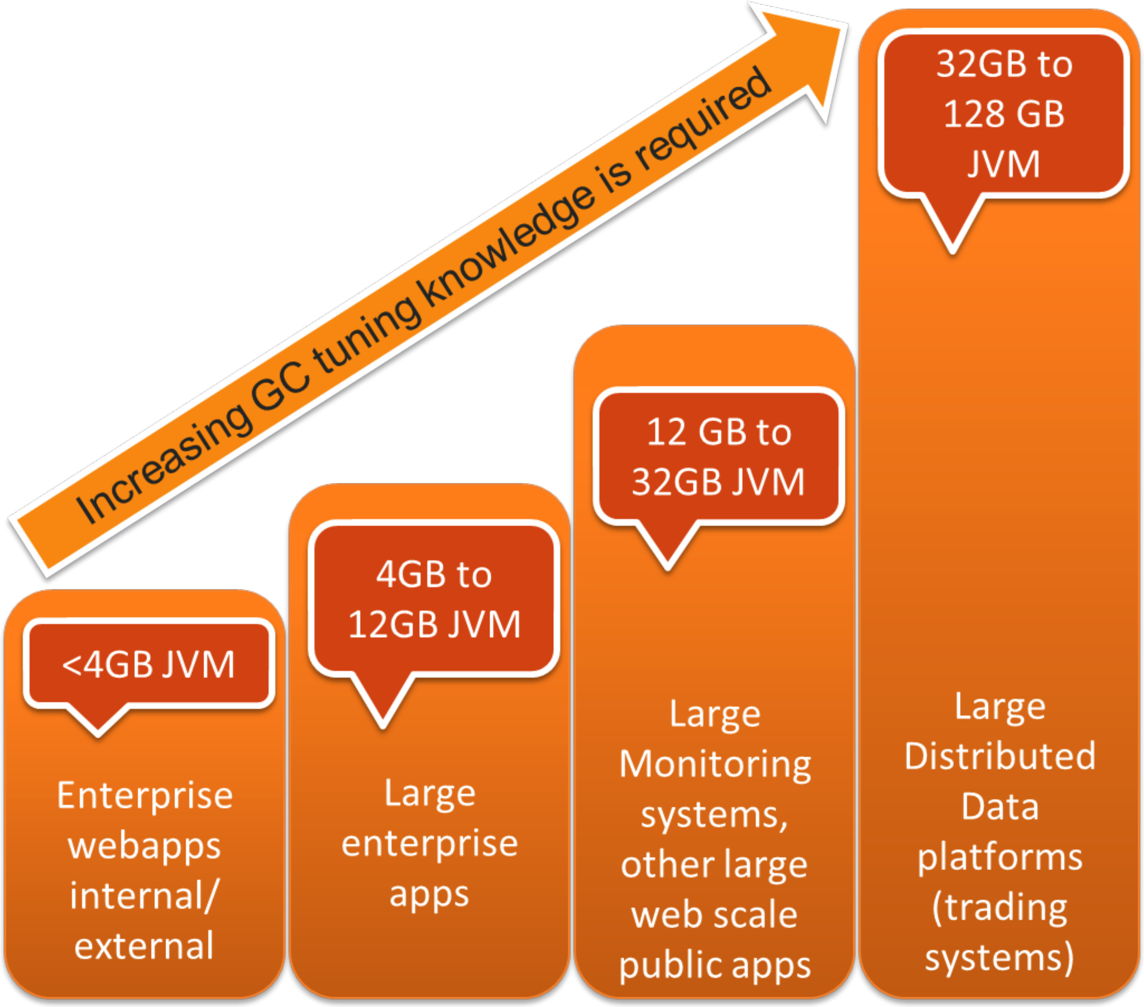
Increasing GC tuning knowledge is required
<4GB JVM
Enterprise webapps internal/ external
4GB to 12GB JVM
Large enterprise apps
12 GB to 32GB JVM
Large Monitoring systems, other large web scale public apps
32GB to 128 GB JVM
Large Distributed Data platforms (trading systems)

Figure 1-2 GC Tuning Knowledge Requirements Increases with Larger JVMs

It is important to keep the following in mind:

- Perhaps in the first case of JVMs less than 4GB in heap size, this the most common in amongst the workloads that exists out there. The 4GB is a special case, as it has the default advantage of using 32-bit address pointers within a 64-bit JVM space, hence very efficient on memory footprint. These require some tuning but not a substantial amount. This workload type falls into the realm of category 1 defined earlier in the chapter. The default GC algorithm on server class machines is that of the throughput garbage collector is adequate. The only time you need tuning is if for some reason the response time measurements are not sufficient and you have to begin a tuning effort. In such cases it is best to follow the guidance on GC tuning later in Chapter 3, “Tuning Large-Scale Java Platforms,” and Chapter 6, “Best Practices.”
- The second workload case, is still within our earlier defined category 1, but it is probably a serious user base internal to the organization. We typically see heavily used enterprise Java web applications on a scale of 1000 to 10,000 users. In these types of environments, GC

tuning and slightly larger than 4GB JVMs are the norm. The DevOps team almost always has decent GC tuning knowledge, and has configured the JVM away from the default GC throughput collector. It is here we start to see the usage of the Concurrent Mark and Sweep (CMS) GC algorithms in order for these type of workloads to deliver decent response times to the user base. The CMS GC algorithm is offered by the Oracle JVM (formerly Sun JVM), for further details and availability of other GC algorithms within Oracle JVM or IBM JVM, refer to Chapter 3 – Tuning Large Scale Java Platforms, and Chapter 6 – Best Practices.

- The third workload type, could be as part of category 2 defined earlier, but a unique case within category 2, since in some cases the larger JVMs are used because the application may not be capable of horizontal scale-out. Now in the generic case of category 2 type of workloads are usually in memory databases, as mentioned earlier in the chapter. In this category no doubt a deep knowledge of JVM GC tuning is required. Your DevOps team needs to be able to articulate all the different GC collectors and select those that are more suitable for improved throughput (Throughput collectors) as opposed latency sensitive workloads that need CMS GC in order to deliver better response times
- The fourth workload type falls into a mix of category 2 and 3, here there could be a large distributed system, where the client enterprise Java applications are consuming data from the backend data fabric where a handful or more of in-memory database JVM nodes are running. Tuning GC at expert level is required here.

Other than having to maintain a very large JVM, you need to be cognizant of the workload choices. We often find that what drives our customers to vertically scale the JVM is usually due to perceived simplicity of deployment and leaving existing JVM process intact. Let’s consider some JVM deployment and usage scenarios, perhaps this is something you have in your environment today, or something you have come across in the past:

- For example, a customer has one JVM process deployed initially, and then as work demand increases for more applications to be deployed, instead of horizontally scaling out, by creating a second JVM and VM, the customer takes a vertical scale up approach. As a consequence, the exiting JVM is forced to vertically scale and carry many different types of workloads with varied requirements.
- Keeping in mind that some workloads, such as, a job scheduler, have a need for high throughput, while a public facing web application has a demand for fast response time. Hence, stacking these types of applications on top of each other, within one JVM, complicates the GC cycle tuning opportunity. We know that when tuning GC for higher throughput, it is usually at the cost of decreased Response Time, and vice versa.
- You can achieve both higher throughput and better response time with GC tuning, but it certainly extends the GC tuning activity unnecessarily. When faced with this deployment choice it is always best to split out the types of Java workloads into their own JVMs. One approach would be to run the job scheduler type of workload in its own JVM and VM, and similarly for the web-based Java application.
- In Figure 1-3, JVM-1 is deployed on a VM that has mixed application workload types, which complicates GC tuning and scalability when attempting to scale up this application mix in JVM-2. A better approach is to split the Web application into JVM-3 and the job scheduler application into JVM-4, i.e. horizontal scaled out and with the flexibility to vertically scale if needed. Certainly if you try and compare the vertical scalability of JVM-3 and JVM-4 vs. vertical scalability of JVM-2 you will find JVM-3 and JVM-4 always scales better and are easier to tune.

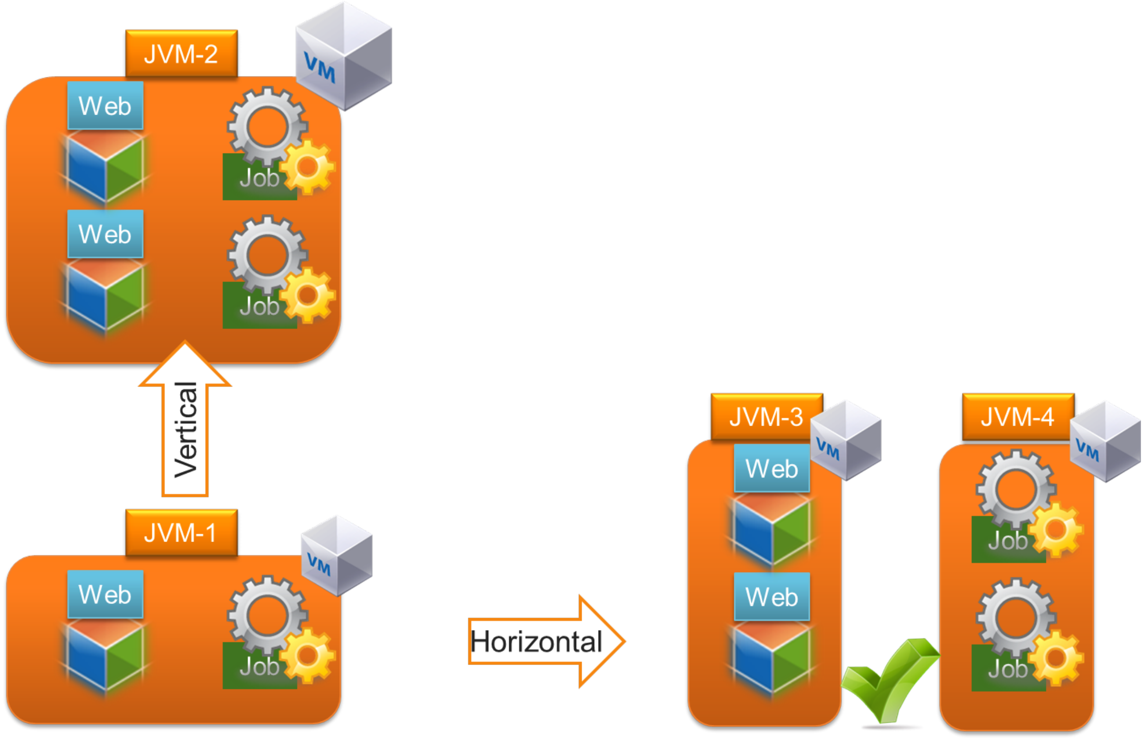

Figure 1-3 Avoiding Mixed Workload Types in the Same JVM

NUMA

NUMA is a computer memory design used in multiprocessors, where the memory access time depends on the memory location relative to a processor. Under NUMA, a processor can access its own local memory faster than non-local memory, that is, memory local to another processor or memory shared between processors.

Understanding NUMA boundaries is critical to sizing VM and JVMs, ideally the VM size should be confined to the NUMA boundaries. Figure 1-4 shows a vSphere host made of two sockets, and hence two NUMA nodes. The workload shown is that of two a vFabric SQLFire VMs, each VM sized to fit within the NUMA node boundaries for memory and CPU. If a VM is sized to exceed the NUMA boundaries, there is potential that it will interleave to the other NUMA node in order to fulfill the request for additional memory that otherwise cannot be fulfilled by the local NUMA node. In the diagram we depict memory interleaving by the red arrows, highlighting that this type of memory interleaving should be avoided as it may severely impact performance.

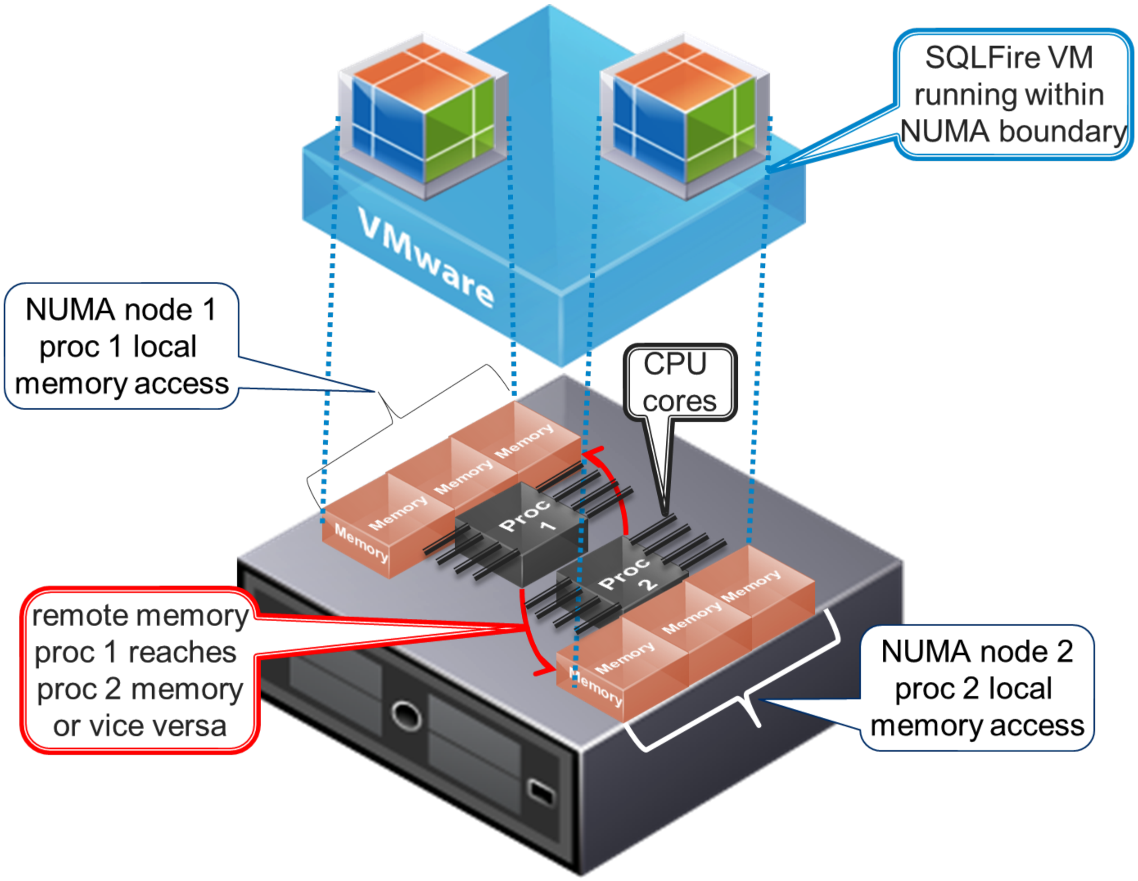
SQLFire VM running within NUMA boundary
VMware
NUMA node 1 proc 1 local memory access
CPU cores
Memory
Memory
Memory
Memory
Proc 1
Proc 2
Memory
Memory
Memory
Memory
remote memory proc 1 reaches proc 2 memory or vice versa
NUMA node 2 proc 2 local memory access

Figure 1-4 Two Socket Eight Core vSphere Host with Two NUMA Nodes, and One VM on Each NUMA Node

In order to calculate the amount of RAM available in each NUMA node you can apply the equation in Formula 1-1

NUMA Local Memory = Total RAM on Server /Number of Sockets

Formula 1-1 Per NUMA Node RAM Size (NUMA Local Memory)

For example, if a server has 128GB of RAM configured on it, and has two sockets (as shown in Figure 1-4), this implies the Per NUMA RAM is 128/2, which equals 64GB. This is not entirely true, however, as ESX overhead needs to be accounted for, and a more accurate approximation would be as per the equation shown in Formula 1-2. The formula accounts for the ESXi memory overhead made of a 1GB as a constant regardless of the size of the server, and a 1% VM memory overhead as 1% of the available memory. The formula is a conservative approximation, and every VM and workload will vary slightly, but the approximation should be pretty close to the worst case scenario.

NUMA Local Memory =

Total RAM on Host- ((Total RAM on Host* nVMs * 0.01)+1GB) /Number of Sockets

Where:

NUMA Local Memory = Local NUMA memory for best memeory throughput an locality with VM and ESXi overhead already accounted for

Total RAM on Host = this is the amount of physical RAM configured on the physical server

nVMs = number of VMs you plan to deploy on the vSphere host

1GB = this is the overhead needed to run ESXi

Number of Sockets = the numbers of sockets available on the physical server, 2 socket or 4 socket

NOTE: the above formula take the most pessimistic end of the overhead range, especially as you increase the number of VMs, clearly as you add more VMs the more overhead. As opposed to lower number of VMs the approximation of formula 1-2 is pretty fair and accurate. Also this assumes non-overcommitted memory situation.

If you don’t have time to crunch through the formula and wish to quickly start configuring, then assume about 6% of overhead due to memory, now clearly there are many times when not all of this is being used. For example:

Example 1- Using 6% Approximation Approach:

This would imply if you have a server that has 128GB physical RAM (2 socket host, 8 core on each socket) and you chose the 6% overhead approach while configuring 2 VMs on the host, then the total NUMA local memory would be => ((128*0.94)-1)/2 => 59.7 GB per VM available for memory. Sicne there are 2 VMs the total memory offered to the two VMs is approximately 59.7*2 => 119.32GB

However, if you apply the approach in Formula 2-1 as shown in Example-2 below:

Example 2 – Using Formula 2-1 to Calculate NUMA Local Available Memory:-

Again assuming 128GB host with 2 sockets (8 core on each socket), and 2 VMs to be configured on it,

NUMA Local memory = (128-(128*2*0.01)-1)/2 => 124.44GB, and note this is for 2 VMs if you decided to instead configure 16 VMs of 1vCPU (1vCPU=1 core), then the NUMA local memory per VM would be NUMA Local memory = (128-(128*16*0.01)-1)/2 => 53.26GB, this is probably overly conservative and a more accurate representation would be around the 6% overhead calculation approach.

For best guidance the best approximation of overhead is the 6% of total physical RAM (plus 1GB for ESXi) approach shown in Example-1.

In the preceding example where we show a calculation based on a server having 128GB of RAM, the true local memory would be: ((128*0.99) – 1GB)/2 => 62.86GB, which is the maximum VM size that can be configured. In this case, you can safely configure two VMs of 62.68GB of RAM and 8vCPUs each, as each of the VMs would be deployed on one NUMA node. You can also alternately deploy four VMs if you wish to deploy smaller VMs of 62.86GB/2=> 31.43GB of RAM and 4vCPU each, and the NUMA scheduling algorithm would still localize the VMs to the local NUMA node.

NOTE On hyper-threaded systems, virtual machines with a number of vCPUs greater than the number of physical cores in a NUMA node but lower than the number of logical processors (typically logical processors are shown as 2.x of physical cores, but more practically logical processors are 1.25x of physical cores) in each physical NUMA node might benefit from using logical processors with local memory instead of full cores with remote memory. This behavior can be configured for a specific virtual machine with the numa.vcpu.preferHT flag. For further details please refer to, http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf and KB article: kb.vmware.com/kb/2003582

It is always advisable to start with vCPUs equal to the number of physical cores and then adjust vCPUs upwards when needed, but less than approximately 1.25x of available physical cores.

To further elaborate on the ESXi NUMA scheduling algorithm, Figure 1-5 shows an example of a two socket six core on each socket server.

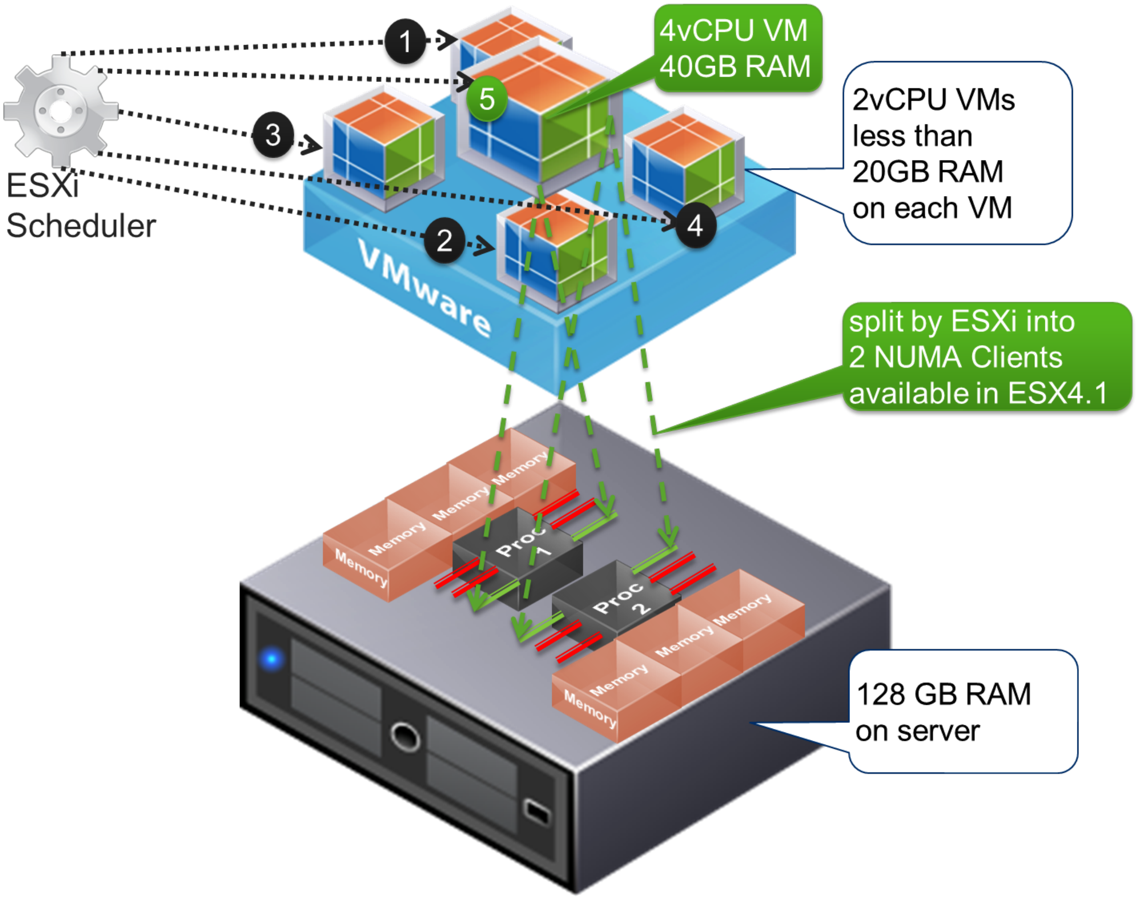

Figure 1-5 ESXi NUMA Scheduling on a Two Socket Six Core Server

On this diagram there are initially four VMs of two vCPU and approximately 20GB RAM on each. The initial ESXi scheduling algorithm will follow a round robin fashion, by step 1 (as shown by the black circle with the number 1), and then the next two vCPU VM will be scheduled on the next available empty NUMA node, and then so on for steps 3 and 4 for scheduling the 3rd and 4th VM. At the point of where all four of the 2vCPU 20GB VMs have been scheduled, and as a result of this scheduling the 4 VMs will occupy the 4 cores on each of the sockets, as shown by red pins on the diagram. Moments later, a 5th VM made of 4vCPU and 40GB RAM is deployed and now ESXi attempts to find space to fit it within one NUMA node. However, because there are not enough CPUs and RAM all within one NUMA node, ESXi scheduler will split the execution of the VM across the two NUMA nodes in such a way as to localize the execution within each NUMA node. This feature has been available since ESXi 4.1, and ESXi will do its best to localize the execution to the NUMA node until such point that it can't and then it will cause memory interleaving of fetching memory from the remote NUMA nodes. ESXi in this case scheduled the 5th VM by allocating 2 CPUs and 20GB from NUMA node 1 and 2 CPUs and 20GB from NUMA node 2. No doubt we want to avoid this type of scenario and ideally size the original 4 VMs to take up the entire compute resources space so that this type of scheduling is avoided.

NOTE In vSphere4.1/ESXi4.1 the underlying physical NUMA architecture is not exposed by the hypervisor to the operating system, and hence application workloads running on such VMs had no ability to take specific advantage of additional NUMA hooks that they may provide. However, ever since vSphere5 the concept of vNUMA has been introduced where through configuration you can expose the underlying NUMA architecture to the OS, and hence NUMA aware application can take advantage of it. In Java the –XX:+USeNUMA JVM option is available however it is only compatible with the Throughput GC, and not the CMS GC. Ironically, in most memeory intensive cases where NUMA is a huge factor, latency sensitivity is a big consideration and hence the CMS collector is more sutiable. This implies you can't use CMS and –XX:+UseNUMA option together. The good news is that vSphere NUMA algorithms are typically good enough to provide locality, especially if you have followed good NUMA sizing best practices – such as sizing VMs to fit within NUMA boundaries for memory and vCPU perspective.

Most Common JVM Size Found in Production Environments

Having discussed thus far the various JVM sizes that you can deploy, and in some cases very large JVMs, it is important to keep in mind that the most common JVMs found in data centers are of 4GB Heap size. This may be a fairly busy JVM with 100 to 250 concurrent threads (actual thread count will vary, as it depends on the nature of the workload); 4GB of heap; approximately 4.5GB for the JVM process; 0.5GB for the Guest OS; hence a total recommended memory reservation for the VM of 5GB with 2vCPU and 1 JVM process, as shown in Figure 1-6.

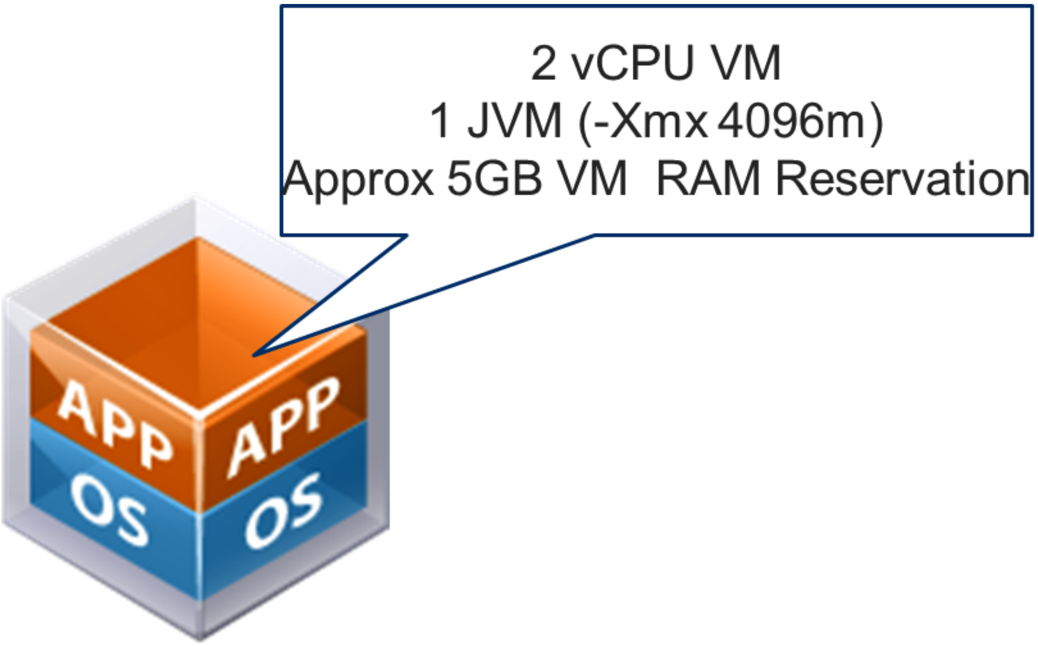

Figure 1-6 Most Common JVM Size Found in Production Environments

Horizontal Scaling versus Vertical Scaling of JVMs and VMs

When faced with horizontal scaling versus vertical choices, three options can be captured as shown in Figure 1-7.

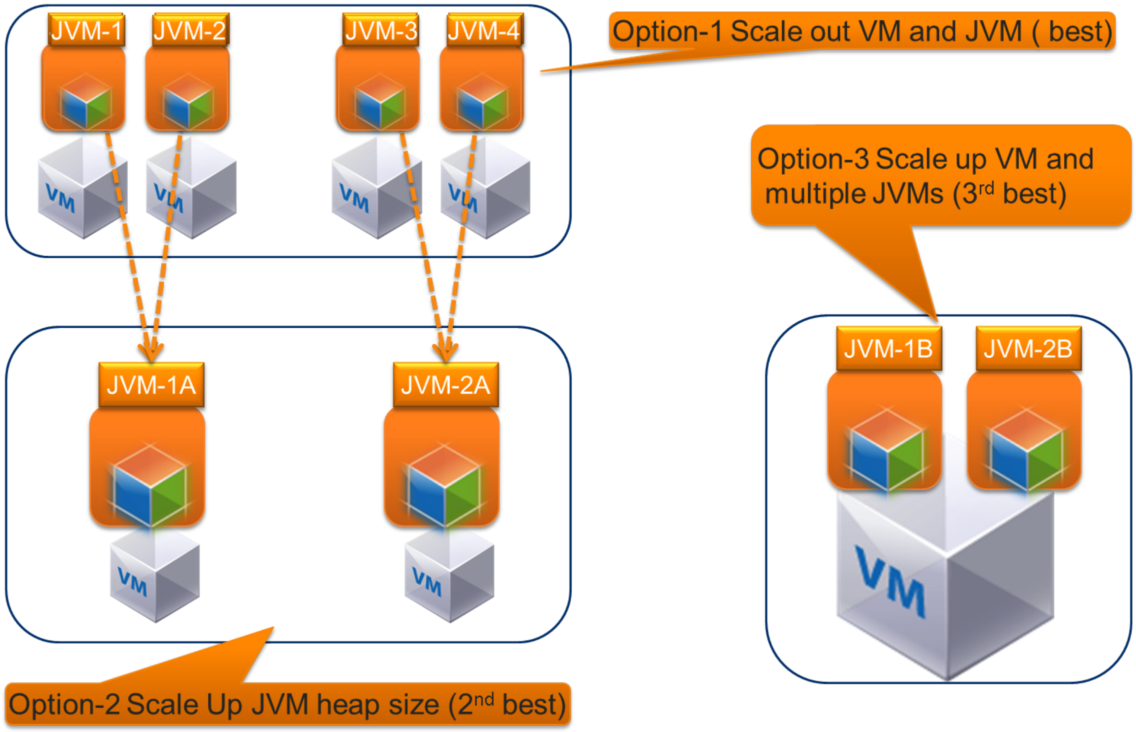
JVM-1
JVM-2
JVM-3
JVM-4
VM
VM
VM
VM
Option-1 Scale out VM and JVM (best)
Option-3 Scale up VM and multiple JVMs (3rd best)
JVM-1A
JVM-2A
VM
VM
JVM-1B
JVM-2B
VM
Option-2 Scale Up JVM heap size (2nd best)

Figure 1-7 Horizontal versus Vertical JVM Scalability Choices

Table 1-1 details the pros and cons of the three options.

Table 1-1 Comparing Pros and Cons of Various Horizontal and Vertical Scaling Options

Option	Pros	Cons

Option 1 JVMs are introduced into the Java platform by creating a new VM and deploying a new JVM on it. Hence a scale out VM and JVM model.	This provides the best scalability since the VM and the JVM are scheduled out as one unit by the ESXi scheduler. It is really the VM that is scheduled by the ESXi, but since there is only one JVM on this VM, then the net effect is the VM and the JVM are scheduled as one unit It also offers the best flexibility in being able to shut down any VM and JVM in isolation without impacting the rest of the Java platform. This is no doubt in relative terms, since most Java platforms are horizontally scalable, and in most cases there are enough instances to service traffic, even though JVM instances are being shut down. The relative comparison in terms of more instances having better scalability is based on when you compare between whether you have 100 JVMs and VMs versus 150 JVMs and VMs for the exactly same system, if for a minute you where comparing and contrasting platform design and you were trying to choose between 100 JVM systems versus 150JVMs, with both cases of 100 and 150 JVMs having the same net RAM. Clearly, the system with the 150 JVMs will have the better flexibility and scalability. In the 150 JVM case it is likely since you have more JVMs the size of the JVM is smaller than compared to a system that has 100 JVMs and in this case if there was a problem with a JVM from the 150 JVM platform, likely the impact is smaller since the JVM holds less data than the 100 JVM case. Hence the scale-out robustness of the 150 JVMs will prove to be more prudent If the system has been refined, i.e. consolidation of VM and JVM best practices have been applied based on a 64-bit JVM architecture having a reasonable size JVM sizes with approximate minimum of 4GB heap space, and not fragmented around a legacy 32-bit JVM limit of 1GB heaps (some legacy 32 bit JVMs could withstand greater than 1GB, but for practical usages 32-bit JVM have a had legacy 1GB limit) space, then the horizontal scalability advantages assumed above are intact	This option is expensive as it leads to having more OS copies and licensing becomes expensive quite quickly Administering such system is more expensive as there are more VMs and JVMs to keep track of There is no technical reason why you need to place 1 JVM on 1 VM, the only exception is in the case of systems that are in memory databases (like category 2) that require high throughput memory from the local NUMA node, in those cases the VMs are sized to fit within the NUMA node and will have only 1 JVM on it. Also important to note that the JVMs in in-memory databases tend to be quite large, sometime as big 128GB, as opposed to category 1 JVM sizes are typically 1 to 4GB heap size JVM range. However, in cases such as in option 1 which are essentially of category 1 (as defined earlier in the chapter, category 1: large number of JVMs category) there are many opportunities to consolidate the JVMs and eliminate wasteful JVMs and VM instances This is a common pattern amongst legacy 32-bit JVMs where the 1GB limit of the 32-bit JVM would have forced Java platform engineers to install more JVM instances in order to deal with increase in traffic. The downside here is that you are paying for additional CPU/licenses, if you consider a consolidation of JVMs by migrating to 64-bit JVM and increasing the heap at the same time, then you will save by having fewer JVMs servicing the same amount of traffic, but of course the JVM size will likely increase form, for example, 1GB to 4GB.

Option 2 Scale Up JVM heap size by consolidating fragmented smaller JVMs and also as a result consolidate VMs	**NOTE:** If you look at option 2 in Figure 1-7, it shows 2 JVMs (JVM-1A and JVM-2A) were consolidated from 4 JVMs (JVM-1, 2, 3 and 4) of option 1. In this process the 4 VMs were also consolidated into 2 VMs as shown by the diagram. For example, if JVM-1, 2, 3 and 4 were all of 2GB heap size each running on VMs of 2vCPU, this implies the total RAM serviced to the heap and in turn to the application is 8GB across all of the JVMs. The total vCPU across all of the VMs is 8vCPU. Now when consolidating down to 2 VMs and 2 JVMs, the JVMs in Option 2 JVM-1A and JVM-2A are each of 4GB heap, for a total of 8GB, and the VMs are 2 vCPU each. This implies total vCPU of 4 across both VMs, a saving of 4vCPU since originally in option 1 there were 4VMs of 2vCPU each. The reason for being able to scale down vCPU while still maintaining equal amount of RAM (Java heap space) is because with larger JVM heap spaces GC is able to scale vertically fairly well without having to excessively consume CPU. No doubt this largely workload behavior dependent, and some workloads may indeed exhibit increased CPU usage when JVMs are scaled-up, but it is true to say that majority of the cases in category-1 type of workloads have exhibited a behavior of releasing the unneeded vCPU when consolidated into larger JVM heap. 64-bit JVMs are highly capable runtime containers, and while there is an initial cost of launching one, they do offer an ability to crunch through massive number of transactions that are within much larger heap spaces. When an intent is made to create a new JVM it should be inspected with the same set of questions as if you were about to create a new VM. If someone needs a new VM, a vSphere administrator would always ask what you need it for. Because a VM is highly capable compute resource, in much the same way the JVM is a highly capable machine, and hence vSphere administrators and DevOps engineers should always inspect closely the validity of having to create a new JVM, as opposed to being able to leverage existing JVM instances and perhaps increase the heap space, within reason, in order to facilitate more traffic. Reduced administration cost due to reduced number of JVMs and VMs Reduced licensing cost due to reduced OS copies Improved response times, likely more transactions are now executed within the same heap spaces as opposed to requiring marshaling across the network	Due to having larger size JVMs if there is no proper redundancy or persistence of transactions, and if a JVM crashes then more data is lost when compared with the case of smaller JVMs in option-1. Due to consolidation you may have fewer HA JVM instances Consolidation is limited to line of business; you can’t mix applications from different lines of business into the same JVM, a crash of the JVM would impact both lines of business. Larger JVMs may require some more GC tuning

Option 3	**NOTE:** If option 1 and option 2 are not possible, then option three could be considered. In this case you are placing multiple JVMs on larger VM. Now JVM-1B and JVM-2B could be JVMs that were consolidated copies, like the ones in option 2, or non-consolidated copies like in option-1. In either case you can stack these JVMs on a larger VM, or multiple large VMs for that matter. If the current platform is similar to that in option 1, it might be an advantage, due to logistical reasons, to keep the current number of JVMs intact in the deployment, but then consider building larger VMs having multiple JVMs stacked on them. Reduced number of OS licenses Reduced number of VM instances Reduced administration cost due to having fewer VMs Ability to have dedicated JVMs to each line of business, but then ability to deploy JVMs from multiple line of business on the same VM. This should only be done if the cost VM consolidation outweighs the dangers in having to impact multiple lines of business during a VM crash. Having large VMs gives the potential of more vCPUs available to JVMs, if for example a VM has 2 large JVMs on it from different lines of business and they peak at different times, it is likely that all of the vCPUs are available to the busy JVM, and then similarly for the next JVM when its peak arrives.	Likely larger VMs are required, scheduling larger VMs may require more tuning than smaller VMs NOTE: that various performance studies have shown that the sweet spot VM size is 2vCPU to 4vCPUs for category-1 type of workloads, but naturally in category-2 workloads larger than 4vCPU are needed. Certainly a starting point of 4vCPU may be needed. But bear in mind the scheduling opportunity form a HA perspective maybe diminished, although in category-2 type of workloads, like in memory data bases, most are fault tolerant, redundant, and disk persistent, and hence may not rely as much on VMware HA and/or automatic DRS Since this option is about trying to consolidate VMs it is highly likely that JVMs from different lines of business may be deployed on the same VM, this needs to be managed correctly, inadvertent restart of a VM will impact multiple lines of business potentially. You can attempt to consolidate JVMs in this case and also stack them up on the same VM, however this forces the JVMs to be much larger in order to fully utilize the underlying memory. If you configure fewer larger VMs, it literally means you have VMs with a lot more RAM form the underlying hardware, and in order to fully consume this it may require larger JVM heap spaces. Having larger JVMs has the potential of losing more data in the event a JVM crashes, especially if the JVM doesn't have adequate redundancy and/or persistence of application data May require large vSphere hosts, and hence larger servers cost more

Chapter Summary

This chapter introduced the concept of large scale Java platforms and categorized it into three categories:

- **Category 1**: Large number of JVMs
- **Category 2**: Smaller number of JVMs with large heap sizes
- **Category 3**: A combination of category 1 and 2.

The chapter also examined the various theoretical and practical limits that exists within the JVM, outlined various workload types and commonly encountered JVM sizes. Finally, the chapter discussed the NUMA and the various pros and cons of horizontal scalability, vertical scalability, JVM consolidation, and VM consolidation.

vmware® PRESS

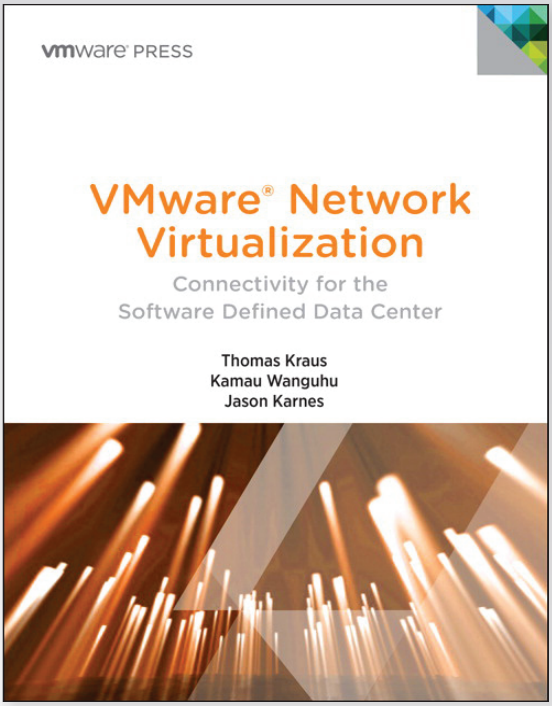

vmware® PRESS

VMware® Network Virtualization

Connectivity for the Software Defined Data Center

Thomas Kraus
Kamau Wanguhu
Jason Karnes

CHAPTER 2

Network Virtualization Defined

AVAILABLE – FEBRUARY/MARCH 2014

vmwarepress.com

Chapter 2

Network Virtualization Defined

This chapter defines *Network Virtualization* in the context of related yet vastly different network technologies. At the conclusion of this chapter, the reader should have a solid understanding of network virtualization, what technical problems it solves, and the various approaches to achieve it. This chapter also explains why a business should care about network virtualization and the benefits it provides.

Future's Past

Beginning in the early 2000s, the modern datacenter underwent a complete transformation when *x86* virtualization was introduced into datacenters around the world. Server virtualization offered customers unprecedented flexibility and cost savings through four fundamental characteristics:

- **Isolation:** Allowed multiple virtual machines to co-exist on the same physical hypervisor but remain completely separate from a compute, storage, and even network perspective.
- **Encapsulation:** Reduced the all too complex hardware and software of what previously was a physical machine into a discreet set of files.
- **Hardware Independence:** Enabled portability through various methods by abstracting the virtual machines from the underlying hardware.
- **Partitioning:** Took a large set of hardware resources and segmented them into virtual machines each running on a *hypervisor*.

Server virtualization radically changed the way compute resources are consumed. This has transformed data center design, the server hardware industry, and how IT Staff operates and manages their environments. When first embarking on server virtualization, businesses managed virtual machines in the same way as physical machines. The cost savings due to more efficient asset utilization and increases in productivity derived from decreased time to provision servers provided a large benefit. To gain further advantage, businesses refined and streamlined their operations to gain the most value from server virtualization.

Cloud Computing has emerged as a way to further reduce costs and decrease project delivery times. To realize the full benefit of a Cloud based *Infrastructure as a Service* (*IaaS*) strategy, each component of the infrastructure must satisfy at minimum, four requirements:

1. Self-service consumption of Infrastructure
2. Resource pooling
3. Elasticity
4. Consumption metering

Servers, storage, and networks must be transformed in such a way as to satisfy these requirements. The fundamental characteristics of virtualization provide a solid foundation upon which to build new capabilities that satisfy the requirements of a Cloud based *IaaS* strategy.

Cloud based *IaaS* strategies focus primarily on offering services to consumers. Based on the consumer demands, the consumers request infrastructure resources that meet certain tier of service requirements such as performance, cost, or security. It is the job of the *Cloud Management Platform* (*CMP*) to determine where best to provide those resources.

To broaden the benefits of a Cloud based *IaaS* strategy, we must begin to think beyond the consumer to the provider. Ideally, an *IaaS* provider should be able to offer services that are seamless and transparent across multiple sites and tiers of service.

The *Software Defined Data Center* (*SDDC*) combines the Cloud *IaaS* benefits to consumers with the advanced seamless and transparent delivery of services by providers.

Why the world needs Network Virtualization

A *SDDC* strategy provides consumers with a complete *Virtual Data Center* (*VDC*) that transparently spans provider sites and tiers of service. Compute, storage, and network services must be available and accessible seamlessly across traditional provider boundaries. Networking is a key requirement to realize this strategy and poses technological and operational limitations.

In order for a *VDC* to appear seamless, the network services provided to end stations must operate in such a way as to appear within the same local network. Since the consumer manages the *VDC* resources, networks may be created, consumed, and destroyed at any time. For example, a consumer may wish to create a single *Layer 2* (*L2*) switch and connect three end stations to the said switch. Those end stations might exist in different physical provider sites that could span the globe. While solutions do exist to carry a *VLAN* across sites, the life cycle management of *VLANs* is challenging from an IT Operations perspective and in no way would meet the self-service requirements of a Cloud based *IaaS*, or *SDDC* strategy.

Networks today are often managed as a standalone system designed to provide services to consumers wanting to connect in a similar fashion as attaching to the power grid. Network designs tend to be rigidly defined and changes to support new business needs are cumbersome at best and often extraordinarily difficult to implement due to security challenges and potential interruption to existing systems. Network administrators are involved in an ever-escalating battle to provide high performance networks and flexible designs that meet security requirements.

Just as virtualization solved many of the challenges faced by server administrators, the four fundamental characteristics of server virtualization can be applied to virtualizing the network. Virtualized networks can provide the flexibility, security, and life cycle management capabilities required to realize a Cloud based *IaaS* or *SDDC* strategy.

What is Network Virtualization?

Network Virtualization creates a comprehensive *logical network* infrastructure in software by abstracting the underlying physical network. Just as physical servers are still required to run virtual machines, a physical network is still required to transport data for virtual networks. Software is used to decouple the Logical from the Physical network providing the abstracted networking functionality needed by end stations. The underlying physical network can now be architected to provide a very resilient and high bandwidth *Layer 3 (L3) IP* forwarding fabric. The use of software to implement higher level networking features provides tremendous value and flexibility. Software can be seamlessly upgraded to extend the capabilities of the network at a much faster rate, as it does not require an upgrade to integrated hardware capabilities. Software can also be automated and programmed easily through an exposed *API* with common standards based development frameworks and languages.

Network Virtualization benefits as derived from the four fundamental characteristics of virtualization:

1. **Partitioning** - Enables an underlying physical network to be divided into isolated *L2* segments or logical networks for individual tenants or application tiers.
2. **Isolation** - Provides separate *L2* broadcast domains in a logical space otherwise known as a logical switch. Logical switches contain logical networks that are completely isolated from other logical and physical networks through tunneling and encapsulation protocols. This enables multiple isolated logical switches to share the same underlying physical transport network without having to configure the physical network for isolation with *VLANs*.
3. **Encapsulation** - Distributed network virtualization technology can encapsulate a complex *L2* to *L7* network configuration into a logical switch. This logical switch encapsulates a logical network that contains the configuration or state information normally scattered across different tiers of a physical network.
4. **Hardware Independence** – The abstraction of the physical (*transport*) network and the ability to reproduce the same functionality in a logical space eliminates the need for much of the functionality (complexity) previously created by hardware vendors at the physical network layer. A physical transport network is always required, however, hardware independence frees the infrastructure from a dependence on the features and capabilities (complexity) of the underlying hardware. The physical network is simplified, and needs only to provide a resilient and high performance *L3* forwarding fabric.

In addition to these four characteristics, fifth fundamental characteristic of network virtualization is automation. To provide this capability, an *API* is exposed for consumption by an *IaaS* Orchestration or *Cloud Management Platform (CMP)* layer. By exhibiting these five characteristics, it becomes possible for network virtualization to provide additional high value benefits.

- Faithful reproduction of networking services. At a minimum, end stations connecting to logical networks have the same network functionality and services available to them as end stations connected to traditional physical networks. Network services are transparent to the end station as there are no changes required on the end station to connect to a logical network.
- Centralized configuration management of all *L2-L7* services for logical networks. In a traditional network infrastructure, configuration elements for networks are scattered across multiple network devices including: switches, routers, load balancers, firewalls, and the hosts themselves.
- The ability to automate the lifecycle of a network. In the traditional network infrastructure, networks exist independently of the workloads attached to them. With

network virtualization, logical networks can be created, consumed, and destroyed as needed by workloads all by an automated *CMP*.

Figure 2.1 Automation

Network Virtualization Use Cases

Now that we have a high-level understanding of what network virtualization is, we can begin to explore common use cases. Each customer has their own unique set of requirements and constraints that lead to endless use case possibilities for network virtualization. While it is not practical to cover every conceivable use case of network virtualization, the most common use cases fall into these primary categories:

- **Programmability/Automation**: Programmatic control of the network enables seamless integration with compute management tools for basic *L2* and *L3* network connectivity as well as *L4* through *L7* services. This use case provides unprecedented flexibility and control over the logical network lifecycle. Administrators now have control over provisioning, policy, services, consumption, and reclamation of logical networks.

- **Self-service/Cloud**: Cloud based *IaaS* strategies expose complex *L3* networking features and software to seamlessly coordinate integration between the network and compute resources. This allows end-users to dynamically create flexible and powerful logical network topologies without reliance on the configuration of the underlying physical network. The ability to decouple the logical from the physical network is what makes this possible.

- **Data center to data center connectivity**: With network virtualization, it is possible to create logical switches that span multiple sites. This allows *L2* networks to appear locally available at each site and end stations to appear as if they are on the same local network segment. Network virtualization also extends the ability to bridge logical and physical networks together to create a single seamless *L2* network.

Figure 2.2 Multi-Site connectivity

- **Server Migrations and Cloud Bursting**: When using network virtualization, it is a simple matter to control or extend *L2* network boundaries of a logical network to bridge to a physical network. This flexibility enables physical and virtual servers to seamlessly communicate over a single *L2* network domain. Physical and virtual servers can then share a single *IP* address space with *L2* network services operating the same way as a traditional physical *L2* network (*VLAN*). This capability provides a highly effective means to convert servers from physical to virtual, or to migrate servers from one data center or cloud provider to another.

How Network Virtualization is Achieved

To provide network virtualization, the physical network must be abstracted and its functions reproduced in a logical space. Workloads connected to logical networks require the same fundamental services as physical networks including switches (*L2*) and routers (*L3*). Abstraction is the key factor of network virtualization as it enables all of the flexibility and physical topology independence inherent to a virtualized network strategy. Encapsulation or tunneling is the most common way of achieving this abstraction along with a programmatic ability to redirect or process packets according to a logical fabric created by a centralized controller. As an example of this abstraction, two virtual machines on a logical network segment can exist in the same logical broadcast domain although the virtual machines may be running on two different hypervisors. Those two hypervisors may be connected to different *L3* physical networks connected by a physical router. While network virtualization provides this abstraction, it must also accurately reproduce all basic network services in a logical space. As with *x86* virtualization, we do not want to modify the guest *Operating Systems* (as para-virtualization requires) in order to have them connect to the logical networks. To achieve this, network virtualization replicates the functionality of physical networking services such as: *L2* broadcast domains, *L2* learning, and delivery of unicast traffic and Broadcast, Multicast, and Unknown unicast (*BUM*) traffic. Other higher-level protocols that depend on these services such as *DHCP*, *DNS*, *HTTP*, or even *IP/TV* must work transparently in a network virtualization environment.

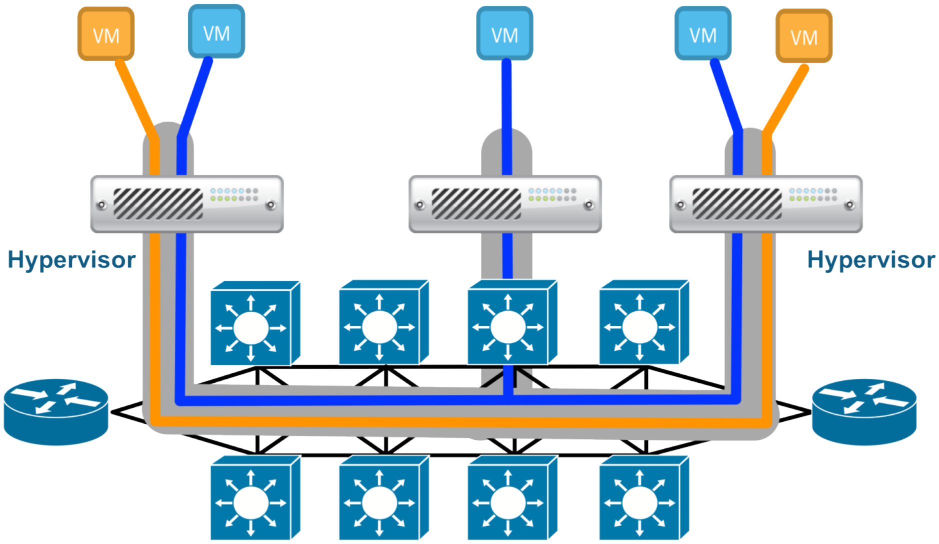

Figure 2.3 Network Virtualization Encapsulation

Network Virtualization Components

Regardless of the specific vendor, a network virtualization solution should typically consist of three key components. The existence of these components should be considered the basic requirements for entry and without all of them, the promise and value of network virtualization will be very difficult to realize. Different vendors provide added benefits and features, which are also critical depending on the use case. Features such as scalability, resiliency, and performance, which are also critical in an infrastructure offering, should be provided by the solution as opposed to being bolted on. The three key components are:

1. Management Plane – The Management plane is responsible for providing programmatic or *UI* based control of the environment. A modern web services based *API* is critical at this layer. The *Cloud Management Platform* and the *Network Virtualization Management Tool* are part of this layer.
2. Control Plane – The control plane is the centralized intelligence that manages the endpoints and network transport nodes in the data plane. This centralized intelligence, or controller, translates intended actions initiated by the management plane and converts them into specific configurations governing the data plane.
3. Data Plane – The data plane is where forwarding decisions are made, packets are analyzed, and processed for delivery. Processing can mean any combination of encapsulation/de-encapsulation, modification of packet headers, dynamic learning, and dropping or forwarding of packets. In a distributed network virtualization solution, the data plane is extended from the network edge switches into the hypervisor. In a hypervisor, this would exist as modules running within the hypervisor to provide virtual switch services. This distribution of the data plane enhances scalability and performance.

NOTE

This degree of separation eliminates many *Denial of Service* vectors, reduces impact of a single plane failure, and provides insulation to prevent information leakage about activities occurring in each plane.

NOTE

While encapsulation is key to providing abstraction from the underlying hardware and creating logical topologies, the specific protocol used is less important. Specific protocols used for tunneling and encapsulation are covered in Part 2. For now, the selection of protocol (VXLAN, STT, GRE, etc.) is not critical to receive the benefits of network virtualization.

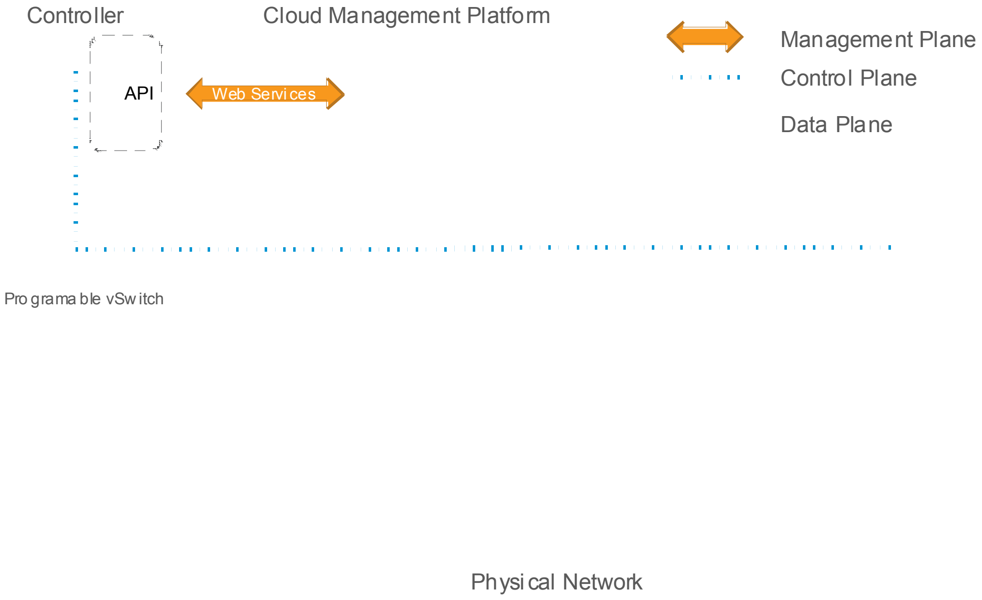

Figure 2.4 Network Virtualization Components

So What is Software Defined Networking?

Now that we have a baseline understanding of what network virtualization is, how does this contrast to *Software Defined Networking* (*SDN*)? *SDN* is an often-overused term that describes the ability of a centralized software program or application to control the entire network. But wait! Isn't software control a characteristic of network virtualization? Yes it is but it is not the only characteristic. You can achieve *SDN* by installing a central controller that will manage physical networking equipment or maybe a combination of physical and software based networking devices. Think of *SDN* as trying to take existing physical networking functionality and protocols and making the management and automation of the systems providing that functionality more dynamic. Contrast this to network virtualization, which provides the higher-level functionality of the network in a logical space using software and leaves the physical network to facilitate fast and resilient delivery of data. Typically this requires separating the control logic, or the control layer, from the underlying physical networking equipment that processes network packets and frames. It can be argued that *SDN* goes back to the 1980s, when telephone carriers began to provide software level management and control of their networked voice switching. The management and control of remote switching equipment was done with in-band tones at different frequencies than those used for voice traffic. This was an attempt at separating the management plane from the data plane. There was no security on the management plane and that made it possible to send control messages to the voice switching equipment and manipulate data plane traffic handling. Steve Wozniak famously exploited this weakness when he developed his blue boxes to make free phone calls across the public telephone network. Software control of the network is a good start, but the separation of the three components: management, control, and data planes are critical to making an *SDN* solution robust and scalable. Network virtualization goes beyond *SDN* by making it possible to add higher-level services at *L4* through *L7*.

NOTE

While *SDN* and network virtualization are very different, the concept of a centralized and programmatic control plane is consistent between both *SDN* and network virtualization.