VMware vCloud® Director™ Infrastructure Resiliency Case Study

VMware vSphere® 5.0, VMware® vCenter™ Site Recovery Manager™ 5.0 and VMware vCloud Director 1.5

TECHNICAL WHITE PAPER
V 1.0 FEBRUARY 2012

vmware®

Table of Contents

Design Subject Matter Experts

The following people provided key input into this design:

NAME	TITLE	ROLE
Duncan Epping	Principal Architect – Technical Marketing	Author
Chris Colotti	Consulting Architect – Center of Excellence	Contributor

Purpose and Overview

VMware vCloud® Director™ 1.5 (vCloud Director) gives enterprise organizations the ability to build secure private clouds that dramatically increase datacenter efficiency and business agility. Coupled with VMware vSphere® (vSphere), vCloud Director delivers cloud computing for existing datacenters by pooling vSphere virtual resources and delivering them to users as catalog-based services. vCloud Director helps build agile infrastructure-as-a-service (IaaS) cloud environments that greatly accelerate the time to market for applications and increase the responsiveness of IT organizations.

Resiliency is a key aspect of any infrastructure—it is even more important in infrastructure-as-a-service solutions. This case study was developed to provide additional insight and information as to how to increase availability and recoverability of a vCloud Director–based infrastructure using VMware® vCenter™ Site Recovery Manager™ (SRM) as well as common disaster recovery methodologies and tools. SRM facilitates fast and reliable recovery and enables you to meet your recovery time objectives (RTOs) by automating the failover process of your vCloud Director management environment.

Target Audience

The target audience of this document is an individual with a technical background who will be designing, deploying or managing a vCloud Director infrastructure, including but not limited to technical consultants, infrastructure architects, IT managers, implementation engineers, partner engineers, sales engineers and customer staff. This solution brief is not intended to replace or override existing certified designs for vCloud Director. It instead is meant to supplement knowledge and provide additional information for implementing a disaster recovery strategy for vCloud Director infrastructures. vCloud Director infrastructure architectural guidance is provided through the VMware vCloud® Reference Architecture Toolkit.

Interpreting This Document

The overall structure of this design document is largely self-explanatory. However, throughout this document several key points will be highlighted to the user by means of the following label:

• NOTE – A point of general importance or a further explanation of a particular section.

This document captures a solution developed for a specific scenario and set of requirements. It is assumed that the reader is familiar with vCloud Director, VMware vCenter Server™, SRM and vSphere reference architectures, technology and terminology.

Case Requirements and Assumptions

Requirements are the key demands on the design. Sources include both business and technical representatives.

ID	REQUIREMENT
R101	Increasing availability of vCloud Director infrastructure
R102	Failover management cluster workload to secondary site
R103	Failover vCloud Director resource cluster workload to secondary site
R104	Must be a fully supported solution

Table 1. Customer Requirements

The following list includes overall assumptions for this particular scenario and considerations to be made before utilizing information contained in this document:

ID	ASSUMPTION
A101	Stretched layer-2 network
A102	Storage-based replication technology
A103	Use of SRM to orchestrate management cluster disaster recovery

Table 2. Assumptions

Infrastructure Logical Architectural Overview

vCloud Director infrastructures must be deployed according to the VMware vCloud Architecture Toolkit (vCAT) defined in version 2.0.1. The vCAT prescribes a scenario where the vCloud Director elements are explicitly separated into two groups, a management cluster and a resource cluster.

- The management cluster contains the elements required to operate and manage the vCloud Director environment. This typically includes vCloud Director cells, vCenter Server(s) (used for resource clusters), VMware® vCenter™ Chargeback Manager™, VMware® vCenter™ Orchestrator™, VMware® vShield Manager™ and one or more database servers.
- The resource cluster represents dedicated resources for end-user consumption. Each resource group consists of VMware® ESXi™ hosts managed by a vCenter Server and is under the control of vCloud Director, which can manage the resources of multiple clusters, resource pools and vCenter Servers.

This separation is recommended primarily to facilitate quicker troubleshooting and problem resolution. Management components are strictly contained in a relatively small and manageable cluster. Running management components on a large cluster with mixed environments can be time consuming and might make it difficult to troubleshoot and manage such workloads.

Separation of function also enables consistent and transparent management of infrastructure resources, critical for scaling vCloud Director environments. It increases flexibility because upgrades for management and resource clusters are not tied. And it prevents security attacks or intensive provisioning activities from affecting management component availability.

Figure 1 depicts this scenario.

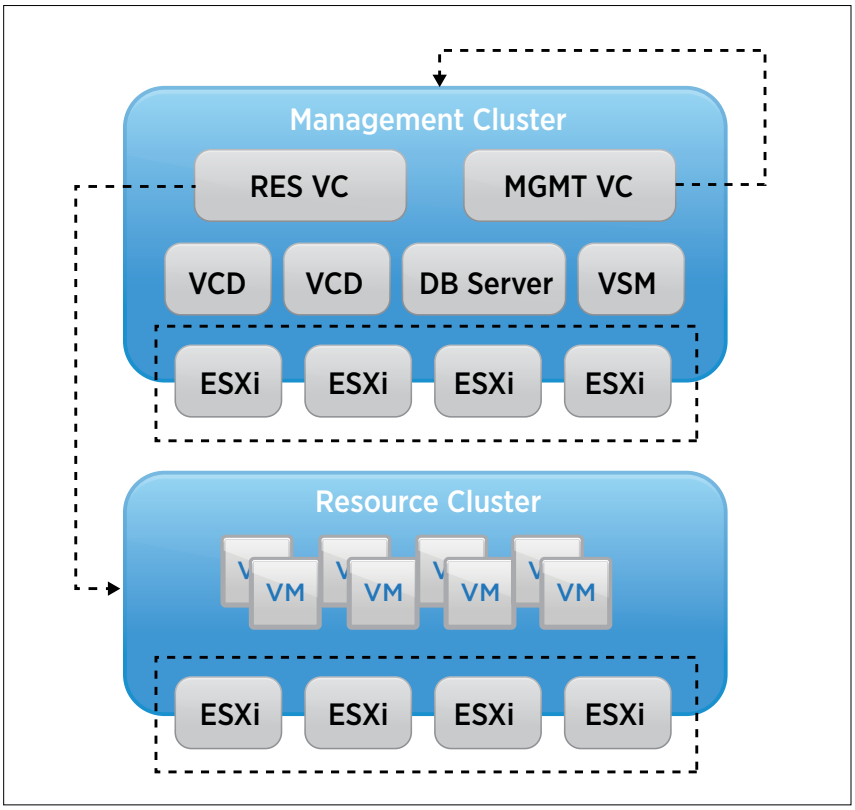

Figure 1. VMware vCloud Infrastructure Logical Overview

Management Cluster Overview

The management cluster hosts all the necessary vCloud infrastructure components. In our scenario, it contains at a minimum the following virtual machines:

- Two vCenter Server systems
  - One vCenter Server instance for management (of the management cluster)
  - One vCenter Server instance for cloud resources managing the resource cluster
- One vCenter Server database
- Two vCloud Director cells
- One vCloud Director database
- One vShield Manager

The management vCenter Server is running as a virtual machine in the cluster it is managing. The black arrow in Figure 2 depicts this.

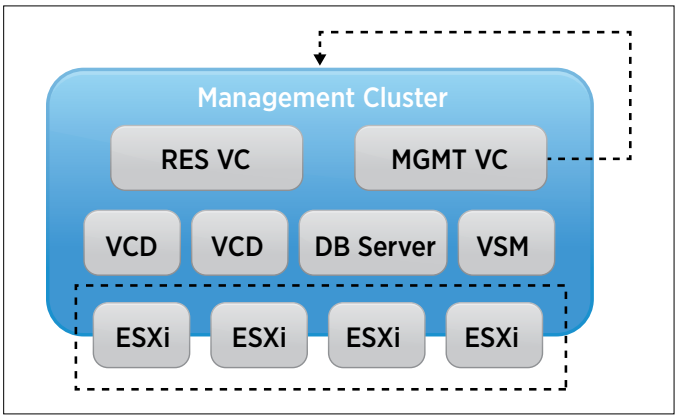

Figure 2. Management Cluster Overview

Resource Cluster Overview

A resource cluster is a set of resources dedicated to end-user workloads and managed by a single vCenter Server instance. vCloud Director manages all the resource clusters through the vCenter Server instances attached to it. All provisioning tasks are initiated through vCloud Director and are passed down to the appropriate vCenter Server instance residing in the management cluster.

A resource cluster contains only virtual machines instantiated by vCloud Director; this includes the VMware® vShield Edge™ appliances for network services. Figure 3 depicts a resource cluster. The vCenter Server virtual machine that manages this environment is not depicted because it resides within the management cluster.

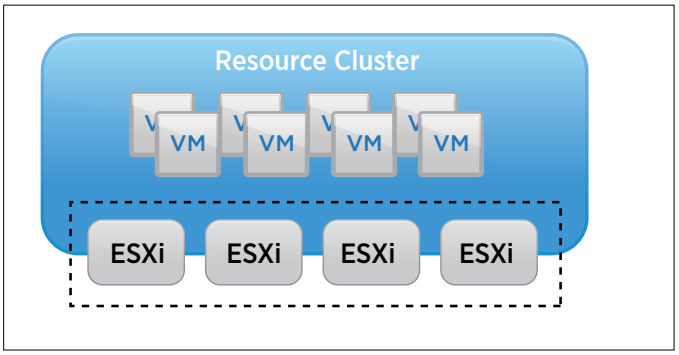

Figure 3. Resource Cluster Overview

Creating a Disaster Recovery Solution for vCloud Director

As of this writing, VMware vCenter Site Recovery Manager 5.0 (or prior) does not support the protection of vCloud Director workloads (resource clusters). To facilitate disaster recovery (DR) in a vCloud Director environment, the proposed solution uses standard disaster recovery concepts that leverage conventional replication technologies and vSphere features for vCloud Director workloads and orchestrate the failover of the vCloud management infrastructure (management cluster) using SRM.

SRM currently does not support vCloud Director because it is designed specifically to control the destination vCenter Server to manage rapid failover of virtual machines to the recovery site. To facilitate controlling failover of one vSphere environment to another, SRM requires some vCenter Server objects—such as resource pools, folders, and port groups—to be precreated on the destination side. Because vCloud Director is designed to fully control a resource cluster, it maintains complete authority and surveillance over all objects created in the cluster. Attempting to use SRM to fail over a resource cluster would result in SRM's precreating placeholder objects for resources at the primary site, and the objects' being unknown to vCloud Director. In testing SRM and vCloud Director together, we have identified the following automatically created objects as posing a challenge when newly created:

- Virtual machines
- Resource pools
- Folders
- Port groups

In addition to those that are automatically created, the following objects are used and referenced by vCloud Director:

- Clusters
- Datastores

Both vCloud Director and vCenter Server heavily rely on management object reference identifiers (MoRef IDs) to correlate the objects between the two platforms. Any unplanned changes to these identifiers will result in loss of functionality, because vCloud Director will not be able to manage these objects. The screenshot in Figure 4 displays the use of a MoRef ID within the vCloud Director database. The use of SRM would result in a change of the MoRef ID on the vCenter Server layer, resulting in an incorrect reference in the vCloud Director database leaving the object (for instance, a virtual machine), which is unmanageable from a vCloud Director perspective.

Results | Messages

	on	primary_nic_id	guest_cust_vm_id	guestos_id	moref
1		0xB5AADAB7938F48D294E2163FA6B20011	88f7a6c6-f803-45f9-85cf-224058c4fce2	3	vm-231
2		0x4211832293F148B9AD0F466EF31B5088	75fbac68-bec5-461b-9c3a-2134ebdd4aeb	3	vm-233
3		0x52E70308BE9F4391891F518965EAD775	75fbac68-bec5-461b-9c3a-2134ebdd4aeb	3	vm-235
4		0xF2002B350676404C82B24376CD152018	e499e8cf-6f92-47cf-bcc3-978bb9c509c3	3	vm-241

Figure 4. Example of a vCloud Director SQL Database Table Containing MoRef IDs

Another common issue when discussing SRM and vCloud Director together is the resignaturing of VMware vSphere® VMFS volumes. In a standard SRM environment, this is a common DR best practice that is done when replication has been stopped and the "replicated" volume is presented to the hosts in the recovery site. It is done to ensure that the volume on the recovery side has a unique ID. This prevents a scenario where two volumes with the same unique ID (UUID) are presented to the same host, which potentially can lead to data corruption. When a datastore is resignatured, there is a requirement to reregister all virtual machines within vCenter Server.

vCloud Director references these virtual machines by MoREF ID and cannot handle these changes. As such, avoiding resignaturing volumes is a requirement.

The proposed solution will prevent changes to any of these objects. This simplifies the recovery of a vCloud Director infrastructure and increases management infrastructure resiliency.

Logical Architecture Overview

vCloud Director disaster recovery can be achieved through various scenarios and configurations. This case study focuses on a single scenario as a simple explanation of the concept, which can then easily be adapted and applied to other scenarios. This case study focuses on an "active/standby" DR approach in which hosts at the recovery site are not utilized under normal conditions.

To ensure that all management components are restarted in the correct order and in the least amount of time, SRM is used to orchestrate the failover. This requires that each site contain a management vCenter Server and an SRM server, as depicted in Figure 5.

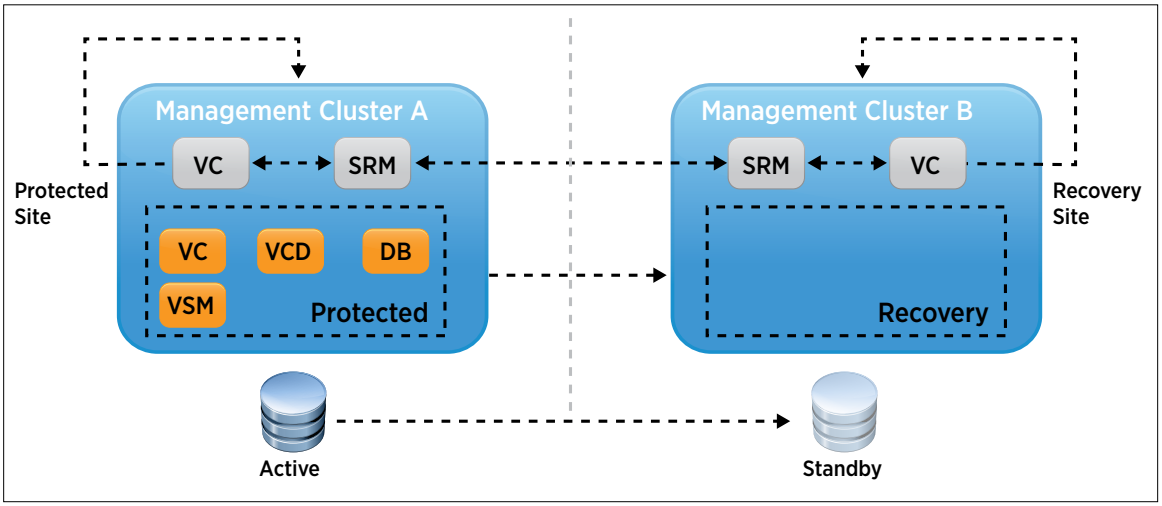

Figure 5. Management Cluster Overview

Figure 6 depicts the full vCloud Director infrastructure architecture used for this case study.

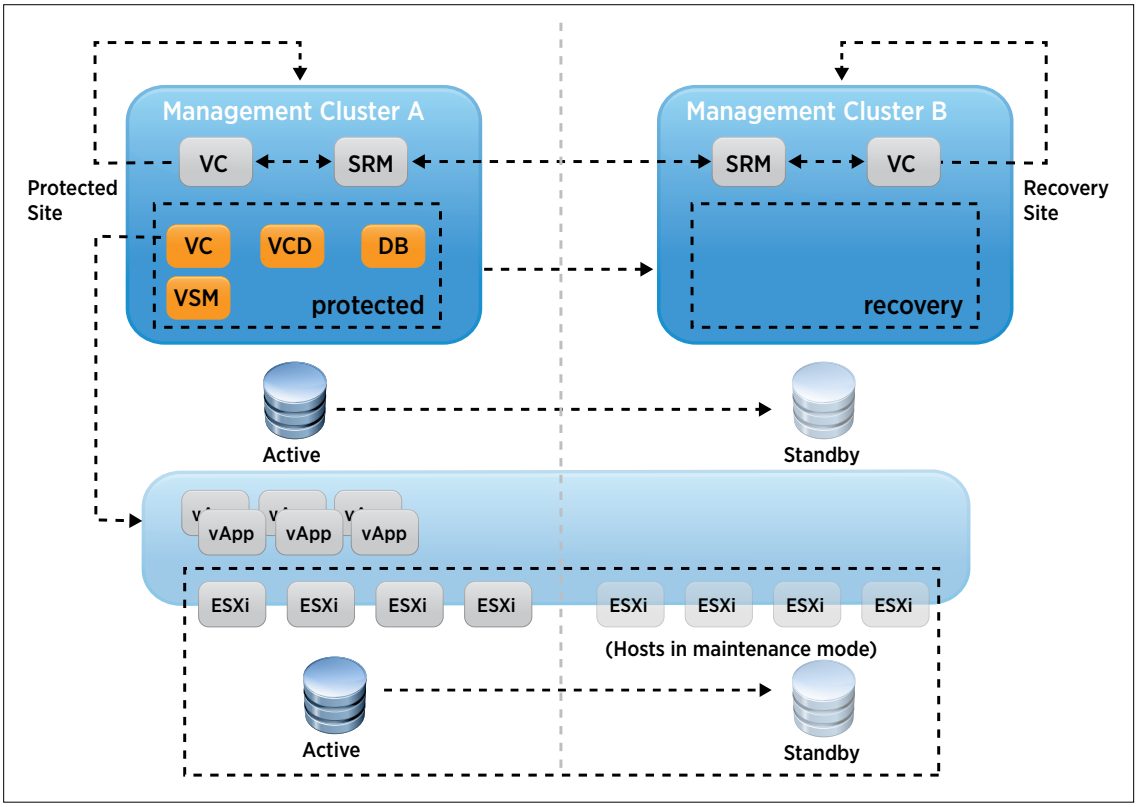

Figure 6. Management Cluster Overview

Both the protected site and the recovery site have a management cluster that contains a vCenter Server instance and an SRM server. These servers facilitate the DR procedures for the other components contained within the management cluster by including them in an SRM protection group. By creating an SRM recovery plan, the virtual machines within this protected group can be failed over to the recovery site. More details regarding the creation of recovery plans, protection groups and SRM configuration in general can be found at http://www.vmware.com/support/pubs/srm_pubs.html.

Storage is replicated and not stretched in this environment. Hosts in the resource cluster at the recovery site are unable to detect the storage at the protected site, so they cannot run vCloud Director workloads in a normal situation. They are depicted as hosts that are placed in maintenance mode. They can also be standalone hosts added to the vCloud Director resource cluster during the failover. For simplification and visualization purposes, this scenario describes the situation where the hosts are part of the cluster and are placed in maintenance mode.

Storage replication technology is used to replicate LUNs from the protected site to the recovery site. This can be done using asynchronous or synchronous replication. It typically depends on the recovery point objective (RPO) determined in the service-level agreement (SLA) and the distance between the protected site and the recovery site. In our case study, synchronous replication was used.

SRM manages the LUNs/datastores on which the vCloud Director management infrastructure is hosted leveraging a storage replication adapter (SRA). Through the use of an SRA, it is possible to fully automate and orchestrate a failover for the vCloud Director management infrastructure.

The LUNs/datastores on which the vCloud Director workloads are running are not managed by SRM because this is currently not supported. As a result, manual steps are possibly required during the failover. Depending on the type of storage used, these steps can be automated leveraging storage system API calls.

Failover Procedure

In this section, the steps required for a successful failover of a vCloud Director environment are described. These steps are pertinent to the described scenario.

It is essential that each component of the vCloud Director management stack be started in the correct order. This is facilitated by SRM. The order in which the components should be started is configured in an SRM recovery plan, and the process can be initiated by SRM with a single button. The following order was used to power on the vCloud Director management virtual machines:

1. Database server (providing vCloud Director and resource vCenter Server databases)
2. Resource vCenter Server
3. vShield Manager
4. vCenter Chargeback Manager (if in use)
5. vCloud Director cell 1
6. vCloud Director cell 2

After the failover of the vCloud Director management virtual machines in the management cluster has succeeded, the vCloud Director workloads can be failed over. The following are the manual steps required for this, but they can be automated using VMware vSphere® PowerCLI™ or VMware vCenter Orchestrator:

1. Validate that all vCloud Director management virtual machines are powered on.
2. Using your storage management utility, break replication for the datastores connected to the vCloud Director resource cluster and make the datastores read/write (if required by the storage platform).
3. Mask the datastores to the recovery site (if required by the storage platform).
4. Using ESXi command-line tools, mount the volumes of the vCloud Director resource cluster on each host of the cluster.
   - esxcfg-volume –m <volume ID>.
5. Using vCenter Server, rescan the storage and validate that all volumes are available.
6. Take the hosts out of maintenance mode for the vCloud Director resource cluster (or add the hosts to your cluster, depending on the chosen strategy).
   - Install the vCloud Director agent (prepare host), if required.
7. Power on the vCloud Director workload virtual machines.
   - In our tests, the virtual machines were automatically powered on by VMware vSphere® High Availability, which detects the situation before the failover and will power on the virtual machines according to the last known state.
   - The power-on procedure can also be scripted leveraging the vCloud API to ensure that all virtual machines that are part of a VMware vSphere® vApp™ are booted in the correct order. This might result in the powering on of vApps that had been powered-off before the failover, because there is currently no way of determining their state.

At this point, a failover of a vCloud Director environment has been successfully completed. Because all vCloud Director management components are virtualized, the virtual machines are moved over to the recovery site while maintaining all currently managed object reference identifiers (MoRef IDs). Leveraging standard vSphere functionality, *mount volume*, eliminates the need to resignature the datastore (giving it a new unique ID) and reregister all the virtual machines.

Figure 7 depicts the previously described failover in further detail with various "power on" options. It should be used as a guide to develop a customized vCloud Director infrastructure resiliency strategy.

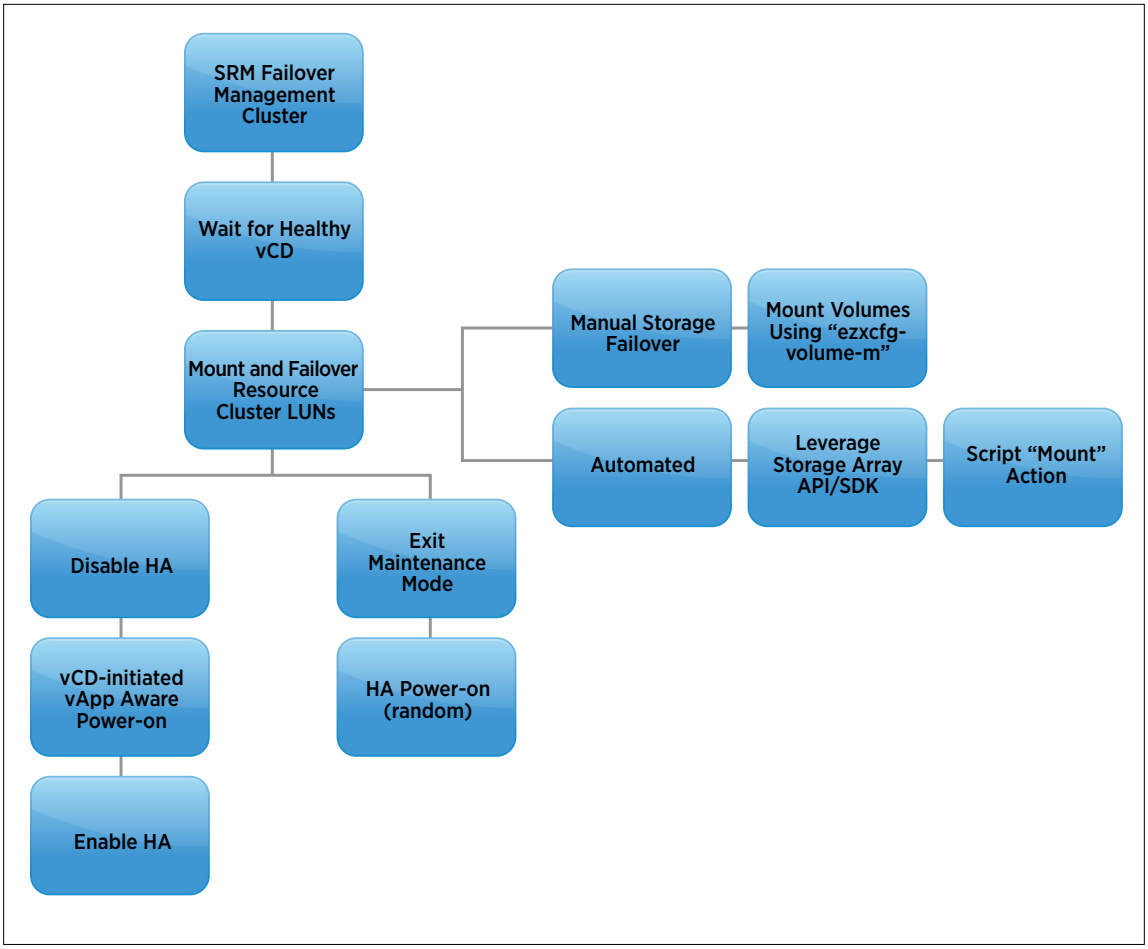

Figure 7. Flow Diagram of vCloud Director Environment Failover

Conclusion

This case study demonstrates how VMware vCloud infrastructure resiliency can be achieved by leveraging basic VMware vSphere and VMware vCenter Site Recovery Manager functionality. This solution enables disaster recovery of the (full) cloud. By virtualizing all management components and following the prescribed procedure, a simple and effective DR strategy can be implemented.

For more details about vCloud infrastructure resiliency, contact your local VMware sales representative.

About the Authors

Duncan Epping is principal architect in the Technical Marketing group at VMware and is focused on cloud infrastructure management architecture. Previously, he worked at Oracle and multiple consultancy companies, where he had more than 10 years of experience designing and developing infrastructures and deployment best practices. He was among the first VMware certified design experts (VCDX 007). He is the co-author of several books, including best sellers *vSphere 5.0 Clustering Technical Deepdive* and *VMware vSphere 4.1 HA and DRS Technical Deepdive*. He is the owner and main author of the leading virtualization blog yellow-bricks.com.

- Follow Duncan Epping's blogs at http://www.yellow-bricks.com and http://blogs.vmware.com/vSphere.
- Follow Duncan Epping on Twitter: @DuncanYB.

Chris Colotti is a consulting architect with the VMware vCloud Delivery Services team. He has more than 10 years of experience working with IT hardware and software solutions. He holds a bachelor of science degree in Information Systems from the Daniel Webster College. Prior to coming to VMware, he served a Fortune 1000 company in southern New Hampshire as a systems architect/administrator, architecting VMware solutions to support new application deployments. At VMware, Chris has guided partners as well as customers in establishing a VMware practice and has consulted on multiple customer projects ranging from datacenter migrations to long-term residency architecture support. Chris is also among the first 50 VMware Certified Design Experts (VCDX37).

- Follow Chris Colotti's blogs at http://www.chriscolotti.us and http://blogs.vmware.com/vCloud.
- Follow Chris Colotti on Twitter: @CColotti.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: VMW-CS-vCLD-DRCTR-INFRA-USLET-101