

vmware® PRESS



# Managing and Optimizing vSphere® Deployments

IT BEST PRACTICES

**Harley Stagner**  
**Sean Crookston**



# **Managing and Optimizing VMware vSphere<sup>®</sup> Deployments**

VMware Press is the official publisher of VMware books and training materials, which provide guidance on the critical topics facing today's technology professionals and students. Enterprises, as well as small- and medium-sized organizations, adopt virtualization as a more agile way of scaling IT to meet business needs. VMware Press provides proven, technically accurate information that will help them meet their goals for customizing, building, and maintaining their virtual environment.

With books, certification and study guides, video training, and learning tools produced by world-class architects and IT experts, VMware Press helps IT professionals master a diverse range of topics on virtualization and cloud computing and is the official source of reference materials for preparing for the VMware Certified Professional Examination.

VMware Press is also pleased to have localization partners that can publish its products into more than 42 languages, including, but not limited to, Chinese (Simplified), Chinese (Traditional), French, German, Greek, Hindi, Japanese, Korean, Polish, Russian, and Spanish.

For more information about VMware Press, please visit  
<http://www.vmware.com/go/vmwarepress>

**vmware® PRESS**

# Managing and Optimizing VMware vSphere® Deployments

Sean Crookston  
Harley Stagner

vmware® PRESS

Upper Saddle River, NJ • Boston • Indianapolis • San Francisco  
New York • Toronto • Montreal • London • Munich • Paris • Madrid  
Cape Town • Sydney • Tokyo • Singapore • Mexico City

## Managing and Optimizing VMware vSphere Deployments

Copyright © 2013 VMware, Inc., Sean Crookston, and Harley Stagner

Published by Pearson Education, Inc.

Publishing as VMware Press

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise.

ISBN-10: 0-321-82047-9

ISBN-13: 978-0-321-82047-1

*Library of Congress Cataloging-in-Publication data is on file.*

Printed in the United States of America

First Printing: July 2012

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. The publisher cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

VMware terms are trademarks or registered trademarks of VMware in the United States, other countries, or both.

### Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The authors, VMware Press, VMware, and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the CD or programs accompanying it.

The opinions expressed in this book belong to the authors and are not necessarily those of VMware.

### Corporate and Government Sales

VMware Press offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales  
(800) 382-3419  
corpsales@pearsontechgroup.com

For sales outside the United States please contact:

International Sales  
international@pearson.com

### VMWARE PRESS

#### PROGRAM MANAGER

Erik Ullanderson

#### ASSOCIATE PUBLISHER

David Dusthimer

#### EDITOR

Joan Murray

#### DEVELOPMENT EDITOR

Ellie Bru

#### MANAGING EDITOR

Sandra Schroeder

#### SENIOR PROJECT

#### EDITOR

Tonya Simpson

#### COPY EDITOR

Karen Annett

#### PROOFREADER

Debbie Williams

#### INDEXER

Tim Wright

#### EDITORIAL ASSISTANT

Vanessa Evans

#### BOOK DESIGNER

Gary Adair

#### COMPOSITOR

Bronkella Publishing

*Sean Crookston—I would like to dedicate this book to my wife and son. Without their support this book would not have been possible.*

*Harley Stagner—I would like to dedicate this book to my wife, Kimberly. Her infinite supply of support and patience during this process made this book possible.*

*This page intentionally left blank*

# Contents

## Preface xi

### 1 Laying the Groundwork 1

- Planning 1
  - Capacity 1
  - Performance 13
  - Management 19
  - Designing 20
- Laying the Groundwork Summary 39

### 2 Implementing the Solution 41

- Following the Design Blueprint 41
  - Reviewing the Design Documentation 42
  - Stakeholder Review 42
  - Functional Requirements 43
  - Constraints 44
  - Technical Review 46
  - Assumptions 47
  - Design Deviations 48
- Automating Implementation Tasks 53
  - PowerCLI 53
  - Host Profiles 53
  - Auto Deploy Server 55
  - vCenter Orchestrator 57
- Verifying Implementation 58
  - Testing Functionality 58
  - Quality Assurance Assessment 76
- Implementing the Solution Summary 84

### 3 Operating the Environment 85

- Backups 85
  - Data Recovery 88
- Disaster Recovery 95
  - Manual Disaster Recovery 96
  - Site Recovery Manager 96

Physical to Virtual Conversions	97
Issues and Troubleshooting	98
Domain Controllers	99
Windows Server with OEM Installations	100
SQL, Exchange, and Other Applications Servers	100
Linux	101
V2V and Other Methods	101
Maintenance	102
Update Manager	103
Monitoring	113
Alerting	113
Defining Actions for Alarms	114
Considerations for Tweaking Default Alarms	118
Verifying Configurations	119
Host Profiles	120
Health Check	121
VMware's Health Check Delivery	121
Operating the Environment Summary	122

#### **4 Managing the Environment 123**

Capacity Management	123
Storage Capacity Management	123
Host Capacity Management	124
Networking Capacity Management	130
Performance Management	149
Storage Performance Management	149
Host Performance Management	152
Networking Performance Management	155
vCenter Operations Management Suite	160
Managing the Environment Summary	165

#### **5 Roadblocks to 100% Virtualization 167**

Political Roadblocks	168
Financial Roadblocks	168
Capex Versus Opex	168
ROI	171
Policies and Culture	175

Technical Roadblocks	177
VM Sprawl	178
Application Roadblocks	181
Desktop Virtualization	187
Roadblocks to 100% Virtualization Summary	191

## **6 Full Case Study 193**

Customer Scenario	193
Planning and Designing	195
Implementation	212
Operating	213
Managing	215
Full Case Study Summary	216

## **A Additional Resources 217**

Chapter 2: Implementing the Solution	217
Chapter 3: Operating the Environment	219
Chapter 4: Managing the Environment	220
Chapter 5: Roadblocks to 100% Virtualization	221

## **Index 223**

*This page intentionally left blank*

# Preface

In our experience as consultants, VMware vSphere is the most robust virtualization solution on the market. The technology is proven and the user base is large.

Although the benefits of virtualization using vSphere are many, proper planning is required to gain these benefits from a vSphere infrastructure. This book is a guide to planning, designing, implementing, operating, and managing a robust vSphere infrastructure. The best practices and advice given in this book come from our own extensive field experience as datacenter architects and implementation engineers.

We wrote this book after noticing a need for the above mentioned areas. Many VMware books tell you how to set up and configure vSphere, but few talk about specific and real use cases. This book provides an in-depth discussion of business drivers and decisions around virtualization, an area that is not often covered.

## Authors' Disclaimer

Although we have made every effort possible, we assume no responsibility for errors or omissions. Neither is any liability assumed for damages resulting from the use of information contained in this book.

## You the Reader

This book is intended for systems administrators with experience using VMware's vSphere products. The products discussed in this book include vSphere, vCenter Operations, VMware Data Recovery, VMware View, VMware Site Recovery Manager, and other third-party additions, such as the Cisco Nexus 1000v distributed virtual switch. While we do expect this level of experience, we have noted resources for further research and learning in Appendix A where appropriate.

## What This Book Covers

Here is a quick rundown of what is covered in *Managing and Optimizing VMware vSphere Deployments*:

Chapter 1, "Laying the Groundwork"

This chapter talks about building a proper foundation for the virtual infrastructure. Many of the design considerations and best practices that show up in later chapters are mentioned here. These decisions will serve as the foundation for a design blueprint.

### Chapter 2, “Implementing the Solution”

This chapter talks about the considerations when implementing a vSphere-based solution. You learn the process of taking a design from the blueprint and bringing it to completion through the implementation process.

### Chapter 3, “Operating the Environment”

This chapter talks about operating a vSphere-based solution. It discusses many of the day-to-day tasks that are sometimes overlooked and provides some great community resources for aiding in these tasks.

### Chapter 4, “Managing the Environment”

This chapter talks about managing a vSphere-based solution after the implementation. It discusses both capacity and performance management. It also looks at how capacity planning and forecasting can be used on an ongoing basis as the infrastructure grows.

### Chapter 5, “Roadblocks to 100% Virtualization”

This chapter talks about the journey toward 100% virtualization. You learn the roadblocks that prevent many organizations from continuing virtualization initiatives and what can be done to help facilitate breaking through these stall points.

### Chapter 6, “Full Case Study”

This chapter brings what you have learned together in a customer case study. You explore a design from its inception to the implementation. You then explore the reasons for the design decisions along the way.

### Appendix A, “Additional Resources”

This appendix includes a list of resources and learning materials mentioned throughout the book to aid the reader in further research.

## Hyperlinks

When necessary, we have provided resources in the appendix of this book to Internet resources that aid or speak further to the content of this book. These resources can also be found online at the following URLs:

[www.seancrookston.com/publications](http://www.seancrookston.com/publications)

[www.harleystagner.com/publications](http://www.harleystagner.com/publications)

## About the Authors

**Sean Crookston** currently is a data center implementation engineer at TBL Networks, a VMware Enterprise Solutions Provider in Richmond, Virginia. Sean holds certifications from VMware, Cisco, EMC, and Microsoft. Throughout his career, Sean has engineered technology solutions for the healthcare, government, manufacturing, publishing/broadcast, and high-tech industries that streamline business processes and reduce operational expenses. Sean has been awarded the VMware vExpert award in 2010 and 2011.

**Harley Stagner** is a VMware Certified Design Expert (VCDX #46) and the first VCDX in Virginia. Harley is currently a datacenter design architect at TBL Networks, a VMware Enterprise Solutions Provider in Richmond, Virginia. Harley is focused on vSphere Architecture utilizing the latest datacenter technologies. Harley maintains a blog about these specialties at [harleystagner.com](http://harleystagner.com). A list of Harley's other publications can be found at [www.harleystagner.com/publications](http://www.harleystagner.com/publications).

## **Acknowledgments**

Before we begin this book, we would like to thank the many people who helped us along the way. First, we would like to thank our wives and families for their dedication to this endeavor. Without their patience throughout the writing process, this book would not be possible.

Second, we would like to thank the production and editorial team at VMware Press/Pearson and our technical editors, Russell Pope and Glenn Drawdy, for their dedication to this book.

Third, we would like to thank TBL Networks, Inc., for providing us use of its lab for testing and developing much of the content in this book.

## We Want to Hear from You!

As the reader of this book, *you* are our most important critic and commentator. We value your opinion and want to know what we're doing right, what we could do better, what areas you'd like to see us publish in, and any other words of wisdom you're willing to pass our way.

As an associate publisher for Pearson, I welcome your comments. You can email or write me directly to let me know what you did or didn't like about this book—as well as what we can do to make our books better.

*Please note that I cannot help you with technical problems related to the topic of this book. We do have a User Services group, however, where I will forward specific technical questions related to the book.*

When you write, please be sure to include this book's title and authors as well as your name, email address, and phone number. I will carefully review your comments and share them with the authors and editors who worked on the book.

Email: [VMwarePress@vmware.com](mailto:VMwarePress@vmware.com)

Mail: David Dusthimer  
Associate Publisher  
Pearson  
800 East 96th Street  
Indianapolis, IN 46240 USA

## Reader Services

Visit our website at [www.informit.com/title/9780321820471](http://www.informit.com/title/9780321820471) and register this book for convenient access to any updates, downloads, or errata that might be available for this book.

*This page intentionally left blank*

## Operating the Environment

This chapter focuses on maintaining and monitoring an active environment. At this point, you might or might not have designed an optimal environment. The environment also might not have been implemented to your standards. After all, sometimes you can't entirely fix what is currently broken and must deal with it for a period of time.

In the field, we see the excitement in customers' eyes at the power that VMware brings to their infrastructures. Cost savings through hardware, high availability, and ease of management are the main things they are eager to take advantage of. However, this excitement sometimes leads to a lack of focus on some of the new things that must be considered with a virtual infrastructure. A lack of maintenance and insufficient or no monitoring are two huge issues that must be considered. Before delving into maintaining and monitoring a virtual infrastructure, this chapter talks about some other operational items that you might not have considered in the design.

### **Backups**

A virtual infrastructure can pose different challenges for backups in terms of a technical understanding of the environment. This is the main reason we see that backups are not being adequately performed. Every organization has its own set of requirements for backups, but consider the following as important items for a backup strategy:

- An appropriate recovery point objective (RPO) or the ability to roll back to a period of time from today
- An appropriate retention policy, or the number of copies of previous periods of times retained

- An appropriate recovery time objective (RTO) or the ability to restore the appropriate backups in a set time
- An appropriate location of both onsite and offsite backups to enable recovery of data in the event of a complete disaster, while still allowing for a quick restore onsite where needed
- The ability to properly verify the validity of your backup infrastructure through regular testing and verification

Furthermore, outside of a technical understanding of the virtual infrastructure, virtualization poses no other significant challenges to maintaining a backup strategy. In fact, it will actually enable easier and quicker restores if properly designed.

When considering your backup strategy, you need to consider your RTO and RPO. You also need to consider your retention policy and proper offsite storage of backups. Properly storing offsite copies of backups is not just about keeping copies offsite that allow a quick restore to a recent restore point. It is also about considering what to also keep onsite so that simple restores are just that. Beyond that, you need to make sure you have all the small details that make up your infrastructure. This includes credentials, phone numbers for individuals and vendors, documentation, and redundancy in each of these contacts and documentation locations.

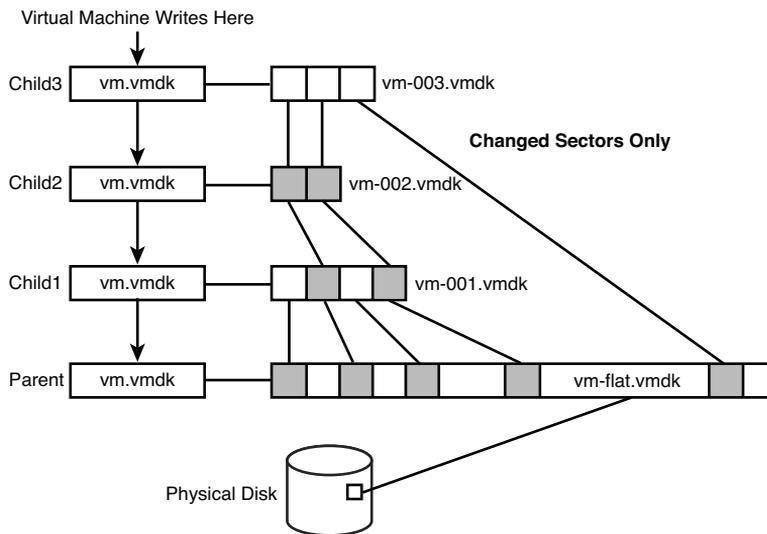
When considering backups, you need to determine the proper mix of file-level backups or virtual machine-level backups. Some organizations continue to do backups from within the guest that can provide a bare-metal restore. This is still a good option, and it might be your only option because of the software you presently use for backups; however, it will not be as quick to restore as a backup product that uses the VMware vStorage APIs to provide a complete virtual machine restore.

Let's take a moment to talk about the verification and monitoring of your backups. Taking backups is not the solution to the task of creating a backup strategy. The solution is the ability to restore the missing or corrupted data to a point in time and within a certain time as dictated by your businesses requirements. Therefore, it is always important to regularly test restoration practices and abilities as well as monitor for issues with backup jobs.

Your backup product should be able to verify the data was backed up and not corrupted; however, you should also schedule regular tests to verify this.

And, finally, let's talk about snapshots. Snapshots are not backups, but in some environments they are used in that fashion. Snapshots are useful when performing updates on a virtual machine as a means of quick rollback; however, they should not be used long term. We've witnessed two main things that occur as a result of snapshots being left behind.

For starters, they result in data needing to be written multiple times. If you have three snapshots, any new data is written to all three. As you can see in Figure 3.1, blocks of data that need to be written are written to each snapshot file, resulting in a performance hit as well as increased space utilization. Multiply this by several virtual machines and possibly even worse by multiple nested snapshots, and it is no wonder that we see datastores fill up because of old snapshots. This can bring virtual machines to their knees and makes rectifying the situation complex. When consolidating snapshots, you need to have space available to write the data to the original virtual machine disk. In this case, you would not have that available, requiring the migration of virtual machines to other datastores.



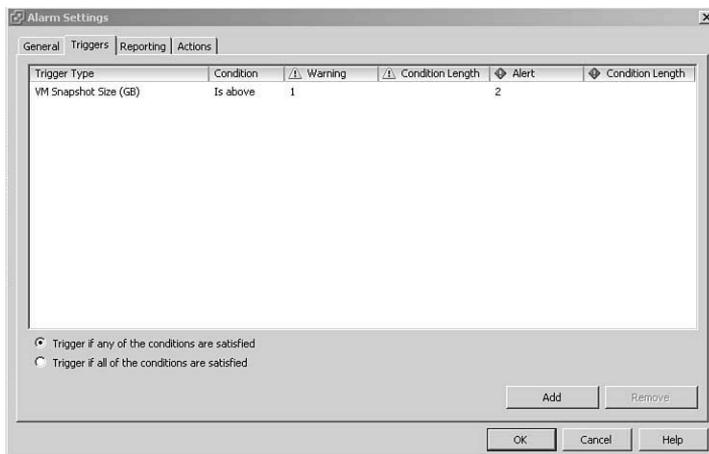
**Figure 3.1** Snapshot Disk Chain

A second problem we have seen many times is often caused by full datastores. Snapshot corruption can occur as a result, leading to the disappearance of any data since the time of the snapshot creation. For example, assume a single snapshot was taken six months ago, right after you installed Windows for your new Exchange server. If that snapshot is corrupted, you will likely be able to repoint to the original VMware Disk (VMDK) file; however, you'll be left with a bare Windows virtual machine. Full datastores are not the only time snapshots can be corrupted. This can also occur as a result of problems during snapshot consolidation or manipulating the original virtual machine disk file from the command line while snapshots are present.

It is important to note that a snapshot itself contains only the changes that occur after the snapshot was taken. If the original virtual machine disk is corrupted, you will lose all of your data. Snapshots are dependent on the virtual machine disk.

VMware's Knowledge Base (KB) article 1025279 discusses in detail the best practices when using snapshots. In general, we recommend using snapshots only as needed and for short periods of time. We recommend configuring alarms within vCenter to notify of snapshot creation and regularly checking for snapshots in your environment. There are many PowerShell scripts available that will accomplish this; however, a great tool to have that includes snapshot reporting is PowerGUI (see Appendix A, "Additional Resources," for reference).

Within vCenter, no default alarms exist to alarm for snapshots. You can, however, create a virtual machine alarm with the following trigger to alarm for snapshots, as shown in Figure 3.2. This will help you with snapshots that have been left behind for some time and have grown to 1GB or larger; however, it will not help until the total amount of snapshot data written for a virtual machine totals 1GB. This chapter discusses alarms later, but you can also check out VMware Knowledge Base article 1018029 for a detailed video demonstration of creating an alarm like this one (see Appendix A for a link).



**Figure 3.2** Configuring Snapshot Alarms

## Data Recovery

Like many products that use the VMware vStorage APIs, VMware's Data Recovery provides the ability to overcome backup windows. That is not to say you might not want to consider backup windows because you also must consider the traffic that will occur on the

network during backups; however, backup windows are of less concern for a few reasons. For starters, Data Recovery provides block-based deduplication and only copies the incremental changes. This occurs from a snapshot copy of the virtual machine that enables virtual machines to continue running while Data Recovery performs the backup from that snapshot copy.

Data Recovery is not going to be the end-all solution to your backup strategy, though. Its intention is to provide disk-based backup storage for your local storage and there is not a native method built in to transfer these backups to tape or other media. Therefore, VMware Data Recovery is best thought of as a complementing product to an existing backup infrastructure. With that said, let's talk about some of the capabilities the product has.

The process to get backups up and running is straightforward:

- Install Data Recovery.
- Define a shared repository location.
- Define a backup job.

### **Installing Data Recovery**

The first thing you need to verify is whether the product will meet your needs. Some of the more common things to consider when implementing Data Recovery are as follows:

- As previously mentioned, Data Recovery is intended to provide a quick method for onsite restores and does not provide offsite capabilities.
- Furthermore, you need to be sure all of your hosts are running ESX or ESXi 4.0 or later.
- Make note that each appliance supports 100 virtual machines with eight simultaneous backups. There is also a maximum of ten appliances per vCenter installation.
- The deduplication store requires a minimum of 10GB of free disk space. When using CIFS, the maximum supported size is 500GB. When using RDM or VMDK deduplication stores, the maximum supported size is 1TB.
- There is a maximum of two deduplication stores per backup appliance.
- Data Recovery will not protect machines with fault tolerance (FT) enabled or virtual machines disks that are marked as Independent.

For a complete list of supported configurations, refer to the *VMware Data Recovery Administration Guide*.

There are two steps to get the appliance installed. First, install the vSphere Client plug-in. Second, import the OVF, which will guide you through where you want to place the appliance. Once completed, you might want to consider adding an additional hard disk, which can be used to store backups.

### **Defining a Shared Repository**

As discussed, each appliance will be limited to two shared repositories and depending on the type of repository, you will be limited to either 500GB (CIFS) or 1TB (virtual hard disk or RDM). You have the following options when choosing to define a shared repository:

- Create an additional virtual hard drive (1TB or less).
- Create a CIFS repository (500GB or less).
- Use a RDM (1TB or less).

If you choose to create and attach an additional virtual hard disk, you need to consider where you are placing it. As mentioned previously, the intention of Data Recovery is to deliver the capability of a quick onsite restore. The use of virtual hard drives provides for the best possible performance. If you use a virtual hard disk, though, you will be storing the backups within the environment they are protecting, so you must consider this carefully. You could store the virtual hard disk on the plentiful local storage that may be present on one of the hosts. You could also store the virtual hard disk on any IP-based or Fibre Channel datastore.

Our recommendation in this case is to use the local storage of one of the hosts if it is available. When given the choice between the two, consider the likelihood of your shared storage failing versus the local storage of a server failing. Additionally, consider the repercussion of each of those failing. If your shared storage were to fail with the backups on them, you would have to use your other backup infrastructure to restore them, which can be quite time consuming. If the local server with your backups on them were to fail, then if a complete disaster occurs you are still going to have the production copies running on shared storage. If you do have a complete site failure, then you are going to need to deploy your disaster recovery strategy. This is discussed further shortly.

Another option is to use a Raw Device Mapping (RDM). If you are using the same storage as your virtual infrastructure, you are taking the same risks. The only way to mitigate such risks is to use storage dedicated for the purposes of backups. Just like the option of using virtual disks, think about where you are going to restore that data to if a disaster occurs. If your storage device is gone, you are going to have to initiate your disaster recovery strategy.

Another option is to use a CIFS share. Remember that CIFS shares are limited to 500GB, so each appliance can only support 1TB of CIFS repositories with its two-repository limit. Although the product lets you configure a CIFS share greater than 500GB, it warns you not to do so. We recommend that you listen to the warning because testing of the product has proven that creating a large CIFS repository can cause Data Recovery to fail to finish its integrity checking, which in turn causes backups to not run.

Another consideration for CIFS is that the share you are sharing out, and for that matter the disk that is being used, should not be used for any other function. Remember that Data Recovery provides for block-based data deduplication. If other data exists on the back-end disk, this can also cause a failure in integrity checking and, in turn, a failure of backup jobs running.

### **Defining a Backup Job**

Now that the appliance is set up and you have set up one or two repositories, it is time to create the backup jobs. Backup jobs entail choosing the following:

- Which virtual machines will be backed up
- The backup destination
- The backup window
- The retention policy

### **Choosing Which Virtual Machines to Back Up**

The virtual machines you choose to back up can be defined by an individual virtual machine level or from vCenter, datacenter, cluster, folder, or resource pool levels as well. Note that when you choose a virtual machine based on the entity it is in, if it is moved it will no longer be backed up by that job.

### **Choosing a Backup Destination**

Your choice of a destination might or might not matter based on the size of your infrastructure or your backup strategy. For sizing purposes, consider that you could exceed the capacity of the deduplication store if you put too many virtual machines on the same destination. For purposes of restoring data, consider the placement of the backups and where it is in your infrastructure.

### **Defining a Backup Window**

Backup windows dictate when the jobs are allowed to run; however, they do not have a direct correlation to the exact time they will execute. By default, jobs are set from 6:00

a.m. to 6:00 p.m. Monday through Friday and all day Saturday and Sunday. Consider staggering the jobs so that multiple jobs do not run simultaneously if you are concerned with network throughput.

### Defining a Retention Policy

When choosing a retention policy, you have the option of few, more, many, or custom. Custom allows specifying the retention of as many recent and older backups as required. The other options have their defaults set, as shown in Table 3.1.

**Table 3.1** VMware Data Recovery Retention Policies

Retention Policy	Recent Backups	Weekly	Monthly	Quarterly	Yearly
Few	7	4	3	0	0
More	7	8	6	4	1
Many	15	8	3	8	3

Changing any one of the settings for these policies will result in the use of a custom policy. When choosing your retention policy, ensure you have the capability to restore data from as far back as you need, but within the confines of the storage you have to use for backups.

At this point, your backups are up and running. You can either initiate a backup now or wait until the backup window has been entered for backups to begin. Once you've seen your first successful backup, you still have a few other items to consider.

### Restoring Data (Full, File, Disks) Verification

When restoring data, you have two key things to consider. When choosing to restore data, you first need to choose your source. A virtual machine can be part of multiple backup jobs, so in addition to having a different set of restore points, you might also have a set of restore points that are also located on a different backup repository. Second, you need to consider where you want to restore the data.

For the purposes of testing the capability to restore, you can perform a restore rehearsal by doing the following from within the Data Recovery interface by right-clicking a virtual machine and then clicking the Restore Rehearsal from Last Backup option. To fully test a restore or to perform an actual restore, you have much more to consider because this option chooses the most recent restore and restores the virtual machine without networking attached. The following sections discuss those considerations further.

## Choosing Backup Source

When restoring, you have the option to restore at any level in the tree, so you can restore entire clusters, datacenters, folders, resource pools, or everything under a vCenter server. When looking at the restore of an individual virtual machine, you can restore the entire virtual machine or just specific virtual disks. You may also restore individual virtual machines from the virtual machine backup, which is discussed shortly.

## Choosing Restore Destination

When restoring, you have several options during the restore, including choosing where to restore the data. When considering restoring an entire virtual machine, you have the following options to consider:

- Restore the VM to a specific datastore.
- Restore the virtual disk(s) to a specific datastore(s).
- Restore the virtual disk(s) and attach to another virtual machine.
- Choose the Virtual Disk Node.
- Restore the VM configuration (yes/no).
- Reconnect the NIC (yes/no).

When restoring, the default setting is to restore the virtual disk in place, so be careful to consider this if it is your intended result. If possible, in all situations you should restore to another location to retain the set of files that is currently in place if further restore efforts are needed on those files.

## File Level Restores

In addition to restoring complete virtual machines or specific disks, you may also restore individual files. File Level Restore (FLR) allows for individual file restoration with an in-guest installed software component. The FLR client is available for both Windows and Linux guests and must be copied off the Data Recovery media locally where it will run. By default, Data Recovery only allows the restoration for files from a virtual machine for which the client is being run; however, if you run the client in Advanced mode, you can restore files from any of the virtual machines being backed up. Note that although you are able to mount Linux or Windows virtual machines regardless of the operating system you are running, you might not be able to read the volumes themselves.

Once mounted, you have the ability to copy files and restore them to locations manually or look through them to find the version you are looking for. The mounted copies are

read-only versions of the files, and any changes made will not be saved, so make sure to copy the files to a desired location before making any changes.

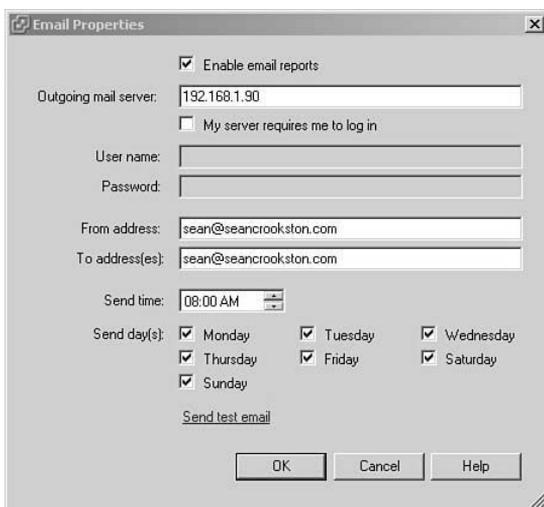
One last note on the use of FLR when using Data Recovery: It is not recommended and Data Recovery should be configured so that File Level Restores are not possible. This is done by configuring the VMware Data Recovery .ini file and setting EnableFileRestore to 0.

### Site Disaster

As mentioned previously, the intended use of this product is for quick restores and is not intended to be your disaster recovery plan. If you were to lose a vCenter server and needed to recover another machine, you would have to stand up a new vCenter server and install the plug-in to use Data Recovery to restore the virtual machine. Additionally, if you lose the appliance itself, you must install a new one and import the repository. Be aware that this can take a long time if a full integrity check is required.

### Monitoring Backup Jobs

Data Recovery allows the configuration of an email notification that can be sent as often as once a day at a specified time. There isn't much to configure with email notification, as shown in Figure 3.3. The important thing is to make sure the appropriate individuals are being notified and that mail is being relayed from the outgoing mail server specified. Remember the server that needs to be authorized is not the vCenter server but rather the Data Recovery appliance itself.



**Figure 3.3** Configuring Data Recovery Email Notifications

## Managing the Data Recovery Repository

The maintenance tasks that run will check the integrity of the data in the repositories and reclaim space in the deduplication stores. By default, Data Recovery is set to be able to run maintenance at any time. This might not be a problem for your environment; however, when integrity check operations are occurring, backups cannot. Therefore, you should change the maintenance window so that it is set to run during a specified period of time. This ensures backups will always have the time to run each day.

When the deduplication store is using less than 80% of the repository, the retention policy is checked weekly to remove any restore points that are outside the specifications. This means that you might have many more restore points than expected as a result. Once 80% of the repository is utilized, the retention policy is checked daily. In the case of the repository filling up, the retention policy is run immediately if it has not been executed in the last 12 hours.

## Disaster Recovery

Disaster recovery is an area of many organizations that has at least some, if not a lot, of room for improvement. When looking at a disaster recovery plan, the following things are important to consider:

- Data is available with an RTO that meets the business's requirements to operate and the data is from a point in time that meets the RPO of the organization.
- Data has been verified as being valid.
- A runbook has been defined for how to and in what order to restore.

The first point is present in most organizations, whereas the second and third are not. It should not be surprising because a failure warranting declaring a disaster is not often needed. Nonetheless, a solid runbook should be defined for your infrastructure. A runbook for restoration for your virtual infrastructure is crucial; however, consider the back-end networking infrastructure first as your virtual machines will be of no use without it.

When looking at recovering your virtual infrastructure, the ideal setup is to replicate among your storage devices and use VMware's Site Recovery Manager (SRM) to automate your restore. Site Recovery Manager is further discussed later in this chapter; however, for those not familiar, it assists in automating the recovery of virtual machine environments during a disaster.

You may also use the set of replicated data to manually configure the virtual machines and power them on at your disaster recovery location. Additionally, you can choose another method of manual restoration. This could be using a copy of the virtual machine files

from some other mechanism or using a backup product to perform a bare-metal copy of the machine and restoring it to a newly configured virtual machine. For the purposes of this discussion, we talk in detail about the use of VMware's Site Recovery Manager as it provides the best mechanism. Before doing that, though, the following sections talk briefly about the other options.

## Manual Disaster Recovery

When looking at implementing a manual data recovery plan, you need to ensure you are doing a few things that Site Recovery Manager would be automatically handling or assisting with. Many times, the use of manual methods is the result of a lack of sponsorship of the initiative in terms of funding; however, that does not mean the process cannot work. If you are creating a manual data recovery plan consider the following.

- Ensure data is being replicated/copied and is current with your RPO.
- Ensure your processes for restoration meet your RTO.
- Ensure the Disaster Recovery (DR) site hardware is supported and will support the load in the event of a disaster.
- Ensure the recovery processes work by performing regular DR tests.
- Ensure the runbook is updated regularly as network, application, and other requirements change.

By keeping these points in mind, your disaster recovery efforts will be successful; however, you will have to perform many of the steps manually.

Whereas storage replication was previously a condition for using Site Recovery Manager, the latest version now supports host-based replication. If you were previously unable to use Site Recovery Manager because of the storage replication requirement, you should reevaluate the product with host-based replication.

## Site Recovery Manager

VMware offers a product called Site Recovery Manager that helps automate most of the process of recovering virtual machines during a disaster. The product allows for isolated testing to ensure recovery is possible in the event of an actual disaster as well as the ability to failback in version 5.0.

When installed at both the production and disaster recovery locations, the product provides for a centralized approach to defining replication and recovery plans. In prior versions, SRM relied on the storage itself to perform the replications and integrated with

the storage using a supported Storage Replication Adapter (SRA). This limited the product for some entities with supported storage. Even those with supported storage devices in both locations might not have had matching storage solutions and, hence, no supported replication infrastructure in place.

SRM 5.0, however, has expanded its market base with the introduction of vSphere Replication (VR). This allows replication from one location to another, regardless of the type of storage on both ends. One or both ends can even be local or directly attached storage. SRM is also protocol independent so you can replicate among Fibre Channel, iSCSI, or NFS storage.

For more information on Site Recovery Manager, check out *Administering VMware Site Recovery Manager 5.0* by Mike Laverick. This book provides an in-depth discussion of the product, using it in a number of scenarios, and is a great read when defining a disaster recovery solution in a virtualized environment.

## Physical to Virtual Conversions

When a virtualization infrastructure is implemented, the first virtual machine installed is typically going to be brand-new installations of Windows for your vCenter and SQL servers. Many times, a project like this also serves as a good time to clean up and move toward the latest server operating systems. Regardless at some point, you need to begin moving some of the existing physical workloads over and retaining their configurations.

VMware provides a free download for a product that will assist in this migration. VMware vCenter Converter Standalone 5.0 allows the conversion of physical systems as well as systems that are already virtualized. When performing physical to virtual conversions, you should be aware of the following things.

For all systems, you should do the following:

- **Prior to conversion**

- Perform a survey of the server and its applications.
- Identify server and application owners for approval and verification testing upon completion.
- Identify performance and configuration of CPU/memory/disk versus actual usage.
- Identify destination for virtual machine (host placement and storage placement).

- Identify and record network configurations.
- Identify downtime and schedule for system(s).
- **During conversion**
  - Place virtual machine on host and storage per design.
  - Adjust configurations of CPU/memory/disk as appropriate.
- **After conversion**
  - Remove nonpresent devices.
  - Remove legacy software.
  - Install VMware Tools.
  - Reconfigure the network.
  - Perform basic testing.
  - Verify functionality with server and application owners.
  - Fully uncable and remove decommissioned physical servers.

## Issues and Troubleshooting

Physical to virtual conversions can fail to start. Typically, we have found this is because of one of the following reasons:

- No Permissions Admin\$
- Firewall exists between server to be converted and vCenter
- Incorrect DNS configurations

In some cases, administrators have removed the admin\$ share on a server, which is required for the vCenter Converter agent installation when installing remotely. You can install the client locally or re-create the share to resolve these issues.

When a firewall exists, it can cause a failure if certain ports are not reachable. VMware Knowledge Base article 1010056 details the required ports to be allowed through the firewall (see Appendix A for a link).

If DNS configurations are not correct on a source virtual machine, this can also cause failures. Ensure DNS is correct or update DNS so that the system to be converted can reach the vCenter server.

Besides these basic considerations, you should also consider a few special cases:

- Domain controllers
- Windows Server with OEM installations
- SQL, Exchange, and other applications servers
- Linux
- Virtual to virtual conversions (V2V)

## Domain Controllers

Active Directory domain controllers have special considerations to take when looking at physical to virtual conversions. Although there are methods to perform a P2V conversion, it is our recommendation you don't and instead create new virtual servers. Domain controllers are extremely sensitive to hardware changes and a failure in following the P2V process at any step can cause major replication issues that will be visible throughout your infrastructure. Creating new servers allows for a clean and safer migration. The recommended process is as follows. Note this may vary depending on how many Active Directory domain controllers you have, their location, and whether they remain physical or not.

- Create one or more brand-new Windows virtual machines.
- Promote these two domain controllers using `dcpromo`.
- Ensure replication via the command line using the `repadmin` tool.
- Transfer FSMO roles.
- Transfer DHCP servers.
- Verify role transfer.
- If the domain controller is also a DNS server, ensure DNS is running on new systems.
- Power down physical domain controller and verify functionality. If no issues exist, power back on and demote existing physical systems using `dcpromo` (FSMO role will be transferred there also).
- Decommission physical hosts.
- Reconfigure DNS settings for servers/workstations to reflect any new IP addresses.

## Windows Server with OEM Installations

Another case for consideration is any physical Windows servers that were originally installed with OEM media. Per Microsoft's licensing agreement, these licenses are tied to the original physical hardware and as a result the right to continue using the installed operating system does not carry over. Microsoft licensing is outside the scope of this book; however, two things can be drawn from this situation:

1. You are not in compliance with your licensing. You need to have or purchase an additional Microsoft server license.
2. You are running a version of Windows you are not licensed for. You need to either install a fresh instance of Windows using volume licensed media or perform an in-place upgrade using the volume licensed media.

If you have ignored these recommendations, you will find that once you bring the newly converted virtual machine online, it will require activation. With OEM installation, the key that is required to activate is not necessarily the one that was on the sticker on the box. With that said, you might be able to activate the software installation; however, you need to ensure you are in compliance with Microsoft licensing as soon as possible thereafter.

Although vendor background screens are typically a clear indication that OEM installations are present, you can also check the product ID for OEM to verify. If it is not present, you are fine. You can script this to check for multiple servers in your environments using various product and key software products that reveal this information. Additionally, you can use code like this PowerShell snippet to gather this information. The following code sets `$Prod_ID` to the `ProductID` of the machine it is run on:

```
$Prod_ID=(get-item 'HKLM:\Software\Microsoft\Windows NT\  
CurrentVersion').getvalue('ProductID')
```

## SQL, Exchange, and Other Applications Servers

Although vCenter Converter will make multiple passes through your data, it is important to consider what could be lost in the transition time in between bringing the physical host down and the last time the data was copied. As a result, it is recommended that you stop all application services.

In addition, consider file shares. You may declare in an email that the server will be down this Saturday, but that doesn't stop someone from changing data. This data could end up being changed at a time of transition and would in effect be lost. Therefore, it is recommended that you unshare any file shares during the conversion process. Additionally, you should remove any temporary files or any data that is no longer needed. This greatly

decreases the time involved when converting and optimizes the amount of physical storage being used.

## Linux

This chapter has talked a lot about virtualizing Windows servers, but now the focus shifts to Linux servers. Linux servers follow a slightly different process and, as a result, there are some things you should be aware of. For starters, the process for Linux does not deploy an agent, but instead a helper virtual machine is deployed on the destination vSphere host. This helper machine will ultimately become the production virtual machine once it has copied all of the data from the physical Linux machine.

You should consider a few important things:

- You must ensure you have SSH access to the Linux machine when doing an online conversion and you must have root access when doing so.
- Only certain flavors of Linux are supported for online conversions. Presently, these are certain versions of Red Hat, SUSE, and Ubuntu.
- Customization during the conversion process is not supported for Linux guest operating systems.

## V2V and Other Methods

Physical machines are not the only machines that may need to be brought into a new infrastructure. For example, you may have existing virtual machines running on storage not accessible to the new environment. You may also have these virtual machines running in a Hyper-V environment. Additionally, you may choose to do an offline conversion by using an imaging product. Regardless of the source, vCenter Converter allows for all of these options when using any of the following supported methods.

The following virtual machine formats are supported for cold conversion:

- Microsoft Hyper-V
- Microsoft Virtual Server
- Microsoft Virtual PC
- Parallels Desktop
- VMware Workstation, GSX Server, Player, Server, Fusion, ESX

The following image formats are supported:

- Symantec Backup Exec System Recovery
- Norton Ghost
- Acronis
- StorageCraft

Additionally, during the conversion process, you can convert any running Windows virtual machine by specifying a powered-on machine as the source.

### **Offline Boot Disc**

When an offline conversion is desired in addition to the mentioned image formats, you can use the VMware vCenter Converter Boot CD. The boot CD is no longer provided as of vCenter Converter 4.3; however, it is still available for download with valid support for a vSphere 4.x Enterprise Edition or greater license. At the time of this writing, there are no current plans to release a vCenter Converter Boot CD for vSphere 5.0; however, version 4.0.1 build 16134 is the latest version and is supported for conversions from a source to a vCenter 5 infrastructure.

The offline boot disc is based on Windows PE and allows the import of network drivers to build a new image if required. The lack of network drivers is the most common reason the offline boot disc does not work.

## **Maintenance**

Maintaining a vSphere-based virtual infrastructure is very important. After all, you have a large number of operating systems now running collectively on a much lesser amount of physical hardware in most cases. A failure to update for and then be exposed to a potential flaw may now put your entire infrastructure at risk instead of only some servers.

Why do organizations not properly maintain their vSphere environments? Everyone agrees with the criticality of maintaining servers whether it is through patches or regular release updates, but still it remains a large problem in many environments. In large part, the main driving force to perform any update is a result of an enhancement release that has added additional features.

## Update Manager

One reason many administrators do not update their infrastructures is due to a lack of understanding of the process. Maybe they are new to VMware and never bothered to even install Update Manager with vCenter. Update Manager is not a requirement to patch systems but the process does become much more involved when using the command-line interface to do so. An administrator must download the update bundle and transfer it to each of the hosts. Then a command-line process must be invoked from each of the hosts. In the days before Update Manager, it is no wonder why some administrators might have chosen to patch less frequently or not at all.

vSphere is a hardened hypervisor and, as a result, needs much less patching and updating for vulnerabilities than a typical operating system. Many administrators, though, take this as a reason not to patch at all.

Some also entirely understand the advantages of Update Manager and have it installed and running. They realize how the effect of an issue with their vSphere infrastructures could now affect all their operating systems instead of just a handful. As a result, they view this increased impact of any updates as possibly negative. This may be the proper viewpoint as certain vulnerabilities may not be a high risk for their environments. They are further justified in their decision in knowing that the impact of any issues that occur in a virtual infrastructure can be huge if not properly planned. Perhaps the feature that is affected is also not something they are using. Being cautious and properly planning and testing for updates is certainly the way to go. To date, I have never worked directly with anyone who has been exploited by a VMware vulnerability. This is a true testament to the ability to harden the hypervisor and keep ahead of the curve with security exploits.

Again, that does not justify not patching. With the ever-increasing deployments of vSphere, it seems pretty reasonable to think the focus will continue to shift toward attacking these consolidated infrastructures powered by VMware. After all, wouldn't it be easier to bring down 10 vSphere hosts running 200 servers than to try to bring them down individually?

Update Manager is a patch-management solution provided by VMware with all versions of vCenter Server. It helps to automate the deployment of patches and updates and provides a means to maintain compliancy among your entire infrastructure. Its capabilities are not just limited to vSphere ESXi hosts either, as you can now patch many virtual appliances as well as extensions such as the Cisco Nexus 1000V and EMC PowerPath.

Formerly, Update Manager was capable of remediating Windows guest virtual machines by providing operating system and application patches. As of vSphere 5, however, this capability is no longer included. Interestingly, they licensed the technology from Shavlik,

which they recently acquired. VMware now offers several other products that offer comparable capabilities. vCenter Protect Essentials and VMware Go both offer abilities to patch and manage guest operating systems. vCenter Protect Essentials provides for an on-premise solution that will patch and manage virtual machines. In the case of VMware Go, it is a cloud-based solution that also offers capabilities for help desk end-user portals. Both products also provide asset and configuration management capabilities that are geared toward the small to medium business market.

A major selling point of utilizing Update Manager to patch your vSphere servers is that when set up and used properly, it requires zero downtime to any of your virtual machines. There is no need to worry about having to have downtime twice for virtual machines for both the vSphere and guest patching. Utilizing DRS in conjunction with Maintenance mode, an administrator can deploy patches to a host with zero downtime to any of the virtual machines in the entire infrastructure.

Update Manager also allows the scheduling of updates. Simply create or attach a baseline to a set of hosts and choose a date and time to run the updates. These baselines can be assigned at the vCenter level or at the datacenter, cluster, or host level. Another useful ability of Update Manager is to stage and schedule virtual machine hardware and tools upgrades.

### **Patching Hosts Using Update Manager**

With a DRS-enabled cluster and the use of Maintenance mode, patching hosts using Update Manager is a straightforward process; however, the following key areas are often overlooked:

- **Sizing the patch repository**—The patch repository can become quite large depending on the versions of vSphere you choose to implement over time. As a result, it is best practice to configure a shared repository outside of the vCenter server or server where Update Manager is installed when separated. VMware offers the vSphere Update Manager Sizing Estimator for download, which will aid in sizing not only the shared repository, but also the database itself.
- **Notification of new patches**—You will have a hard time knowing when to install updates if you do not know when they come out. The easiest way to be notified of new patches is by configuring email notifications under the Download Schedule of Update Manager.
- **Failure to consider compatibility and support**—There is a lot to consider when choosing to install updates. If you are running a solution where the vendor will only support virtual machines on a certain revision of the software, then you should clarify how these support policies are affected by updates. This is a rarity these days as

solutions such as Cisco's unified communications on top of UCS software are fully supported by all updates at the time of release.

- **Failure to disable HA during an update**—If you have a smaller cluster of hosts, you might run into a failure if you do not disable HA during a host update. By default, this is not set but can be if you are going to run into this issue. Without doing so, if your cluster cannot support HA and you attempt a remediation, it will fail.
- **Failure to properly configure DNS**—If DNS is not properly configured, you will spend a lot of time troubleshooting why Update Manager is not working. It is highly dependent on DNS to be configured properly on both the vCenter and vSphere hosts. Failing to do so causes Update Manager to fail during the remediation.

## Upgrading Hosts

VMware periodically releases new versions of vSphere that require an upgrade to vSphere. If an environment is healthy and no issues exist, we recommend using Upgrade Manager to upgrade the hosts in place. If, however, there are issues with your environment, consider wiping away each host and starting fresh.

You should also consider downtime in your environment for the upgrade. If your virtual machines are on shared storage backed by hosts that can vMotion among one another, you will be able to have much less downtime than an environment with virtual machines on local storage, for example. Different circumstances warrant different paths, so let's talk about some of the key items to consider when planning your upgrade.

### Planning for vSphere Upgrades

Planning for vSphere upgrades requires investigating your environment from top to bottom to ensure you are presently free of any issues and have the appropriate pieces to perform a successful upgrade of your environment. In fact, before an upgrade is the perfect time for a health check to be performed by a VMware authorized partner. Although this section describes the important steps to consider, you might also want to take a look at the *vSphere Upgrade Guide* provided by VMware.

### *Upgrade Entitlement*

Before you get too far along, you need to ensure you are eligible to upgrade. Upgrades are not at an additional cost when you have a valid support contract with VMware for your purchased licenses. If you have an eligible support contract, you can find both the software and licenses available in your VMware software and licensing portals. If you don't have an eligible support contract, you need to either renew your support or purchase additional

licenses. In addition to support for vSphere, this is also a good time to ensure your hardware is still supported by your vendor before proceeding with an upgrade.

### *Feature Changes*

Another consideration is changes that might have occurred between the old and new versions of vSphere. An example of this is an organization using Update Manager as part of vSphere 4 to patch its Windows guest. This functionality is no longer included as of vSphere 5; however, it can be acquired as part of vCenter Protect Essentials or Essentials Plus.

### *Hardware Compatibility*

If your hardware is older, there is a chance that the hosts, storage, or IO devices might not be supported with the new release. Regardless of how new your equipment is, you should reference the *VMware Compatibility Guide* online to ensure the hardware will be supported after an upgrade. Although some people might not be concerned with hardware being fully supported, they should be advised that if it is not supported, there is a chance it will not work in some fashion. Be sure to check compatibility for your specific host. You will find that there may be several versions and revisions for popular brand models. Additionally, be sure to check your I/O devices and storage as well.

### *Database Compatibility*

You need to be sure your database is supported with the new version of vCenter to which you will be upgrading. Many versions of SQL 2000 and 2005 are no longer supported despite their use, and you should consider upgrading the database servers if yours are not supported. Check out the VMware Product Interoperability Matrix online to verify your database software support before proceeding with any vCenter upgrades.

### *vCenter Support*

You also need to ensure your vCenter server will support any older versions for hosts that are going to run for any period of time on an older version as part of that vCenter server. You can again check the VMware Product Interoperability Matrix online to verify this information. Pay close attention to the matrix, as shown in Figure 3.4. At the time of this writing, there is a known issue with VMware 4.0 U2 that does not allow it to be managed by a VMware vCenter server. This is a good example where a lot of people continue to make assumptions of the support to find out later that it does not work correctly.

<input type="checkbox"/>	Platform	VMware vCenter Server 5.0	VMware vCenter Server 4.1 U2	VMware vCenter Server 4.1 U1	VMware vCenter Server 4.1	VMware vCenter Server 4.0 U4	VMware vCenter Server 4.0 U3	VMware vCenter Server 4.0 U2	VMware vCenter Server 4.0 U1	VMware vCenter Server 4.0	VMware vCenter Server 2.5 U6
<input type="checkbox"/>	VMware ESXi 5.0	<input checked="" type="checkbox"/>									
<input type="checkbox"/>	VMware ESX/ESXi 4.1 U2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>						
<input type="checkbox"/>	VMware ESX/ESXi 4.1 U1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>						
<input type="checkbox"/>	VMware ESX/ESXi 4.1		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>						
<input type="checkbox"/>	VMware ESX/ESXi 4.0 U4	<input checked="" type="checkbox"/>									
<input type="checkbox"/>	VMware ESX/ESXi 4.0 U3	<input checked="" type="checkbox"/>									
<input type="checkbox"/>	VMware ESX/ESXi 4.0 U2		<input checked="" type="checkbox"/>								
<input type="checkbox"/>	VMware ESX/ESXi 4.0 U1	<input checked="" type="checkbox"/>									
<input type="checkbox"/>	VMware ESX/ESXi 4.0	<input checked="" type="checkbox"/>									
<input type="checkbox"/>	VMware ESX/ESXi 3.5 U5	<input checked="" type="checkbox"/>									
<input type="checkbox"/>	VMware ESX/ESXi 3.0.3 U1		<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>				

**Figure 3.4** vCenter Compatibility and Support Matrix

Additionally, you need to make sure vCenter is installed on a 64-bit operating system. If your existing vCenter server is older, you might not be able to directly upgrade anyway; however, if it is installed on a 32-bit operating system, you definitely need to install a fresh operating system.

In addition to support, you must also see if an upgrade is possible. As shown in Figure 3.5, you can see that, in general, there is direct upgrade available from 4.0 U1 up to vCenter 5.0 with the exception of 4.0 U4. Both 4.0 and 4.0 U4, and even 2.4 U6, however, can be upgraded to 4.1 U2. Once at 4.1 U2, they can be updated to vCenter 5.0 directly. Always check the VMware Product Interoperability Matrixes and *vSphere Upgrade Guide* for the most up-to-date support information. Again, remember there is a 64-bit requirement, so if you don't have a 64-bit server, you need to install a new version of Windows to support your new vCenter Server installation.

<input type="checkbox"/>	Platform	VMware vCenter Server 5.0	VMware vCenter Server 4.1 U2	VMware vCenter Server 4.1 U1	VMware vCenter Server 4.1	VMware vCenter Server 4.0 U4	VMware vCenter Server 4.0 U3	VMware vCenter Server 4.0 U2	VMware vCenter Server 4.0 U1	VMware vCenter Server 4.0	VMware vCenter Server 2.5 U6
<input type="checkbox"/>	VMware ESXi 5.0	☑									
<input type="checkbox"/>	VMware ESX/ESXi 4.1 U2	☑	☑	☑	☑						
<input type="checkbox"/>	VMware ESX/ESXi 4.1 U1	☑	☑	☑	☑						
<input type="checkbox"/>	VMware ESX/ESXi 4.1		☑	☑	☑						
<input type="checkbox"/>	VMware ESX/ESXi 4.0 U4	☑	☑	☑	☑	☑	☑	☑	☑	☑	
<input type="checkbox"/>	VMware ESX/ESXi 4.0 U3	☑	☑	☑	☑	☑	☑	☑	☑	☑	
<input type="checkbox"/>	VMware ESX/ESXi 4.0 U2		☑	☑	☑	☑	☑	☑	☑	☑	
<input type="checkbox"/>	VMware ESX/ESXi 4.0 U1	☑	☑	☑	☑	☑	☑	☑	☑	☑	
<input type="checkbox"/>	VMware ESX/ESXi 4.0	☑	☑	☑	☑	☑	☑	☑	☑	☑	
<input type="checkbox"/>	VMware ESX/ESXi 3.5 U5	☑	☑	☑	☑	☑	☑	☑	☑	☑	☑
<input type="checkbox"/>	VMware ESX/ESXi 3.0.3 U1		☑	☑	☑	☑	☑				☑

**Figure 3.5** vCenter Upgrade Compatibility and Support Matrix

### Dependencies

Outside of the core functionality in the vCenter server and the vSphere hosts, there exist some other pieces that need consideration as well. These are just some examples and you need to also consider any additional software or plug-ins that are used in your environment. Make sure to consider these pieces by verifying support by the vendor or within the *vSphere Upgrade Guide*:

- vCenter Update Manager
- vCenter License Server
- VMware View
- VMware Data Recovery
- Site Recovery Manager
- Third-party plug-ins like PowerPath

- Use of Nexus 1000V
- Any PowerShell or other scripting used for troubleshooting and reporting

### *Upgrade Paths*

vSphere 5 is the first version of vSphere that has been released in only the ESXi flavor, so there is only one destination when upgrading to vSphere 5. You must also consider the source of the server and whether you have the option to upgrade.

The following is true about upgrading older versions of ESX and ESXi to vSphere 5.0. Note there are conditions where these items might not apply, so be sure to check the VMware Product Interoperability Matrixes and *vSphere Upgrade Guide* for the most up-to-date support information.

#### **ESX & ESXi 3.5**

- No direct upgrade available
- Upgrade to 4.x first
- Note that the partition layout might be incompatible with vSphere 5, so this can prohibit such an upgrade to 5.0

#### **ESX & ESXi 4.0**

- Direct upgrade available with Update Manager, interactively, or scripted
- Might not be compatible with all environments
- For example, ESX 4 hosts on SAN/SSD might not have optimal partitions and disks with multiple VMFS partitions cannot be upgraded
- Additionally, note that a host with any third-party vSphere Installation Bundles (VIB) may require using the ESXi Image Builder CLI to create a customized ESXi install ISO

And one last note on upgrading hosts. As of vSphere 5, the advanced version no longer exists and any customers with active support agreements for vSphere 4 Advanced are entitled to vSphere Enterprise.

### *Order of Operations*

When laying out your plan for an upgrade, you must consider the order in which you are going to do so. Outside of the vCenter and vSphere hosts themselves, you need to make sure you upgrade to supported code and firmware for your storage and other devices ahead of time. Additionally, be sure you have proper backups of the necessary components. For

vCenter, you need at minimum a backup copy of the database as well as Secure Socket Layer (SSL) certificates from the server. For the hosts themselves, you need to have good documentation on their configuration as well as a backup copy of all virtual machines. This holds especially true if you are upgrading a host with virtual machines running on local storage. For virtual machines on shared storage, you need to ensure backups exist as you will be upgrading our virtual hardware and VMware Tools later on.

In general, follow these steps to perform an upgrade:

1. Run the vCenter Host Agent Pre-Upgrade Checker. This can be found on the vCenter installation media and is a great verification tool to ensure the likelihood of a successful upgrade.
2. Upgrade or install a new vCenter server.
3. Upgrade or install a new Update Manager.
4. Upgrade or install other plug-ins and third-party packages.
5. Upgrade or install vSphere on hosts.
6. Upgrade VMFS.
7. Upgrade virtual machine tools and hardware.

### Methods for Upgrading vSphere

As discussed previously, to perform your vCenter upgrade, you can either upgrade the software in place if supported or install a fresh vCenter server. You can then choose to either start completely fresh, redefine roles and other vCenter configurations, or import the database and continue from there.

For vSphere hosts, you not only have the option of upgrading or starting fresh, but you also have several methods to perform the upgrade. When possible, we recommend building new hosts and bringing configurations over.

In previous versions of vSphere, the Host Update Utility was included on the vCenter installation media for performing host upgrades on a host-by-host basis. Note that this is no longer the case and you must upgrade your hosts by either using vSphere 5 media or through Update Manager.

### *Manual Upgrade*

You may manually perform an upgrade to a host using the ESXi installation media by performing an interactive or scripted upgrade. It is recommended you disconnect all storage from the host as this greatly reduces the amount of time required for the upgrade.

When upgrading a host, you have three options:

- Upgrade ESXi, Preserve VMFS Datastore or Force Migrate ESXi, Preserve VMFS Datastore
- Install ESXi, Preserve VMFS Datastore
- Install ESXi, Overwrite VMFS Datastore

The first option will vary if any custom VIBs are not included with the vSphere 5 media. If that is the case, Force Migrate ESXi replaces Upgrade ESXi. Make sure to back up any items on the local VMFS datastore beforehand and especially when choosing to overwrite the VMFS datastore.

In addition to performing an interactive upgrade, you may also choose to perform a scripted installation. For full details on creating a scripted installation, including adding custom drivers and third-party VIBs, check out the *vSphere Upgrade Guide*.

### *Update Manager*

When using Update Manager to upgrade hosts, an orchestrated host upgrade can occur that allows not only for vSphere host installation, but the installation of VMware Tools and the upgrade of virtual hardware.

Update Manager does have some limitations that you may encounter. Recall from the earlier discussion of upgrade paths that there are some limitations even when following a supported path. Update Manager cannot be used to upgrade an ESX 4.x host if it was previously upgraded from 3.x as a result of insufficient space in the /boot partition. This problem is not unique as it is possible an ESX 4.x host may also not have the proper amount of space.

If you are not installing a fresh version of vSphere, it is recommended to use Update Manager because it greatly eases the upgrade process. The use of Update Manager does a better job of preventing erroneous actions and disallows things such as upgrading the virtual machine hardware before installing VMware Tools.

**Host Upgrades**    Upgrading a host requires the creation of a host upgrade baseline. Additionally, you are required to import the ESXi image to be used for upgrades.

You may choose to have separate baselines and separate images in the repository. For example, you may have different images based on hardware for the hosts, which may be of different vendor types and contain different third-party VIBs.

You cannot roll back to the previous version of ESX/ESXi when upgrading with Update Manager, so, as always, make sure you have the configuration of your host documented and the proper backups of all virtual machines in place before proceeding with any upgrades.

**Virtual Machine Upgrades**    Upgrading virtual machines after an upgrade requires using an existing baseline or the creation of a baseline group. You cannot upgrade VM hardware until the virtual machine is running the latest version of VMware Tools. Update Manager makes sure this happens to avoid these operations not occurring in sequence.

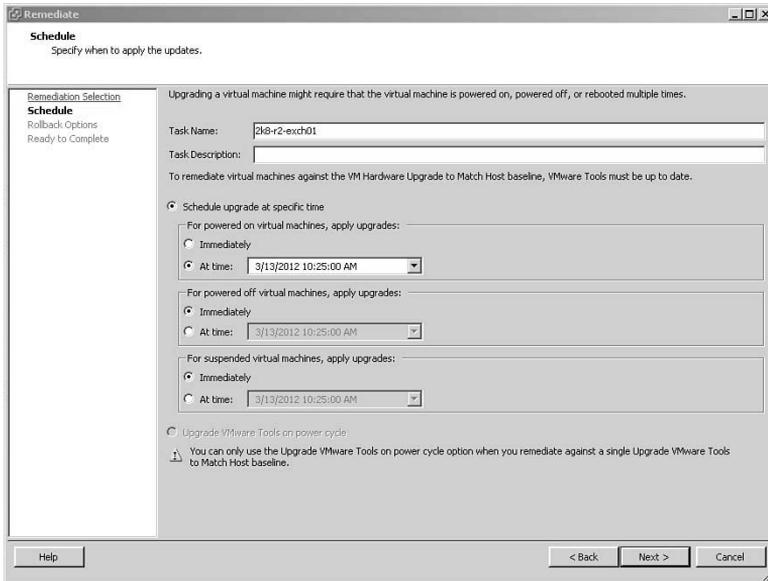
By default, the following two baselines are created:

- VMware Tools Upgrade to Match Host
- VM Hardware Upgrade to Match Host

When scheduling the update, you can granularly schedule separate virtual machines depending on the following power states. For example, you might want to schedule any powered-on machine later because they will require downtime, as shown in Figure 3.6. Your options for scheduling virtual machine updates include the following:

- Powered On
- Powered Off
- Suspended

An orchestrated upgrade of virtual machines is not required but greatly reduces the time it takes to remediate a large number of virtual machines at the same time. If you would rather manually remediate the virtual machines, simply upgrade VMware Tools on each virtual machine and then power off the virtual machine to perform the virtual hardware upgrade.



**Figure 3.6** Scheduling Update Manager Updates

## Monitoring

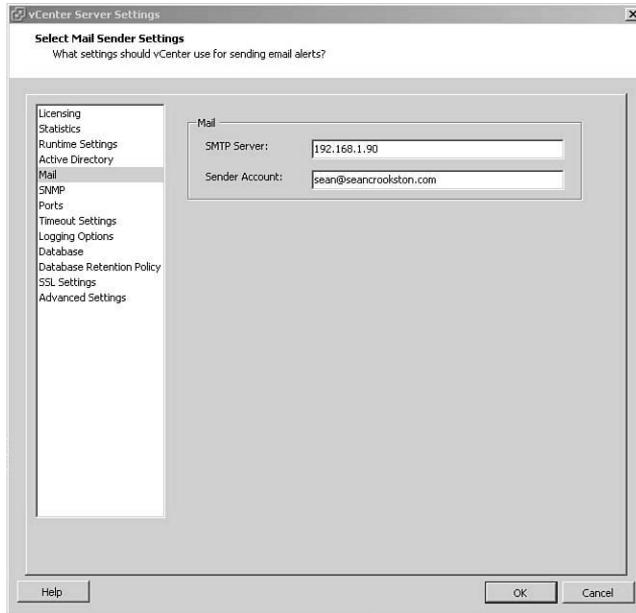
Like maintenance, monitoring is sometimes forgotten with a virtual infrastructure. Many organizations continue their monitoring of their guest virtual machines without a consideration for the hosts themselves. Others consider the hosts but don't have the proper monitoring software, licensing, or understanding of how or what to monitor in the virtual infrastructure. Regardless of the reason, the need to monitor the underlying components of a virtual infrastructure remains high.

## Alerting

I once had a customer contact me who did not understand why he didn't receive an email notification that one of his storage paths had lost redundancy. He had logged in to his vCenter server and noticed the down host, which had been offline for two days. Although this showed off how well the cluster handled the failure of the host, it was a major point of concern for him because he didn't know the host had failed. In this case, the customer had not fully configured the alarms in vCenter. This section discusses the process required to set up alarms as well as some common issues encountered.

For starters, you need to configure the mail setting in vCenter Server.

To do this, go to Administration, vCenter Server Settings from the vSphere Client. Next, configure the SMTP server and appropriate sender account, as shown in Figure 3.7.



**Figure 3.7** Configuring vCenter Email Settings

You need to configure both an SMTP and a sending account. Additionally, you need to ensure your SMTP server can accept relayed messages from your vCenter server.

This is a step that nearly everyone configures during the default install. A common problem, though, is this is where many people stop. By default, vCenter 5 has 54 alarms defined; however, to set up any type of SNMP or email alerting, actions must be individually defined for each alarm.

## Defining Actions for Alarms

For most alarms, only three actions can be defined. You may define an action once or multiple times for each alarm, and you may define multiple types of actions for a single alarm. The actions that are available to be configured are as follows.

- Send a Notification Email
- Send a Notification Trap
- Run a Command

Two monitor types, however, have the capability of performing specific actions. The Alarm Type Monitor for Virtual Machines may take the following actions in addition to sending an email, sending an SNMP trap, or running a command:

- Enter Maintenance Mode
- Exit Maintenance Mode
- Enter Standby
- Exit Standby
- Reboot Host
- Shutdown Host

The Alarm Type Monitor for Hosts may take the following actions in addition to the three actions mentioned—sending an email, sending an SNMP trap, or running a command:

- Power On VM
- Power Off VM
- Suspend VM
- Reset VM
- Migrate VM
- Reboot Guest On VM
- Shutdown Guest On VM

For the following Alarm Type Monitors, the only three actions are to send a notification email, send a notification trap, or run a command:

- Clusters
- Datacenters
- Datastores
- vSphere Distributed Switches
- Distributed Port Groups

- Datastore Clusters
- vCenter Server

The process for defining actions for alarms is pretty straightforward; however, there are a few things to be aware of.

First, as mentioned, 54 alarms are defined by default. Defining all 54 alarms individually would take a long time and would likely result in a few of them being configured incorrectly due to an occasional keystroke error. Don't worry, though, because PowerShell can be used to automate the creation of these actions and is discussed shortly.

Second, when you are defining actions, you must define when the action will occur and how often notification will occur for issues that persist. By default, you receive an email notification only when going from a yellow to a red state. There are four configurable options to consider:

- Green→Yellow
- Yellow→Red
- Red→Yellow
- Yellow→Green

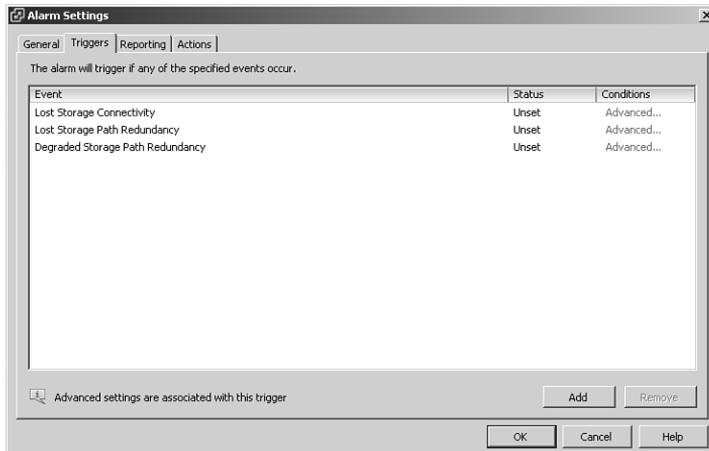
Let's stop for a moment to talk about which of these four you will want to be notified of. If you are relying on SNMP traps being sent to your existing monitoring software, you may choose to have very little to no email notifications. Many smaller environments do not rely on SNMP notifications or still may require email notifications outside of their existing monitoring solutions. For environments with no other monitoring, it is best to configure all of the default alarms and some additional ones as well. These additional recommendations as well as automating the process are discussed in just a bit.

So you now have defined actions for all of your desired alarms as well as the severity changes you would like to be notified of and the amount of times you would like to be notified if the issue persists. That brings us to another common thing to consider for a new implementation.

We have witnessed some environments that simply forgot to allow the vCenter server to use the mail server as a relay. After all, the vCenter server may be a new addition to an environment and would not have been previously configured to relay email messages from the SMTP server. If you are unsure if the mail server is allowing relay for the host and do not have access to the email server to check, you may try the following:

```
telnet mailservername.vmware.com 25
helo vmware.com
```

There is still one more thing to be aware of. Even after all of this, you might find you are not being notified of some issues, for example when storage path redundancy is lost. This is because some triggers are left unset by default, as shown in Figure 3.8. When set to Unset, alarms do not show in vCenter; however, they are sent to email or as SNMP traps if configured. As you can see for the case of lost storage path redundancy, the status for each event is not set.



**Figure 3.8** Unset vCenter Alarms

The following is a list of the other default alarms that are not set up:

- Unmanaged workload detected on Storage I/O Control (SIOC)-enabled datastore (this is disabled by default)
- VMkernel NIC not configured correctly
- Network uplink redundancy degraded
- Health Status Changed Alarm
- License Error
- Exit Standby Error
- Migration Error
- Host Connection Failure
- Virtual Machine Error
- Host Error

- No Compatible Host for Secondary VM
- Timed Out Starting Secondary VM

Two of the default alarms also are not configurable. These alarms are triggered via the vSphere API and can only be modified as such:

- Datastore Capability Alarm
- Thin-Provisioned LUN Capacity Exceeded

When creating actions, you just need to select an SNMP action in addition to or instead of an email notification so that a trap is sent. You may also enable SNMP traps for each individual host if desired. This may be beneficial in the event of a vCenter server outage as the individual hosts themselves will not communicate any status back otherwise.

## Considerations for Tweaking Default Alarms

Some of the default alarms may have some notification options that are less than desirable for your environment. For example, you may have an environment that is strictly testing for internal IT staff. You may decide you still want all the alarms but fully accept that the vSphere hosts in question will likely be pegged pretty hard in terms of memory at certain times of the day. After all, this may be older hardware with lesser memory. You still, however, want to know if there is a consistent condition where memory is steady at 95% or greater for 30 minutes or more.

In this case, by default, the Host Memory Usage alarm warning triggers a warning when host memory usage is above 90% for 5 minutes. Also by default, an alert triggers when host memory usage is above 95% for 5 minutes. By setting both values to 5% higher and to lengths of 30 minutes, you do not get repeated alerts for expected high memory conditions, but do get notified when the issue becomes persistent enough where it may warrant finding additional memory for these hosts.

In closing, you can see that there is a lot to consider even when looking specifically at just vCenter alarms. Walking away from this discussion on alarms, remember the following key points:

- Consider that the alarms can be defined at many levels. Depending on your infrastructure, you might want to define alarms at the vCenter, datacenter, cluster, or individual host level. For that matter, you may also want to get even more granular and enable alarms on specific virtual machines, datastores, datastore clusters, and virtual distributed switches.

- Consider that triggers may have multiple actions that trigger based on both actions happening or one or the other.
- Consider how often you want to be notified and of what state changes you would like to be notified. Too many alerts can become just as big of a problem as not enough alerts at times if you begin tuning them out.

Before moving on to the next section, some assistance in setting up these alarms using PowerShell was promised. With just a few modifications, the provided PowerShell script allows you to easily set up all or as many of the default alarms as you would like. Note that you need to configure the alarms mentioned that are not configured by default to your liking for your environment. Although this still leaves some manual configuration, you no longer have to enter an email address for any of the alarms. It is our recommendation that you start by configuring all vCenter alarms and remove alarms that are not necessary for your environment.

You can download this script from [http://www.seancrookston.com/set\\_alarms.ps1](http://www.seancrookston.com/set_alarms.ps1) (see Appendix A for a link).

## Verifying Configurations

Another important component of operating a vSphere infrastructure is configuration management. When talking about configuration management, the concern is with ensuring configurations are not unknowingly changed or drift from their intended configurations. You want configurations to match their intended configuration and be consistent across the environment. For example, you want your hosts to be running on a certain build of vSphere and to be consistent with the other hosts within the same cluster.

vCenter Operations Enterprise versions include vCenter Configuration Manager, which provides the ability to monitor configuration virtual infrastructure configurations. vCenter Operations are discussed in further detail in Chapter 4, “Managing the Environment.”

Even with a product like vCenter Operations, you still need to implement the most important part of a solid configuration management strategy. Policies and procedures for documenting configuration are the foundation to maintaining an environment with consistent and desired configurations.

When thinking about the configurations, the goal is to maintain many items that might not seem obvious initially. The following is a list of some of the pieces in your virtual infrastructure that might have configurations—in terms of software or firmware—to track and ensure are desired and consistent. Keep in mind this list is brief and we could easily drill even deeper.

- vCenter Server configuration
  - Cluster configuration
  - High Availability configuration
  - DRS configuration
  - Update Manager configuration
- vSphere host configuration
  - vSphere drivers and operating software
  - HBA drivers and firmware
  - NIC drivers and firmware
- IP network configuration
- Storage network configuration
- Storage firmware and software versions
- Virtual guest configuration

## Host Profiles

If a product like Operations Manager is not a fit, you may also use Host Profiles. This feature is included with the Enterprise Plus level of vCenter licensing and is of great assistance with managing the delivery of consistent configurations. Additionally, you may use Scheduled Tasks in vCenter to define a scheduled compliance check that will notify you daily of any configuration drift.

After you've configured your first host to the desired gold state, you can simply create a profile using this host as a reference host. Then you can take your baseline profile and apply it to other hosts or clusters. You will be prompted to enter dynamic information, such as network information during the application, but other configuration settings will be applied consistently to your hosts.

At any time, you can check the host's compliance against the profile or receive notification via email when a drift in configuration occurs. When the time comes to make a change to your standard configuration, the process is just as easy. Simply update your reference host's configuration and then update the profile and reapply the configuration to your other hosts.

Even if you do not have Enterprise Plus licensing, you should consider using Host Profiles during your setup as part of your 60-day evaluation licensing.

## Health Check

So far, this chapter has discussed ways to operationally maintain the environment through updates and alerts. Another important operational step is to perform regular health checks of your environment. This may consist of a physical inspection as well as checking configurations. You may also be ensuring your configured alarms are configured as expected and manually checking for issues just in case. You may also be looking for drifts in configuration based on your organization's standardized configuration.

These are all important things to do and there are many community resources that can assist in these efforts. One such resource is a daily health check script developed by Alan Renouf called vCheck, detailed further in Appendix A.

This script creates a daily report that gives a great report of the environment, including items such as snapshots and new virtual machines that have been created. The setup process has been made easy with an install script, and a great demo video is included on the site for guidance in setting the script up.

Continuing the discussion of performing health checks, another reason to do a health check might be to get a new perspective on the current state of the environment. You might think to yourself, "Well, nothing has changed in this environment in the last three months." Considering that perhaps nothing has changed in the environment, you also need to consider what has changed externally to your environment. This doesn't strictly refer to the storage or networking attached to your vSphere hosts, although checking on these is equally important. Technology is often updated or at times has vulnerabilities due to security flaws in the product.

Bugs, workarounds, patches, and best practices are regularly released and updated. Many individuals barely have the time to perform their regular day-to-day duties, and this information can be difficult to find at times. This is where the aid of someone focused on vSphere technologies is of great advantage.

## VMware's Health Check Delivery

VMware offers a Health Check service that can greatly aid in this need. Any of the information that is used during this process is available to anyone and you could use scripts like the ones mentioned to verify much of the same information. The time to do so could be substantial and unless you have significant experience across many environments, there

may be the risk that you are missing something. The health check delivery has many big advantages, such as the following:

- Consultants will add in their experiences recently as well as perform additional checks.
- Consultants will have at minimum a VCP.
- Quick collection of data for analysis will be performed by an expert.

The result of the engagement is a report and in-depth analysis of the environment with suggestions and remediation. The suggestions are based on best-practice configuration and known issues across a wide range of industries and environments.

---

## **Operating the Environment Summary**

This chapter focused on the operating phase of a virtualization infrastructure. You learned about many tools and methods to operate and monitor on a daily basis in addition to best practices and methods for continuing to bring existing physical workloads into your virtual infrastructure.

This chapter also discussed methods for monitoring and alerting of issues in the virtual infrastructure. From an operational perspective, you have covered the grounds of day-to-day management of your virtual infrastructure.

Moving forward, you need to monitor your environment's performance and capacity for growth. This is discussed in the following chapter.

# Index

## Numerics

---

- 1GbE network architecture, capacity planning, 9-10
- 802.1q VLAN tagging, 136-137

## A

---

- actions for alarms, defining, 114-118
- Administering VMware Site Recovery Manager 5.0, 97
- administrator role as roadblock to virtualization, 177
- admission control policy, 23-25
- advantages of scale-out architecture, 125-126
- affinity rules (DRS), 31-32
- alarms
  - default alarms, configuring, 118-119
  - defining actions for, 114-118
- alerts, 113-114
  - alarms
    - default alarms, configuring, 118-119*
    - defining actions for, 114-118*
- anti-affinity rules (DRS), 31-32
- APIs, Storage APIs-Array Integration, 149-151
- application dependency mapping, vCenter Operations, 163
- application servers, P2V conversions, 100
- applications
  - existing applications, capacity planning, 3-7
  - new applications, capacity planning, 2
  - roadblocks to virtualization
    - business-critical, 185-187*
    - real-time, 181*
    - voice, 182-184*
- array operations (Storage APIs—Array Integration), 150
- assumptions in vSphere design, 47-53

- Auto Deploy Server, automating the implementation, 55-57
- automating the implementation
  - Auto Deploy Server, 55-57
  - Host Profiles, 53-54
  - PowerCLI, 53
  - vCenter Orchestrator, 57-58
- Automating vSphere with VMware vCenter Orchestrator, 58
- auto-tiering
  - and Storage DRS, 34-36
  - storage layout, 18
- availability of storage networking, verifying functionality, 71

## B

---

- backup jobs
  - defining, 91-92
  - monitoring, 94
- backups, 85-86
  - snapshots, 86-88
  - verification and monitoring, 86
- beacon probing, 144
- best practices for virtual switches, 142-144
- binding types for VDS, 145-146
- blades, capacity planning for converged blade architecture, 12-13
- Bunch, Cody, 58
- business-critical applications as roadblock to virtualization, 185-187
- business-driven initiatives as roadblock to virtualization, 175

## C

---

- calculating storage performance, 16-18
- capacity planning, 1-2
  - for customer scenario, 195-212
  - existing applications, 3-7
  - growth, planning for, 123
  - host capacity management
    - scale-out*, 124-130
    - scale-up*, 125
  - networking capacity management
    - planning*, 131-135
    - VDS*, 140-142
    - VSS*, 135-139
  - for networks
    - 1GbE architecture*, 9-10
    - converged blade*, 12-13
    - converged rackmount*, 10-11
  - new applications, 2
  - vCenter as virtual machine, 145
  - with vCenter Operations Management Suite, 161-162
- capex (capital expenditures), decrease in as roadblock to virtualization, 168-171
- case studies
  - business-critical applications, 186-187
  - for design assumptions, 47-48
  - on design constraints, 44-46
  - deviations in design blueprint, 49-51
  - protecting virtual machines with fault, 69
- CIFS shares, 91
- Cisco Nexus 1000v distributed switch, 146-148
- Cisco VNTag, 191
- collecting resource data for existing applications, 4

- commands
    - ping, 72
    - vmkping, 72
  - configuration
    - DRS, 31-32
    - HA, selecting, 23-27
    - host-specific, verifying, 80-81
    - network-specific, verifying, 83
    - profile-driven storage, 36-38
    - selecting, 23
    - storage-specific, verifying, 82-83
    - Storage DRS, 32
      - affinity rules for virtual disks*, 33
      - and auto-tiering*, 34-36
      - I/O latency load balancing*, 33
      - space utilization load balancing*, 33
      - usage recommendations*, 34
    - vCenter-specific, verifying, 78-79
    - virtual machine-specific, verifying, 81-82
    - vMotion
      - multi-NIC vMotion*, 28
      - QoS options*, 30, 39
      - selecting*, 27
  - configuration management, 119
    - host profiles, 120
    - vCheck script, 121
  - configuring
    - CPU affinity, 183
    - default alarms, 118-119
    - iSCSI port binding, 72
    - multi-NIC vMotion, 28
    - vCenter email settings, 114
  - connectivity, verifying, 72
  - constraints on design, 44
    - case study, 44-45
    - common constraints, 46
  - converged blade architecture, capacity planning, 12-13
  - Converged Networking, effect on datastore heartbeating, 27
  - converged rackmount architecture, capacity planning, 10-11
  - conversions
    - P2V
      - application servers*, 100
      - domain controllers*, 99
      - Linux*, 101
      - Windows Server with OEM installations*, 100
    - physical to virtual, 97-99
    - V2V, 101-102
  - cost savings of virtualization, 172
  - CPU affinity, configuring, 183
  - creating virtual machine alarms, 88
  - customer scenario, 193-195
    - implementation, 212-213
    - managing the infrastructure, 215-216
    - operating the infrastructure, 213-215
    - planning and designing, 195-212
- 
- ## D
- Data Recovery, 88
    - backup jobs
      - defining*, 91-92
      - monitoring*, 94
    - installing, 89-90
    - repository, managing, 95
    - restoring data, 92-94
    - shared repository, defining, 90-91
  - datastore clusters, 33
  - datastore heartbeating, 26-27
  - decrease in capex as roadblock to virtualization, 170-171

decrease in opex as roadblock to virtualization, 170

default alarms, configuring, 118-119

defining

- actions for alarms, 114-118
- backup jobs, 91-92
- shared repository, 90-91

defragmentation, 180

dependencies, upgrading hosts, 108

design, planning virtualization projects for, 20, 39

- configuration, 23-38*
- hardware, 22*

design blueprint, 41

- assumptions, 47-48, 52-53
- constraints, 44
- case study, 44-45*
- common constraints, 46*

deviations, 48

- case study, 49-51*
- changes in functional requirements, 48*

functional requirements, 43

reviewing, 42

stakeholder review, 42-43

technical review, 46

desktop virtualization, 187-189

deviations in vSphere design, 48

- case study, 49-51*
- changes in functional requirements, 48*

Directpath I/O, 158-159

disadvantages of scale-out architecture, 126-130

disaster recovery, 95-96, 173

documentation, design blueprint, 41

- assumptions, 47-48, 52-53
- constraints, 44-46
- deviations, 48-51
- functional requirements, 43

- reviewing, 42
- stakeholder review, 42-43
- technical review, 46

domain controllers, P2V conversions, 99

DPM (Distributed Power Management), 61

DRS (Distributed Resource Scheduler), 31, 173

- affinity/anti-affinity rules, 31-32
- functionality, testing, 61
- Storage DRS, 32
- affinity rules for virtual disks, 33*
- and auto-tiering, 34-36*
- I/O latency load balancing, 33*
- space utilization load balancing, 33*
- usage recommendations, 34*

dynamic binding, 145

dynamic resource sharing, 173

---

## E

eager zeroed thick disks, 68

email settings (vCenter), configuring, 114

ephemeral binding, 145

EPT (Extended Page Tables), 183

EVC (Enhance vMotion Capability), 66

existing applications, capacity planning, 3-7

explicit failover, 139

---

## F

failback of virtual machine networking, verifying, 64

failover, explicit failover, 139

fault tolerance, 172

- host failures, simulating, 69-70
- networking functionality, verifying, 66-68
- virtual machines, protecting, 68-69

## features

DRS, 31-32

HA, selecting, 23-27

profile-driven storage, 36-38

selecting, 23

Storage DRS, 32

*affinity rules for virtual disks, 33**and auto-tiering, 34-36**I/O latency load balancing, 33**space utilization load balancing, 33**usage recommendations, 34*

vMotion

*multi-NIC vMotion, configuring, 28**QoS options, 30, 39**selecting, 27*Fibre Channel, verifying functionality of  
storage redundancy, 74-75

file level restores, 93-94

financial roadblocks to virtualization, 168

capex, 168-171

limited funding, 174

opex, 168-170

ROI, 171

forcing kernel panics, 60

forecasting with vCenter Operations  
Management Suite, 161-162

functional requirements, 43

changes in during design, 48

gathering, 41

functionality, testing, 58

of DRS, 61

of fault tolerance networking, 67-70

of HA, 59-60

of IP storage networking, 70-76

of management networking, 62-63

of networking, 62

of quality assurance assessment, 77

of virtual machine networking, 63-65

of vMotion networking, 65-67

future of desktop virtualization, 188-189

**G**

gathering functional requirements, 41

growth, planning for, 123

host performance management, 152

*CPU, 154**memory, 155**resource, 153*network performance management,  
156-157**H**

HA (high availability)

admission control policy, 23-25

datastore heartbeating, 26-27

DRS, testing functionality, 61

functionality, testing, 59

host failures, testing functionality, 59-60

VM Monitoring, 26

halting the implementation case study, 50-51

hardware

networking, selecting, 174

selecting, 22

servers, selecting, 174

storage

*selecting, 174-175**sizing guidelines, 180-181*

Health Check tool, 77

health checks, performing, 121

High Availability, 172

- host capacity management
  - scale-out architecture, 124-130
  - scale-up architecture, 125
- host failures
  - HA functionality, testing, 59-60
  - simulating, 69-70
- host performance management, planning for growth, 152
  - CPU, 154
  - memory, 155
  - resource, 153
- host profiles, 120
  - automating the implementation, 53-54
- host-specific configurations, verifying, 80-81
- hosts
  - patching with Update Manager, 104
  - upgrading with Update Manager, 105-109
    - methods of upgrading*, 110-112
    - order of operations*, 110

## I

---

- implementation
  - automating
    - Auto Deploy Server*, 55-57
    - Host Profiles*, 53-54
    - PowerCLI*, 53
    - vCenter Orchestrator*, 57-58
  - of customer scenario, 212-213
  - halting, case study, 50-51
  - of quality assurance assessment, verifying, 77
  - testing functionality, 58-76
  - verifying, 58
- incorrect assumptions during implementation, 52-53

- infrastructure of customer scenario
  - managing, 215-216
  - operating, 213-215
- installing Data Recovery, 89-90
- IP storage networking functionality
  - availability, 71
  - Fibre Channel, 74-75
  - iSCSI storage, 72-74
  - NFS storage failover, 71
  - performance, verifying, 75-76
  - verifying, 70-71
- iSCSI
  - multipathing, verifying functionality, 73
  - port binding, configuring, 72
  - storage failover, verifying functionality, 72-74
- isolation of virtual machine networking, verifying, 63

## J-K-L

---

- jumbo frames, 157
- kernel panics, forcing, 60
- latency, impact on I/O, 152
- Laverick, Mike, 97
- lazy zeroed thick disks, 68
- LBT (Load Based Teaming), 155
- licensing
  - Microsoft OEM, 172
  - OSE, 127
- limited funding as roadblock to virtualization, 174
- Linux, P2V conversions, 101
- load balancing policies
  - VDS, 141
  - VSS, 138-139

---

**M**

- maintenance, Update Manager, 102-103
  - hosts, patching, 104
  - hosts, upgrading, 105-112
- management, planning virtualization projects for, 19-20
- management networking, testing functionality, 62-63
- managing
  - configurations, 119
    - host profiles*, 120
    - vCheck script*, 121
  - Data Recovery repository, 95
  - infrastructure, customer scenario, 215-216
- manual disaster recovery, 96
- manual host upgrades, performing, 111
- Maritz, Paul, 188
- memory, sizing guidelines, 179-180
- metrics for storage performance, 15-18
- Microsoft OEM licensing, 172
- Microsoft OSE licensing, 127
- migrating VSMs to VEM, 148
- misconfigurations
  - host-specific, 80-81
  - network-specific, 83
  - storage-specific, 82-83
  - vCenter-specific, 78-79
  - virtual machine-specific, 81-82
- monitoring, 113
  - backup jobs, 86, 94
- multipathing (iSCSI), verifying functionality, 73

---

**N**

- NetFlow for VDS, 141
- Netqueue, 158
- network failure detection options for VDS, 142
- network shares, 156
- network-specific configurations, verifying, 83
- network virtualization, 190-191
- networking
  - capacity planning
    - 1 GbE architecture*, 9-10
    - converged blade architecture*, 12-13
    - converged rackmount*, 10-11
    - planning for growth*, 131-135
    - VDS*, 140-142
    - VSS*, 135-139
  - fault tolerance networking
    - functionality*, 67-68
    - host failures*, 69-70
    - virtual machines*, 68-69
  - functionality, testing, 62
  - hardware
    - selecting*, 174
    - voice application recommendations*, 184
  - IP storage networking
    - functionality, testing*, 70-75
    - performance, verifying*, 75-76
  - management networking, testing functionality, 62-63
  - performance management
    - planning for growth*, 156-157
    - planning for peak*, 157-159
  - virtual machine networking, testing functionality, 63-65
  - vMotion networking, testing functionality, 65-67

new applications, capacity planning, 2  
NFS storage failover, verifying  
    functionality, 71  
NIOC (vSphere Network I/O Control),  
    28-29  
normalizing resource data for existing  
    applications, 4-7

## O

---

Octopus, 189  
operating the infrastructure, customer  
    scenario, 213-215  
opex (operational expenditures)  
    decrease in, 170  
    as roadblock to virtualization, 168-169  
organizational cohesiveness, lack of as  
    roadblock to virtualization, 176-177  
OSE (Operating System Environment), 127  
OVA (Open Virtualization Appliance), 2  
OVF (Open Virtualization Format), 2

## P

---

P2V (physical to virtual) conversions  
    application servers, 100  
    domain controllers, 99  
    Linux, 101  
    Windows Server with, 100  
patching hosts with Update Manager, 104  
peak utilization, planning for, 149  
    latency, 152  
    network performance, 157-159  
    SIOC, 151  
    Storage APIs—Array, 149-151

performance. *See also* performance  
    management  
    planning virtualization projects for, 14-18  
    of storage networking, verifying, 75-76  
    of virtual machine networking, verifying,  
        64-65  
    of vMotion networking, verifying, 66-67  
performance management  
    host performance, 152-155  
    network performance, 156-159  
    storage performance management,  
        149-152  
    with vCenter Operations Management,  
        163  
physical to virtual conversions, 97-99  
ping command, 72  
planning virtualization projects  
    capacity planning, 1-13  
    design, 20, 39  
        *configuration*, 23-39  
        *hardware*, 22  
    growth, 123  
        *networking capacity management*,  
            131-135  
peak utilization, 149  
    *latency*, 152  
    *SIOC*, 151  
    *Storage APIs—Array*, 149-151  
performance, 14-18  
    management, 19-20  
political roadblocks to virtualization, 175  
    administrator role, 177  
    business-driven, 175  
    lack of, 176-177  
port binding (iSCSI), configuring, 72  
port mirroring, VDS, 141  
PowerCLI, automating the implementation,  
    53

PowerPath, 75  
 processors, sizing guidelines, 179-180  
 profile-driven storage, 36-38  
 Project Octopus, 189  
 protecting virtual machines with fault tolerance, 68-69

## Q-R

---

QoS  
 shaping, 30  
 traffic policing, 30  
 vMotion, 30, 39  
 quality assurance assessment, Health Check, 77

RAID, storage performance, 16-18  
 RDM (Raw Device Mapping), 51, 90  
 read rate, 15  
 real-time applications as roadblock to virtualization, 181  
 recommendations for virtual switches, 142-144  
 redundancy  
 backups, 85-86  
   *restoring data, 92-94*  
   *snapshots, 86-88*  
   *verification and monitoring, 86*  
 disaster recovery, 95-96  
 verifying, 72  
 virtual machine networking, verifying, 64

Renouf, Alan, 121  
 resource data, existing applications  
 collecting, 4  
 normalizing, 4-7  
 resources, 217-222  
 restoring data, 92-94

retention policy, defining, 92  
 reviewing design documentation, 42  
 roadblocks to virtualization  
 applications, 183-187  
   *CPU, 182-183*  
   *memory, 182-183*  
   *real, 181*  
   *voice, 182*  
 financial, 168  
   *capex, 168-171*  
   *limited funding, 174*  
   *opex, 168-170*  
   *ROI, 171*  
 political, 175  
   *administrator role, 177*  
   *business-driven, 175*  
   *lack of, 176-177*  
 technical, VM sprawl, 178-179  
 ROI as roadblock to virtualization, 171  
 RTO (recovery time objective), 51  
 RVI (Rapid Virtualization Index), 183

## S

---

scale-out architecture, 124-130  
 scale-up architecture, 125  
 scope creep, 167  
 selecting  
 configurations, 23  
   *DRS, 31-32*  
   *HA, 23-27*  
   *profile-driven storage, 36-38*  
   *storage DRS, 32-36*  
   *vMotion, 27-39*  
 hardware, 22  
 networking equipment, 174

- server hardware, 174
- storage hardware, 174-175
- server hardware, selecting, 174
- shared repository, defining, 90-91
- shares, 156
- simulating host failures, 69-70
- SIOC (Storage I/O Control), 151
- Site Recovery Manager, 96, 173
- sizing guidelines
  - memory and processor configuration, 179-180
  - storage configuration, 180-181
- snapshots, 86-88
- SRA (Storage Replication Adapter), 97
- stakeholder review, 42-43
- static binding, 145
- Storage APIs—Array Integration, 149-151
- storage configurations, verifying, 82-83
- Storage DRS, 32
  - affinity rules for virtual disks, 33
  - and auto-tiering, 34-36
  - I/O latency load balancing, 33
  - space utilization load balancing, 33
  - usage recommendations, 34
- storage equipment
  - performance, planning virtualization, 14-18
  - profile-driven storage, 36-38
  - selecting, 174-175
  - sizing guidelines, 180-181
  - verifying IP storage networking, 70-76
  - voice application recommendations, 183
- storage performance management, 149-152
- storage virtualization, 190
- Sub-LUN Tiering, 18

---

## T

---

- technical review of design blueprint, 46
- technical roadblocks to virtualization, VM sprawl, 178-179
- terminating the implementation case study, 50-51
- testing
  - functionality, 58
    - of DRS*, 61
    - of fault tolerance networking*, 67-70
    - of HA*, 59-60
    - of IP storage networking*, 70-75
    - of management networking*, 62-63
    - of networking*, 62
    - of virtual machine networking*, 63-65
    - of vMotion networking*, 65-67
  - performance of IP storage networking, 75-76
  - quality assurance assessment, Health Check delivery, 77
- thick disks, 68
- thick provisioning, 180
- thin provisioning, 124, 180
- traffic policing, 30
- troubleshooting
  - vCenter Operations Management Suite, 160
  - voice application recommendations, 184
- TSO (TCP segmentation offload), 158

---

## U

---

- Update Manager, 102-103
  - hosts
    - patching*, 104
    - upgrading*, 105-112

upgrading hosts with Update Manager,  
105-106, 109  
dependencies, 108  
methods of upgrading, 110-112  
order of operations, 110

## V

V2V (virtual to virtual) conversions, 101-102

VASA (vStorage APIs for Array Awareness),  
38

vCenter

configurations, verifying, 78-79  
as virtual machine, 145

vCenter Operations Management Suite

application dependency, 163  
capacity planning, 161-162  
performance management, 163  
troubleshooting with, 160

vCenter Orchestrator, automating the  
implementation, 57-58

vCheck script, 121

VDS (vSphere distributed switch), 140

binding types, 145-146  
Cisco Nexus 1000v, 146-148  
load balancing options, 141  
NetFlow, 141  
network failure detection, 142  
port mirroring, 141

verifying

backups, 86  
configurations, 119  
*host profiles*, 120  
*host-specific*, 80-81  
*network-specific*, 83  
*storage-specific*, 82-83  
*vCenter-specific*, 78-79

*vCheck script*, 121

*virtual machine-specific*, 81-82

connectivity, 72

implementation, 58

*of quality assurance assessment*, 77

*testing functionality*, 58-76

virtual machines

alarms, creating, 88  
networking, testing functionality, 63-65  
upgrading, 112  
vCenter as, 145

virtual machine-specific configurations,  
verifying, 81-82

Virtual Machine Total Disk Latency, 152

Virtual Storage Pools, 18

virtual switches

recommendations, 142-144

VDS

*binding types*, 145-146

*Cisco Nexus 1000v*, 146-148

virtualization

cost savings from, 172  
desktop virtualization, 187-189  
disaster recovery, 173  
dynamic resource sharing, 173  
High Availability, 172  
network virtualization, 190-191

planning projects, 1

*capacity planning*, 2-13

*design planning*, 20-39

*management planning*, 19-20

*performance planning*, 14-18

roadblocks to

*application roadblocks*, 181-187

*financial*, 168-171, 174

- political, 175-177*
  - technical, 178-179*
- storage virtualization, 190
- VM Monitoring, 26
- VM sprawl as roadblock to virtualization, 178-179
- VMDKs, 124
- vmkping command, 72
- vMotion, 27
  - EVC, 66
  - multi-NIC vMotion, configuring, 28
  - networking, testing functionality, 65-67
  - QoS options, 30, 39
- VMware Capacity Planner, 4
- VMware Horizon Application Manager, 189
- VMware Horizon Mobile, 189
- VMware Project Octopus, 189

- voice applications
  - as roadblock to virtualization, 182-184
    - CPU, 182-183*
    - memory, 182-183*
    - storage, 183*
  - troubleshooting, 184
- VSMs, migrating to VEM, 148
- vSphere Replication, 173
- VSS (standard vSwitch)
  - 802.1q VLAN tagging, 136-137
  - load balancing policies, 138-139
  - networking capacity management, 135

---

## **W-X-Y-Z**

- Windows Server with OEM installations, P2V conversions, 100
- write rate, 15