

vmware® PRESS



# Virtualizing Oracle® Databases on vSphere®

**Kannan Mani**  
**Don Sullivan**

Foreword by **Chris Williams**,  
Global Practice Director Data Center Consulting, Dimension Data



# **Virtualizing Oracle<sup>®</sup> Databases on vSphere<sup>®</sup>**

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, Chinese (Simplified), Chinese (Traditional), French, German, Greek, Hindi, Japanese, Korean, Polish, Russian, and Spanish.

For more information about VMware Press, please visit [vmwarepress.com](http://vmwarepress.com).

# vmware® PRESS



[pearsonitcertification.com/vmwarepress](http://pearsonitcertification.com/vmwarepress)

Complete list of products • Podcasts • Articles • Newsletters

**VMware® Press** is a publishing alliance between Pearson and VMware, and is the official publisher of VMware books and training materials that provide guidance for the critical topics facing today's technology professionals and students.

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 completing the VMware certification exams.



Make sure to connect with us!  
[informit.com/socialconnect](http://informit.com/socialconnect)

vmware®

PEARSON  
IT CERTIFICATION

Safari®  
Books Online

*This page intentionally left blank*

# Virtualizing Oracle<sup>®</sup> Databases on vSphere<sup>®</sup>

Kannan Mani, Don Sullivan

vmware<sup>®</sup> PRESS

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

# Virtualizing Oracle® Databases on vSphere®

Copyright © 2015 VMware, Inc.

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: 0-13-357018-5

ISBN-13: 978-0-13-357018-2

Library of Congress Control Number: 2014914987

Printed in the United States of America

First Printing: October 2014

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 DVD or programs accompanying it.

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

## Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at [corpsales@pearsoned.com](mailto:corpsales@pearsoned.com) or (800) 382-3419.

For government sales inquiries, please contact [governmentsales@pearsoned.com](mailto:governmentsales@pearsoned.com).

For questions about sales outside the U.S., please contact [international@pearsoned.com](mailto:international@pearsoned.com).

**VMWARE PRESS**  
**PROGRAM MANAGER**  
Dave Nelson

**ASSOCIATE PUBLISHER**  
David Dusthimer

**ACQUISITIONS EDITORS**  
Joan Murray  
Mary Beth Ray

**DEVELOPMENT EDITOR**  
Dan Young

**TECHNICAL EDITORS**  
Greg Loughmiller  
Marlin McNeil

**MANAGING EDITOR**  
Sandra Schroeder

**PROJECT EDITOR**  
Mandie Frank

**COPY EDITOR**  
Keith Cline

**PROOFREADER**  
Kathy Ruiz

**INDEXER**  
Brad Herriman

**DESIGNER**  
Chuti Prasertsith

**COMPOSITOR**  
Bumpy Design

**EDITORIAL ASSISTANT**  
Vanessa Evans

*This book is dedicated to my wife, Mohana, and my two sons,  
Sricharan and Akshay, for their continued support in all ways,  
and to my parents for their inspiration.  
I also dedicate this book to the Oracle and VMware communities.*

*—Kannan Mani*

*The success of this book should be shared. First, Donald E. Sullivan,  
my dad, who taught business computing before the first billionaires,  
introduced me to computing many years ago.  
My friend Ron Sparagoski taught me how to coach and how to lead.  
Two of my professors at Arizona State University, Roger Eck and  
Michael Goul, are responsible for my expertise in the subject of data.  
My friend Jeff Francis worked tirelessly to help me understand  
operating systems. My colleagues at Polyserve inspired me to  
brave the world of the start-up, and Kannan Mani believed  
in me enough to convince me to co-author this book.  
And finally, my friend and mentor in all things Oracle, Scott Gossett.*

*—Don Sullivan*

*This page intentionally left blank*

# Contents

**Foreword** xv

**Introduction** xix

**About the Authors** xxiii

**Acknowledgments** xxv

**About the Reviewers** xxix

## **Chapter 1 Introduction to Oracle Databases on Virtual Infrastructure** 1

Virtualization with ESXi and vSphere and the Software-Defined Datacenter 3

Virtualizing Oracle Databases on vSphere: Benefits and Examples 7

Oracle Databases and DBA Fundamentals 8

Understanding Oracle Database Architectures 11

Summary 12

## **Chapter 2 Virtualization and High-Performance Oracle Workloads** 15

Virtualized Oracle Environments on vSphere Key Benefits 15

Consolidating Platforms to Reduce Datacenter Costs 17

Enhancing Database Availability and Cost-Effective Disaster Recovery 19

Provisioning Rapid and New Database Server Environments 21

Reducing Planned Downtime with Migration of Live Oracle Database Servers 22

Guaranteeing Resources in a Shared Environment 23

Achieving IT Compliance 24

Zeroing In on Key Trigger Events 24

Solving Oracle Database Deployment and Management Issues Using VMware 25

Implementing Dynamic Oracle Datacenter Resource Management 26

Minimizing Server Sprawl 27

Meeting SLA Demands for Database Performance, Availability, and Disaster Recovery 27

Supporting a Dynamic Business Environment 27

Minimizing License Costs 28

Maximizing Oracle Workloads and Sizing 28

Option 1: Sizing the Oracle DB Workload 30

Option 2: Sizing the Oracle DB Application Vendor Recommendations 32

Option 3: Sizing Oracle DB Server Vendor Guidelines 35

Testing the Limits: Performance Studies and Stress Tests 36

Summary 36

**Chapter 3 Oracle Databases and Applications in Virtual Infrastructure: Architectural Concepts 37**

- VMware ESXi Hypervisor 38
- Designing Databases on VMware 41
  - Designing for Scalability on Demand 42
  - Designing for High Availability 44
  - Maintaining Compliance 48
  - Consolidating Database Servers 48
  - Virtualizing Oracle RAC 51
- Identifying Key Stakeholders 53
- Summary 54

**Chapter 4 Oracle on vSphere Best Practices 55**

- Implementing ESX Host Best Practices 57
  - Maximizing Performance Using BIOS Settings 58
  - Operating System Processes 59
  - Upgrading the Version of ESX/ESXi and vSphere 60
  - Maximizing Support for a Hardware-Assisted Memory Management Unit 61
- Implementing Memory-Related Best Practices 61
  - Supporting Large Pages 64
  - Implementing Compute (vCPU)-Related Best Practices 65
- Configuring Storage-Related Best Practices 68
  - Categorizing Storage Virtualization Technologies 71
  - Understanding Storage Protocol Capabilities 71
  - Understanding Database Layout Considerations 73
  - Comparing VMFS to RDM: Performance and Functionality 76
- Networking Guidelines 78
- Monitoring Performance on vSphere 79
- Timekeeping in Virtual Machines 81
- Summary 82

**Chapter 5 Oracle Database High Availability: Planned and Unplanned Downtime 83**

- Protecting the Virtualized Environment with vSphere High Availability 84
- Protecting Applications with vSphere and Symantec AppHA 86
- Understanding Oracle RAC in Virtual Machines 88
  - Implementing Oracle RAC One Node 88
  - Implementing Multinode RAC 90
- Deploying Oracle RAC on vSphere 92
- Protecting Oracle Databases Against Downtime 98

---

Transitioning RAC Nodes Between Hosts Using VMware vMotion	100
To RAC or Not to RAC	103
Summary	104

## **Chapter 6 Performance Workload and Functional Stress Test Studies 105**

Oracle Single-Instance Workload Study	106
Test Methodology	106
Test Result Details	109
Oracle RAC Workload Characterization Study	121
vMotion and VMware HA	121
Large-Scale Order Entry Benchmark Kit (Swingbench)	122
Architecture	123
Network Configuration	125
Oracle RAC Installation Overview	130
24-Hour Workload Test	131
Oracle RAC Node vMotion Test	132
Mega vMotion-RAC Functional Stress Test	135
Summary	139

## **Chapter 7 Support and Licensing 141**

Contemplating Oracle Software Support and Licensing	141
Understanding Oracle Certification and Support for VMware Environments	143
Certification of Oracle on VMware vSphere	144
Licensing Oracle	147
Advising VMware Customers	149
Summary	154

## **Chapter 8 Performance Management and Monitoring 155**

Performance Management Terminology	157
The Role of the DBA in Performance Management	158
Processing Power: CPU or vCPU	159
CPU Ready Time (%RDY)	162
Memory	164
System Huge Pages	167
Transparent Page Sharing	172
Non-Uniform Memory Access	172
Networking	175
Network Load Testing	176
Dropped Packets	177

Storage Configuration and Utilization	181
SCSI Queues	182
NFS Storage	185
Storage Access Latency	187
Spindle Busy Average	190
Understanding SCSI Queue Depth on an ESX/ESXi Host and Virtual Machine	191
Storage Path Throughput	192
Storage Benchmarking VMDK	193
Benchmarking and Ongoing Maintenance	197
Iometer	198
Oracle ORION	203
Comparing Storage Types	219
Block Alignment	220
Using pvSCSI and LSI Controllers	222
ASM Is Comparable to an LVM	225
Understanding the Oracle Enterprise Manager vCOPS Adapter	228
Using Oracle Database Server Metrics	230
Installing Oracle Enterprise Manager Adapter	232
Validating the OEM Adapter	233
Creating Oracle Database Custom Dashboard	233
Configuring a Metric Graph (Rollover View) Widget	235
Configuring a Generic Scoreboard Widget	235
Finalizing the Oracle Database Dashboard	236
Summary	239

## **Chapter 9 Business Continuity and Disaster Recovery 241**

VMware vCenter Site Recovery Manager	243
vSphere Replication	245
Storage Array-Based Replication	247
Storage Replication Adapters	247
Application-Based Replication	248
Oracle Data Guard	248
Repairing Logical Data Block Corruption with Oracle Data Guard	249
Combining vSphere Replication and Data Guard	250
Testing SRM vSphere Replication	251
Using Storage Array-Based Replication with vSphere	253
Virtual Provisioning for Oracle ASM Disk Groups	255
Solution Findings	256
Creating a Disaster Recovery Plan	257
Configure Connections	257
Break the Connection	257

---

Export System Logs	257
Using Array-Based Replication	258
Summary	264

## **Chapter 10 Backup and Recovery 267**

Backup and Recovery Principles	269
Backing Up Data Using In-Guest Software Solutions	270
Oracle Database Backup Methods	270
Classic Oracle Database Backups	270
Listing of Storage Vendor Backup Tools	272
Other Backup Tools	273
Storage Vendor Backup Solutions	273
Working with NetApp Backup Solutions	274
NetApp Backup and Restore Solution Overview	274
Integrating NetApp with vSphere	274
Working with NetApp Snapshot	275
Backing Up a Virtualized Oracle Database with NetApp Snapshot	275
Tools Available for Backups of Oracle Using NetApp	276
Step-by-Step Solution for Backing Up a Virtualized Oracle Database with NetApp Storage and NFS Datastores	277
Restoring a Database Using NetApp Snapshot	277
Backup and Restore Use Case with Snap Creator	278
EMC Avamar Backup and Restore Solution Overview	279
Backing Up the Oracle Database	280
Restoring the Oracle Database	281
VMware Data Protection Advanced	281
Comparing VMFS and RDM	282
Backups	283
Understand the Functionality of VMFS Versus RDMs	283
Oracle Data Guard for Backup	284
Oracle Database Backup Strategy Matrix	286
Summary	286

## **Chapter 11 Provisioning and Automation 289**

Migrating Oracle Database from a Physical to Virtual Environment	291
Viewing Oracle Migration from a Physical to Virtual Solution	293
Facilitating Deployments	294
Understanding the Business Scenario	294
Lab Architecture	295
Migrating Oracle Database from Physical to Virtual	296
Configuring Application Blueprints Using vCAC	302

Building a Database-as-a-Service Platform	311
Listing the Benefits of DBaaS	312
Allocating Storage as Part of the DBaaS Paradigm	313
Choosing the Components of a DBaaS Architecture	313
Summary	314

## **Chapter 12 Case Studies 317**

Indiana University	318
American Tire Distributors	320
EMC Information Technology	321
Green Mountain Power	323
The Idaho Supreme Court	324
The University of British Columbia	326
VMware Information Technologies	327
So Many Others	330
Working with Events	330
Summary	332
Book Conclusion	332

## **Index 341**

# Foreword

One of my favorite quotes comes from motivational speaker and business leader Harvey Mackay. He recently wrote, “‘Genius’ is sometimes just not realizing that something is impossible.” That most certainly is the story behind my success as one of the earliest people to virtualize Oracle databases on VMware vSphere. My first experience with this impressive combination of technology was virtualizing Oracle 10g on Linux using VMware ESX 2.5 for the production environment and VMware GSX Server for the non-production systems.

The system was SharePoint Portal Server in a medium server farm configuration, with an in-house application using custom Web Parts. SQL Server was used for the SharePoint database. For reasons tied to “developer preferences,” Oracle 10g was used as the back-end database that the custom Web Parts used for their work. We loaded and tuned the systems following the spirit of much of the best practices we know and are now documented in this book today.

It worked. Flawlessly.

In fact, it worked so well that something as simple as a virtual machine (VM) snapshot on our GSX server systems literally saved the entire project one day from an overzealous developer who accidentally trashed the entirety of the production and development code bases. A simple snapshot rollback saved both of them in a matter of minutes. Nobody questioned whether it would work. Nobody knew enough to question this. Nobody worried about licensing per se. Of course, this client’s very deep pockets gave us “all you can eat” enterprise site licenses for everything, so we wouldn’t have worried anyway. Nobody questioned whether anything we did was supported by the vendors. It was early enough in those days that the vendors provided support because they didn’t know, or weren’t concerned, about using a hypervisor to deploy their systems and applications on virtualized infrastructure.

A few years later, some vendors (arguably more for business as opposed to technical reasons) had begun to care. Once again, my new team and I were challenged to break a new set of rules—ironically with another division of the same client. This time, we deployed a production vSphere platform on converged infrastructure, complete with an iSCSI over lossless 10GbE SAN using jumbo frames (no Fibre Channel anywhere), disk-based backup, and offsite archive to Cloud storage. We deployed multiple three-node Oracle RAC 11.2.0.1 clusters with grid control (and later 11.2.0.2) running in a configuration where the individual RAC nodes were vMotion capable. All of that was running with VMware vSphere 4.0 (and later 4.0.1). This was the first known production system of its

kind ever deployed for a client. While this configuration is considered routine today, back then vMotion capable Oracle RAC nodes were considered impossible.

But, for us, this was a non-negotiable client requirement. So that inevitable combination of 5% inspiration and 95% desperation drove us to ignore the idea of “impossible” and instead find a way to make it work. We broke all kinds of established rules along the way—and set industry firsts and new best practices in the process.

Again, it worked flawlessly, and even more impressive, it was easily fast enough to serve as a backup platform for the Oracle Exadata V2 system we had integrated into our overall solution. I’ll never forget the phone calls I had with Kannan back then when we’d realized what we had just accomplished.

A short time later, we adapted what we learned with Oracle RAC on that fully converged architecture to create a configuration that enabled both RAC and SQL Server clusters to be virtualized and be vMotion capable on any supported vSphere configuration—not just native iSCSI storage. The way we did it is still considered ground breaking today all of these years later.

All throughout this engagement and others since, we fought through the mislaid perceptions and even outright objections to virtualizing Oracle databases being possible. We learned that it is tough for certain software vendors to argue against something being possible to do when faced with the direct evidence of it working perfectly right before their eyes. Amazingly, a few were not convinced even when faced with this evidence. We then learned the importance of negotiating our way through both the business and technology sides of the Oracle database virtualization proposition. That meant having sometimes-heated discussions about Oracle licensing and support with people who were part of the technical side as well as sales. Most amazing were the discussions where some clients were absolutely convinced of their (mis)perceptions of the capabilities of virtual infrastructure. One client I worked with actually considered licensing every single ESX server they had in their building, including production and development, for Oracle RAC—because they were concerned that it might someday run an Oracle database (even though Oracle had never been running or installed anywhere even close to the vast majority of those systems).

You, the reader of this book, have a distinct advantage over early adopters and those of us who learned these lessons the hard way. In this book, Don and Kannan have provided a treasure trove of information and time-saving tips to get the most out of your Oracle on vSphere deployments. As a part of that, they have included ways to overcome the most common issues (and several less-common ones), as well as organizational and even political objections you might encounter. It’s all presented in an easy to understand and easy to use format by authors who, by the way, happen to be the foremost experts on the subject you can find anywhere today.

---

You'll learn about the four V's: viability, value, versatility, and vision. These four tenants cover the full spectrum of issues most everyone will face when virtualizing Oracle on vSphere. Building upon this foundation, the authors deliver a practical set of technical and business best practices for compute, networking, and storage topics, which are presented in detail. You'll learn how to plan for, install, and properly optimize Oracle and Oracle RAC from the leading technical minds on the subject. In addition, there are discussions and tips on how to plan for and properly license Oracle to fit your business needs to ensure that you get maximum value out of your investment. Most important, you'll learn how to get support for Oracle on your vSphere platform, and along the way, you'll see that Oracle on vSphere is arguably supported in a way that's even better than what is available when running Oracle on physical infrastructure alone.

If you're looking to get the most out of Oracle on vSphere, keep this book in an easy-to-reach, prominent place on your desk. You will want to come back to it again and again. Oracle and vSphere together make up one of the most powerful and compatible combinations of technology you can find anywhere. I never cease to be amazed at what we've been able to accomplish with these tools and, as they continue to mature, the future looks very bright. I look forward to seeing you on this incredible journey.

—Chris Williams,  
Global Practice Director Data Center  
Consulting, Dimension Data

*This page intentionally left blank*

# Introduction

The idea of virtualization of infrastructure for all levels of workload is a settled concept in the information technology industry. Because the subject of this book is VMware, we will examine the various factors that have driven the industry toward this inevitable outcome from the perspective of VMware, but the concept of virtualization is a much older and broader concept than the company that has changed the world over the past decade, VMware. The term *virtual* as well as the basic concept of virtualization was probably first used by IBM in the 1960s, along with the concept of a hypervisor, which was derived from the technical concept of a supervisor.

As the decades have progressed, the term virtual has been used, and overused, and as so many other terms in this industry, the term has often been abused. But most significantly, the term has come to depict a true and comprehensive abstraction of the server from a physical resource to a logical resource. As the chapters in the book progress, we discuss the ideas of Type 1 and Type 2 hypervisors as well as paravirtualization versus non-paravirtualization, but regardless of the specific architectural precepts, the idea of true virtualization allows for ubiquitous resource abstraction and all the benefits that are implied therein.

There is a trend in the technology industry of various companies monopolizing or at least claiming certain cultural ownership over letters of the alphabet. Technology professionals can easily guess which companies effectively have laid claim over certain letters. Facebook has claimed the letter *F*, Google *G*, Oracle *O*, Apple *i*, and Twitter has the letter *t*. VMware can therefore claim the letter *v*, specifically in lowercase. Sticking with the theme and labeled with the letter *v*, it is useful to have the discussion about virtualization categorized with that same letter, four times. Viability, value, versatility, and vision constitute the most significant headings of any VMware discussion, especially a discussion focused on business-critical applications (BCAs). Most prominent among those BCA or Tier 1 (maybe even Tier 0) apps are SAP, Microsoft apps such as SQL Server and Exchange, and of course, Oracle.

As time progresses, it is impossible to ignore the pervasive trends in the industry. You may resist, but you will eventually have to at least adapt to and recognize those trends or have them render you obsolete. Years ago, database administrators (DBAs) were responsible for managing databases and only databases. As the more sophisticated relational database management systems (RDBMS) were developed, the role of the DBA was innately expanded. Oracle Parallel Server and subsequently Oracle Real Application Clusters (RAC) forced the database professional to become adept at managing certain network functions; otherwise, the RAC interconnect would not be defined adequately, and the ensuing instability would return unfavorable results. Automatic Storage

Management (ASM) forced the Oracle professional to become a storage administrator, because most professional storage administrators did not embrace the idea of managing an ASM instance. In this decade, we have observed the concepts of virtualization being imposed on the database professional. Reporters in the blogosphere, speakers at conferences, and frequent chat room residents readily use terms such as *vDBA* and *vRAC-DBA* to depict the new set of skills that this decade's database professional must possess or risk obsolescence.

## Prerequisites

This book will enhance the overall work and academic experience of anyone on any level who considers themselves to be an Oracle professional. The deep technical considerations are appropriate for the DBAs as well as the developers in the audience, whereas the higher-level architectural concepts will help information architects of all disciplines to build out elegant and effective systems architectures. Management personnel will find this book invaluable with regard to the nontechnical areas, particularly the sections on licensing and support. They will also find the high-level technical explanations both revealing and confirming.

Anyone who has responsibility for any part of an application stack that includes Oracle software, from the most junior administrator to the seasoned veteran, will find something within these pages to enhance their overall effectiveness as an IT professional.

## Who Should Read This Book

This work has been crafted to include subject matter that is pertinent to not only each level of an application stack but also to each professional discipline. Many technical books are crafted as technical manuals or as academic texts. This book, in part, is a compilation of stories and analogies taken from many conversations over many years and therefore, is written in a manner best suited to be used as the focus of dialogue of a group of disparate IT professionals loitering around someone's cube at the end of a day. It is written to be conversation starter, and it is written in a conversational style.

## Book Overview

The topics this book covers vary from the deep minutia to elegant architecture and from the profoundly obvious to the subtle and elusive. This book is neither a textbook nor a

technical manual; it is literary work. In the interest of that literary effort, we have made substantial use of allegories, metaphors, and analogies for the purpose of both attaining and maintaining the attention of the reader and creating indelible images that have some lasting effect. We have also endeavored to create a conversational tone, which is similar to the approach that we have both used at countless customer meetings, conference speaking sessions, and executive briefings over many years. Our intent is to use this literary work to bring to you, the reader, the essence and the substance of the conversations that we have had over many years. In respect of this effort, we have included web links to many of the documents that we cite so that the reader may immediately access pertinent details. The compilation of the graphs, tables, links, and lists is an essential element of philosophy that this book is based on.

The initial chapters (1–3) focus on the basic ideas of Oracle and vSphere, as well as Oracle on vSphere, and the various different roles involved with that effort. Small, medium, and high workloads are discussed as well as the respective architectures and architectural concepts that should be utilized to optimize the capabilities of the virtualized infrastructure.

Chapter 4 focuses on long-developed best practices for all high workloads, with particular focus on Oracle running on vSphere. Chapter 5 transitions into the realm of high availability and all the options available to meet the requirements of every service level agreement (SLA). Chapter 6 digs deeper into the technical details necessary to grasp when optimizing Oracle performance on vSphere. This subject includes each of the various methods of implementing Oracle, including Oracle RAC.

Chapter 7 takes a slight deviation from the technical to discuss the always-intense areas of Oracle Licensing and Support, with special emphasis on the specialized Oracle support team, which is part of VMware's Global Support Services.

The later chapters, beginning with Chapter 8, focus on performance monitoring and management as well as on infrastructure and application management, starting with vCenter Operations Management (vCOPS) and the Oracle Enterprise Manager (OEM) plug-in. A plethora of other tools and management methodologies are either introduced or discussed in detail in this chapter. Chapter 9 focuses on disaster recovery, discussing both VMware's Site Recovery Manager (SRM) and Oracle's Data Guard. Backup and recovery follows in Chapter 10, which covers concepts such as snapshots both from a vSphere and storage perspective. Major storage paradigms, such as Fibre Channel (FC) and network-attached storage (NAS), are given ample coverage.

Chapter 11 encapsulates the Oracle and business-critical applications discussion by transcending the application layer and focusing on the infrastructure management. Provisioning and automation are becoming more important aspects of systems management, and VMware has all the tools to accomplish these tasks and meet the requirements of the provisioning SLAs. vCenter Automation Center (VCAC), vFabric Application Director (vFAD), and vFabric Data Director will soon be coalesced into a single

automation and provisioning system, but they are discussed here in terms of their individual functionality.

Finally, the success stories and many case studies are discussed in Chapter 12, “Case Studies.” From the massive government institutions to the small start-ups, and from the largest universities to the smallest local school systems, customers are using vSphere as their platform of choice for business-critical and Tier 1/0 applications. ESXi is a hypervisor, but vSphere is a platform of virtualized hardware, and companies of all sizes and styles, and institutions of every possible configuration and purpose, are recognizing the viability and value of running Oracle on vSphere. Everyone who reads the following chapters will also come to recognize vSphere as the premier platform in existence to run Oracle.

## About the Authors



**Kannan Mani** (@kantwit) is currently a Staff Architect - Oracle Solutions for VMware. Kannan has been with VMware for more than 4 years, involved in developing and architecting business critical Oracle solutions on VMware platforms, and helping customers and partners successfully virtualize Oracle on VMware vSphere platform globally. Kannan was previously Reference Architecture Specialist at NetApp, where he architected and developed Oracle solutions on NetApp Storage. Prior to NetApp, Kannan was an Architecture Specialist at Unisys, where he led Oracle Center of Excellence. Kannan is the domain expert in Oracle technologies on various platforms (Storage and Virtualization) and published numerous customer-facing technical documents on Oracle and Database technologies. Kannan has over 17 years in the IT industry experience, and his expertise includes Oracle Real Application Clusters (RAC), Automatic Storage Management (ASM), clustering, customer relationship management (CRM), enterprise resource planning (ERP), business intelligence, performance and scalable enterprise architectures, benchmark and performance, technical solutions marketing and management, virtualization, and Cloud solutions. Kannan is a regular speaker at IOUG, VMworld, VMware Partner Exchange, Oracle Open World, EMC World, NetApp Insight, SNIA, and he is also an evangelist of Oracle technologies. Kannan has been recognized by Oracle as an Oracle ACE, and by VMware as CTO Ambassador and vExpert. Kannan holds a Master's degree in Computer Applications and a Master's degree in Business Administration focused on technology.



**Don Sullivan**, an Oracle Certified Master, a vExpert, and a VMware CTO Ambassador joined VMware in June of 2010 as a Systems Engineer Database Specialist and Oracle Solution Architect for the entirety of the Americas. In that capacity, he has worked with numerous customers and partners focused on the proposition of running Oracle, SQL, and other high-workload systems on vSphere. Presently, the Product Line Marketing Manager for Business Critical Applications at VMware, Don is a frequent speaker at conferences focused on databases and virtualization.

After finishing his Master's thesis at Arizona State University in 1996, Don focused on logical database design with Sybase TxBSQL, and he moved to Denver to work as a contract DBA. Don subsequently worked for AT&T as a contract DBA with both Sybase and Oracle. In 1998, he joined Oracle and Oracle University and became a Senior Principal Instructor for Oracle University, focusing on server products. He taught all server-based

classes for 6 years, which included all New Features classes, OPS/RAC, Backup & Recovery, Performance Tuning, SQL Tuning, Data Guard, and the Data Server Internals (DSI) classes from 7.3 through 10g. He is a co-author of the Oracle Certified Master Practicum, and he is an original Oracle Certified Master. He also co-authored a performance-tuning class text for MySQL. In 2004, he became a consultant with Oracle's Advanced Technology Services (ATS) and spent the next 18 months involved in a number of proofs of concept (POCs) and other post-sales engagements. In 2005, Don joined Polyserve Corporation as the primary customer-facing Oracle Solution Architect. Although his role was primarily pre-sales, he was involved with all Polyserve customers who had Oracle implementations at every step of their implementation, both pre- and post-sales. In 2007, Polyserve was acquired by HP, and he stayed with HP. In that capacity, Don spent the majority of 2009 through 2013 delivering seminars and workshops to large customer groups focused on Oracle over Network File System (NFS). In 2010, Don joined VMware as a customer-facing Systems Engineer Database Specialist with both Sales and later PSO. In addition, Don is also a project manager for many projects to include cross-corporate functional stress tests. Finally, Don manages the virtualizing applications sub-track at VMworld and VMware's series of select database workshops.

# Acknowledgments

Any attempt at listing each and every individual or institution that has influenced our development in the subjects of computing, databases, Oracle, and virtualization (and thus the essence of this book) would prove embarrassingly inadequate. We will, however, endeavor to acknowledge those who directly contributed to the writing of this book by either providing material or explicitly helping us describe the many disparate yet interconnected technical concepts in these pages. Material contributions came in a number of forms. Some individuals contributed to the vast supply of VMware best practices and deployment guides as well as Knowledge Base articles. A few folks directly edited our work, whereas others simply helped us understand nuances of specific areas of technology that are necessary to recognize and discuss if one is to compose a comprehensive book on a subject as deep and broad as Oracle on vSphere.

We should start by recognizing a few companies other than VMware. Oracle, EMC, Cisco, NetApp, and Pure Storage all have many individuals who have significantly influenced our understanding of this technology. VMware has a series of elite Oracle implementation partner companies such as House of Brick, VLSS, Ntirety (now part of Hosting), Viscosity North American, the Yucca Group, and others. These partners work with us at conferences, on panels, and most importantly, with customers on a daily basis.

We will also mention some institutions of higher learning such as Stanford University whose land VMware lives upon but has also provided a number of the great minds that made and keep VMware a great pioneer in technology. Don graduated from both the University of Virginia and Arizona State University, and Kannan studied at Madras University and Walden University. Each of these institutions provided the intellectual framework for both of us on our individual journeys in computing.

Our formal editors from Pearson as and our individual tech editors, Greg Loughmiller and Marlin McNeil, spent significant time with us developing approaches to subtle concepts and composing each and every word. Marlin easily joined us on 25 conference calls, all of which lasted between 1 and 3 hours. Others who contributed to the editing include Mark Achtemichuk, who worked with us on the deep ideas of performance tuning and found the time despite the responsibilities of real life and the arrival of his second child, Luke Patrick. Mark helped us understand that this was an endeavor that needed to be complete and that we were the individuals that needed to complete it. Anoop Jalan stepped up to help us at a time when we were very uncertain as to where, when, or how we would finish. He provided both technical insights and a sense of calm encouragement that helped us attain a degree of serenity during that difficult time. Amanda Blevins has provided much support both professionally and personally over many years. Jonathan Nimer

helped us describe the subtleties of the multifaceted requirements of running Oracle on vSphere with clarity. And Tracie Giovanni helped over multiple decades with the creation of a vast array of analogies, many of which are used in this book.

Todd Muirhead, Bob Goldsand, Mohan Potheri, and Vas Mitra all contributed mightily to the VMware SAP collateral, which was the source of our application examples. Avi Nayak contributed to the intellectual direction of our discussion on the subject of database-as-a-service (DBaaS). Jeff Browning referred to vSphere as a platform of virtualized hardware in many public forums and conferences years before the world accepted that profound notion. Shruti Bhat, Chris Rimer, Dave Welch, Will Monin, and the author of our Foreword, Chris Williams, collectively paved the path toward the messaging many years ago that still to this day defines the fundamental parameters of the Oracle on vSphere discussion. Sam Lucido, Darryl Smith, Jeff Szastak, and Mike Webster have all helped develop the industry accepted approaches and methodologies. Kevin Closson, one of the great authorities on Oracle performance, has been a remarkable and generous source of deep and true understanding and discovery for many years.

Jeremy Kuhnash should be acknowledged for his contributions in the area of automation and provisioning. Emad Benjamin is responsible for development of the majority of the Java content, and that underscores much of the application virtualization approach that we have espoused. K. Gopalakrishnan lent his experience around authoring Oracle books, and Sudhir Balasubramanian continuously contributes to all areas of Oracle-related customer-centric and academic-based activities that we engage in. Andy Nelson, through his calm and thoughtful approach to all things technical, has educated both of us in the area of disaster recovery concepts more than anyone.

Jeff Margolese, Scott Salyer, Mike Adams, Stephen Beck, Matt Stepanski, and John Steiner all provided leadership as well as tolerance of the time dedicated to writing this book. Mike Matthies led the Global Support Team that specializes in Oracle with the unambiguous and enthusiastic promise that all Oracle on vSphere-related problems would be addressed by his team and that the service requests would not be closed until resolution was achieved.

All these individuals, named and unnamed, have profoundly assisted us in developing this text which comprehensively covers everything from technical details to philosophical approaches. Maybe even more important, these individuals kept us from being skewered by audiences everywhere by noticing subtle errors and helping us write with a greater degree of linguistic clarity than we could have done alone.

Other contributors who may not even realize that they have implicitly contributed to the body of knowledge that has been drawn from to create this book include the testing teams from Principled Technologies, especially Corey Bunch, who successfully managed the tests and report generation that is often cited in this book. Also Tushar Patel helped

complete the hardware configuration and generate the data for the original Oracle workload studies that we describe. Samir Shah for his contribution towards Oracle on VCE.

All the panelists on all the panels at VMworld and other events run over the years, including both customers as well as partners, have contributed enormously to the body of knowledge that made this book possible. Some of those panels and panelists are referenced in this book, and some of the actual panel discussions were videoed and are linked in the various chapters.

Overall, these small but impactful groups of Oracle and VMware professionals that have influenced and contributed to this book constitute a loosely connected team that has no restrictive boundaries to entry and extend well beyond the lawns of the VMware campus and past the temporary residences of our professional careers. The commonality that binds this group is our belief that the best approach for any twenty-first century company to implement their Oracle-based business-critical applications and respective databases is on VMware virtualized infrastructure with vSphere. The introduction of the new Independent Oracle Users Group (IOUG) VMware SIG at VMworld-US in San Francisco in August 2014 is indicative of the worldwide adoption of this approach. In closing, we believe that this book constitutes a triumph of the committed with both Oracle and VMware with their shared customers being the ultimate victors.

*This page intentionally left blank*

## About the Reviewers

**Greg Loughmiller** is currently a member of NetApp's Enterprise Ecosystem Organization with a focus on Database Solutions and Architecture using NetApp Storage Systems. He works with the NetApp field community to assist with Solutions and deployments of Oracle databases in the Unix infrastructure space. He also provides assistance with customers for their Oracle deployments on NetApp storage. Prior to working in the Enterprise Ecosystem Organization, Greg was part of the NetApp Professional Services Organization, responsible for designing and implementation of Oracle solutions to meet the needs of those customers across the East Coast of the United States for five years. Greg has been part of organizations responsible for Oracle Database Architecture and deployments for 19+ years. Prior to joining NetApp in 2006, he spent 15 years with a wireless telecommunication provider. All of this time was in the Oracle RDBMS technology space, from Operations DBA, Management of Databases, to a Database Infrastructure Architect.

**Marlin McNeil** is Managing Partner and co-founder at the Yucca Group. There he champions virtualizing all kinds of Oracle on VMware and is helping clients establish innovative business practices through virtualizing their business-critical computing infrastructure. Formerly, he worked as Technical Services Director at Justice Systems, Inc. (1996–2011) and as an Analyst at RE/SPEC, Inc. (1991–1996) contracted to Sandia National Laboratories.

*This page intentionally left blank*

## **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.

We welcome your comments. You can email or write to let us 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 we cannot help you with technical problems related to the topic of this book.*

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

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

Mail: VMware Press  
ATTN: Reader Feedback  
800 East 96th Street  
Indianapolis, IN 46240 USA

## **Reader Services**

Visit our website at [www.informit.com/title/9780133570182](http://www.informit.com/title/9780133570182) 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*

## Introduction to Oracle Databases on Virtual Infrastructure


Oracle databases and software run successfully on vSphere and provide significant scalability, availability, and performance benefits. In fact, virtualization quite simply makes Oracle better. What makes the transition smooth is that the Oracle database administrator's (DBA) skill set, deployment technique, and responsibilities do not change when transitioning from a physical to virtual environment. However, it is important that the DBA's scope of responsibility does increase in breadth.

In years past, the DBA's concerns moved into the network realm as Oracle introduced horizontal scalability with Oracle Cluster (later to become Oracle Parallel Server and finally Oracle Real Application Cluster [RAC]). Over the past decade, the focus of the DBA grew to include storage as Oracle introduced Automatic Storage Management (ASM), and thus, the DBA was confronted with a lack of willingness on the part of the storage administrator to manage the ASM instance. Similarly today, we see the realm of the DBA extend into the virtualization arena. Importantly, it is axiomatic that basic database administration skills do not change when virtualization is included in the stack. This is because ESXi does not alter the kernel of any guest operating system (OS). Likewise, it is equally important for the DBA to embrace the fact that some components of the stack that affect the database have been extended into the virtualized infrastructure, such as networking, storage access, processing capability, and memory. Consequently, the areas of concern for the DBA have been extended, respectively.

Throughout this book, we maintain a conversational tone along with a thematic approach to the organization centered on the idea of the four V's. Often, it can be both entertaining and memorable to point out certain technology industry trends. The trend of the monopolization of specific letters by certain well-known companies is an example. We all know

who dominates the use of the letter *f* or *t* or *i* or even *O*. Ironically, VMware prominently uses the uppercase letter *V* to begin the name VMware, and VMW is the acronym most often associated with VMware. However the main product vSphere begins with a lowercase *v*, and it is commonplace for the individual features to begin with the lowercase letter *v*. The thematic approach referred to earlier will be centered on that letter *V*, although no adherence to the case will be necessary.

The four V's are viability, value, versatility, and vision and are shown in Figure 1-1. The first step in any early implementation process is to convince the critical stakeholders such as the DBAs and information technology (IT) managers that virtualized infrastructure is a *viable* alternative to nonvirtualized infrastructure. Notice the subtlety in the reference to “nonvirtualized infrastructure” as opposed to the more common reference of “physical infrastructure.” Subsequent to the viability concerns being satisfied, we move on to the *value* discussion. The value of vSphere as a platform of virtualized hardware for business-critical applications (BCAs) is discussed in great detail as the chapters of this book unfold. The discussion eventually leads to the *versatility* of vSphere and VMware broadly. VMW field personnel will proudly state to all prospective customers that “We are agnostic to both logical and physical architectures.” And finally, no discussion on this subject is complete without acknowledgment of the future *vision* of vSphere's capabilities and VMware as a corporation. The choice of transitioning a company's entire IT architecture to this wonderful “Platform of Virtualized Hardware” is succinct in description, profound in consequence, but colossal in potential. Therefore, every potential stakeholder in the company should have an understanding of this vision.




**Figure 1-1** Four V's approach

## Virtualization with ESXi and vSphere and the Software-Defined Datacenter

ESX or ESXi is the world's leading x86 hypervisor. Hypervisors were first introduced at IBM in the 1960s by abstracting the machine's supervisor state and allowing multiple virtual machines (VMs) to run simultaneously in separate VM contexts. A hypervisor, sometimes scientifically referred to as a virtual machine monitor (VMM), is software that allows for the creation, management, and runtime execution of independent VMs running their own guest operating systems. The physical machine that the hypervisor runs on is referred to as the host machine.

A number of hypervisor types exist. A Type 1 hypervisor, such as ESXi, runs on the bare metal of the computer. The VMs are created on the layer above ESXi and the guest OS runs within that second layer. ESXi has a minimal memory footprint (144M for vSphere version 5.x). A Type 2 hypervisor runs within a base OS, and therefore the guest operating systems run on the third layer above the hardware. See Figure 1-2.




**Figure 1-2** Type 1 and Type 2 hypervisor

It is also important to point out that ESXi is nonparavirtualized. Paravirtualization is discussed later in the book, but it is important to understand that this means that no guest OS kernel is altered and that there therefore exists a perfect state of abstraction between the guest OS and the hardware.

ESXi is a hypervisor, but vSphere is a “platform of virtualized hardware.” It is logical hardware and should always be described and understood as such. And from this point on, we refer to virtualized hardware and nonvirtualized hardware.

The software-defined datacenter (SDDC) is a philosophy of architecture, not an actual product, although it does imply comprehensive full-stack virtualization. In a complete SDDC, all elements of the datacenter are virtualized. VMware has been very successful virtualizing the server components to include processing with virtual CPUs (vCPUs) and memory allocated to the VM (we avoid using the phrase *virtual memory* because it has other well-accepted connotations), but virtualization of the network and storage have been elusive. It is true that common terms such as virtual disk (VMDK), virtual network interface card (vNIC), and virtual distributed switch (vDS) all imply virtualization, but that is not the reality. An important attribute of true virtualization is abstraction. See Figure 1-3.




**Figure 1-3** VMware SDDC (high level)

As VMware moves into the next phase of technology, the SDDC will include not only the virtualization of the server but also the virtualization of the network and storage. In 2012, VMware acquired Nicira Corporation, whose technology has led to the Network Virtualization and Security platform (NSX), which does meet anyone's strict definition of true network virtualization. Figure 1-4 shows vSphere with NSX incorporated. Subsequent chapters cover specific networking recommendations.

For more information on VMware NSX, refer to the link below.

[www.vmware.com/products/nsx](http://www.vmware.com/products/nsx).



**Figure 1-4** VMware NSX

Storage virtualization comes in many flavors. Often, the phrase is used to describe the storage paradigms used only with “stretch clusters,” in which a single logical unit (LUN) of storage will exist in two different physical locations but synchronicity will be maintained through disk mirroring. Systems such as IBM SAN Volume Controller (SVC), shown in Figure 1-5, HP 3PAR Peer Persistence, and EMC VPLEX, shown in Figure 1-6 (formally Yotta-Yotta technology), correctly come to mind when the phrase *storage virtualization* is used. In the VMware context, we understand storage virtualization as an intrinsic part of the SDDC, and we include advanced VMware storage capabilities such as virtual storage-area network (vSAN), the soon-to-be released virtual volumes (vVols), and vFlashReadCache (vFRC) among others as the essential components of the storage virtualization paradigm.

For more information on HP 3PAR, refer to [hp.com](http://hp.com) or to the link below.

[www8.hp.com/us/en/products/storage-software/product-detail.html?oid=5335710#!tab=features](http://www8.hp.com/us/en/products/storage-software/product-detail.html?oid=5335710#!tab=features).




Figure 1-5 IBM SVC stretch cluster




Figure 1-6 EMC VPLEX stretch cluster

Together, the tried-and-true virtualization of the server through vSphere when combined with the more recent ideas of network and storage virtualization constitute the SDDC VMware style. And throughout this book, we refer to each component of this evolutionary leap in datacenter design. Paul Maritz, the former CEO of VMware and present CEO of Pivotal Solutions, once referred to VMware as “the magic pixie dust that was changing the world.” We agree, although we constantly have to point out to customers that despite the magical illusion presented by VMware products, the laws of physics still apply, and so each implementation must consider the limits of those pesky rules of Newton and Einstein otherwise risk disappointment. However, it is not overheated rhetoric to state that VMware is changing the world. For example, it would be difficult to find a single corporate entity that has facilitated more reduction in power consumption throughout its customer base. More importantly, the name VMware itself implies a tectonic-like shift in the world of technology. Computing started with hardware, and then intrepid minds developed software to effectively use that hardware; but to comprehensively tie software to hardware, the world needed VMware.


For more information on EMC VPLEX stretch clusters, refer to the link below.

[www.emc.com/storage/vplex/vplex.htm](http://www.emc.com/storage/vplex/vplex.htm).

## **Virtualizing Oracle Databases on vSphere: Benefits and Examples**

There are many benefits to virtualizing infrastructure for Oracle databases on vSphere. Among those are the reduction of the number of physical systems your organization requires and the more efficient use of existing systems. However, the most important benefits are the resource management capabilities innate in vSphere and the features that facilitate the guarantee of adherence to service level agreements (SLAs) such as availability, disaster recovery, performance, security, and provisioning. This section describes these benefits and then illustrates examples.

Virtualizing database workloads on vSphere significantly reduces the number of physical systems your organization requires, while achieving more effective utilization of datacenter resources. Clients realize tangible savings from this consolidation along with operational cost savings from reduced datacenter floor space, power, and cooling requirements. Figure 1-7 illustrates an example of Oracle database servers on vSphere architecture with application services and infrastructure services.




**Figure 1-7** Oracle databases on VMware vSphere

## Oracle Databases and DBA Fundamentals

DBAs wear many hats and play many roles within an organization. A DBA is the administrator who designs, implements, tests, operates, and maintains databases for an organization. Figure 1-8 illustrates the DBA's general tasks, and the list that follows provides further explanation.

- **Database design, storage, and capacity planning:** DBAs play a major role in designing the database along with determining disk storage requirements and future database growth. Monitoring database growth trends is important so that the DBA can advise management on long-term capacity plans.
- **Install, configure, upgrade, migrate, and provisioning:** Although system administrators are generally responsible for the hardware and OS on a given server, installation of the database software is typically done by the DBA. This role requires knowledge of the hardware prerequisites and requirements so that the database server runs efficiently, and then communicating those requirements to the system administrator. The DBA installs the database software and selects from various options to configure it for the purpose for which it is being deployed. As new releases and patches are made available, it is the DBA's role to determine which are appropriate and to

complete the installation. If the server is a replacement server, it is the DBA's role to transfer the data from the old server to the new one. The more seasoned DBA will take a highly cautious approach to any data transfer or migration operation because data loss is a potential consequence of sloppy transitions. DBAs are tasked to provision database servers on demand for development, testing, QA, and reporting.



**Figure 1-8** General tasks for DBAs

- **Database security:** Databases centralize the storage of data and are attractive targets for hackers. DBAs must understand the particular security model that the database product uses as well as the security requirements of the application and how to effectively control data access. The three basic security tasks are authentication (setting up user accounts to control logins to the database), authorization (setting permissions on various schemas and database objects), and auditing (tracking user movements and actions within the database). The auditing task is particularly important as regulatory laws, such as Sarbanes-Oxley (SOX), the Healthcare Insurance Portability and Accountability Act (HIPAA), and Payment Card Industry (PCI), have security and reporting requirements that require adherence.

- **Backup and recovery, high availability:** DBAs are responsible for developing, implementing, and periodically testing a backup and recovery plan for the databases they manage. Even in large organizations where a separate system administrator performs server backups, the DBA has final responsibility for ensuring that the database backups are done as scheduled and that they include all the files necessary to make database recovery possible after a failure. When failures do occur, the DBA needs to know how to use the backups to restore the database to operational status as quickly as possible, without losing any transactions that were committed. There are several ways a database can fail, and the DBA must have a strategy to recover from each type of failure. From a business standpoint, there is a cost to doing backups, and the DBA makes management aware of the cost/risk trade-offs of various backup methods. DBAs use techniques such as online backups, clustering, replication, and standby databases to provide higher availability.


**TIP**

“When running on vSphere Oracle remains the same” is one of the main pillars of Oracle on vSphere. This nugget of wisdom applies to all backup philosophies. Upon migration to virtualized infrastructure, DBAs can maintain all backup process and techniques previously used. However, DBAs should consider incorporation of the capabilities of virtualization into their back strategy. The classic approaches to include the use of Recovery Manager (RMAN), storage vendor snapshot database tools, and the manual copying of database files remain the most popular methods.

- **Performance tuning and monitoring:** DBAs are responsible for monitoring the database server on a regular basis to identify bottlenecks and remedy them. Database server tuning is performed at multiple levels. The capacity of the server hardware and the way the OS is configured can become limiting factors, as can the database software configuration. The way the database is physically laid out on the disk drives and the types of indexing chosen also have an effect. The way queries against the database are coded can dramatically change how quickly results are returned. A DBA needs to understand which monitoring tools are available at each of these levels and how to use them to tune the system. Proactive tuning involves designing performance into an application from the start, rather than waiting for problems to occur and fixing them. It requires working closely with developers of applications that run against the database to make sure that best practices are followed so that good performance will result.
- **Troubleshooting and support:** When things go wrong with the database server, the DBA needs to know how to quickly ascertain the problem and to correct the issue without losing data or making the situation worse. DBAs provide 24x7 support, 365 days a year.

## Understanding Oracle Database Architectures

The Oracle Server is a relational database management system (RDBMS) that provides an open, comprehensive, and integrated approach to information management. An Oracle server consists of an Oracle database and an Oracle instance. Figure 1-9 describes the relationship between the database and the instance.



**Figure 1-9** Oracle database server architecture

However in Oracle RAC, there may be more than one instance accessing the same database. An instance and a database may have a many-to-one relationship when using RAC, and a one-to-one in case of single-instance non-RAC Oracle deployments.

The following steps describe a basic Oracle configuration where the user and associated server process are on separate machines connected via a network:

- An instance is running on the computer that is executing Oracle, often called the host or database server.
- A computer used to run an application (a local machine or client workstation) runs the client as a user process. The client application attempts to establish a connection to the server using the appropriate SQL\*Net/Oracle network driver.

- The server is running the proper SQL\*Net/Oracle network driver. By default, the server detects the connection request from the application and creates a (dedicated) server process request on behalf of the user. Other types of database connections are possible, such as “shared server” connections as well as connections coming from application servers by proxy through connection pools.
- The user executes a SQL statement and commits a transaction. For example, the user changes a name in a field or row of a table.
- The server process receives the statement and checks the shared pool for a SQL statement resident in the shared SQL area that contains an identical SQL statement. If a shared SQL statement is found, the server process checks the user’s access privileges to the requested data and the previously existing shared SQL statement is used to process the new statement; if not, a new shared SQL area is allocated for the statement so that it can be parsed and processed.
- The server process retrieves any necessary data from the actual data file (tables) after checking for the respective data blocks in the buffer cache.
- The server process may modify data in the buffer cache, which is a primary component of the system global area (SGA). Once the transaction is committed, the log writer (LGWR) process immediately records copies of the transcription of the transaction from the log buffer to the online redo log file. At this point, the database writer (DBWR) process writes modified data blocks permanently to the data files on disk when doing so is efficient.
- If the transaction is successful and disk acknowledgment is received, the server process sends a message across the network to the application. If it is not successful, an appropriate error message is transmitted.

Throughout this entire procedure, the other background processes run, watching for conditions that require intervention. In addition, the database server manages other user transactions and attempts to minimize contention such as locking, deadlock conditions, and bottlenecks on processing resources.

## Summary

The most important advancement in twenty-first century computing technology has been the inception and maturity of virtualization, more precisely VMware vSphere. At this point in the evolution of the industry, it is indisputable that 99.9% of all database or data management systems should be considered candidates for virtualization on vSphere. Oracle databases and software are prime candidates to consider migrating to virtualized

infrastructure. Subsequent chapters discuss the ever-evolving role of the DBA, which will be delivered with a heartfelt admonition: to consider the direction of technology growth and the projection of your own position in the next decade. Only one conclusion is rational because only one path is sustainable: Embrace virtualization as the next phase of the ever-expanding province of the DBA or face obsolescence.

*This page intentionally left blank*

# Index

## A

---

- access latency, NFS storage, 187-190
- ACID (Atomic, Consistency, Integrity and Durability), 269
- actions, application blueprints, 307
- Adapter for Hyperic Oracle Plug-in, 229
- adapters, paravirtualized SCSI, 78
- affinity, 158
- aligning file partitions, 77-78
- alignment, partitions, 220-222
- allocation, memory, 336
- American Tire Distributors (ATD) case study, 320
- AMM (Automatic Memory Management), 168-169
- AppD (Applications Director), 26
- application blueprints
  - actions, 307
  - adding application components, 304
  - adding dependencies, 305-306
  - adding multiple NICs, 308
  - adding services, 303
  - binding properties, 306-307
  - configuring, vCAC, 302-310
  - disk profiles, 308-309
  - external services, 309-310
  - rapid provisioning, 310
- application vendor recommendations, databases, sizing, 32-34
- application-based replication, 248
- applications
  - accelerating delivery, 41
  - BCAs (business-critical applications), 15
  - DHW (data warehouse), 15
  - legacy rehosting, 50
  - QoS, improving, 41
  - RAC (Real Application Clusters), virtual machines, 89-98
  - reducing costs, 41
- Applications Director (AppD), 26
- architectures
  - databases, 11-12
  - ESXi, 38
  - Optimal Flexible Architecture (OFA), 73
  - RAC VMs, 123-124
- array-based replication, 258-264
- ASM (Automatic Storage Management), 1, 55, 224
  - disk groups, virtual provisioning, 255-256
  - managing files, 73-74
  - versus LVMs, 225-228
- ATD (American Tire Distributors) case study, 320

Atomic, Consistency, Integrity and Durability (ACID), 269  
 Automatic Memory Management (AMM), 168-169  
 Automatic Storage Management (ASM).  
*See* ASM (Automatic Storage Management)  
 automation, 289-291  
 availability, databases, enhancing, 19-21  
 Avamar, backup and recovery, 279-281  
 AWR (Automatic Workload Repository), reports, 155

## B

back-end storage connectivity, 127  
 backup, 286, 338  
   Avamar, 279-280  
   cold/offline, 270  
   data pump/export/import, 272  
   databases, 10  
     Database Backup strategy matrix, 287  
   DG (Data Guard), 284-286  
   essential components, 268  
   Flashback, 272  
   hot/online, 271  
   in-guest software solutions, 270  
   methods, 270-273  
   NetApp, 274  
     integrating with vSphere, 274  
     Snap Creator, 276  
     Snap Manager, 276  
     Snapshot, 274-276  
     storage, 277  
   Netbackup, 273  
   principles, 270-273  
   RMAN, 271  
   storage vendor solutions, 273-274  
   VDPA (VMware Data Protection Advanced), 273, 282  
 ballooning, memory, 171-172  
 BCAs (business-critical applications), 15, 18  
 Benchmark Factory for Databases, 30  
 best practices, 55, 82, 334  
   configuring storage-related, 68-76  
   ESX, implementing host, 57-58  
   implementing memory-related, 61-64

  implementing vCPU-related, 65-68  
   networking guidelines, 78-79  
 binding properties, application blueprints, 306-307  
 BIOS settings, maximizing performance, 58  
 block alignment, 220-222  
 Blue Medora, 157  
 Browning, Jeff, 331  
 Buono, Eric, 331  
 business-critical applications (BCAs), 15, 18

## C

capacity planning, databases, 8  
 case studies, 339  
   ATD (American Tire Distributors), 320  
   EMC IT, 321  
   GMP (Green Mountain Power), 323  
   Indiana University, 318-319  
   ISC (Idaho Supreme Court), 324-325  
   University of British Columbia, 326  
 certification, VMware environments, 143-147  
 chief information officers (CIOs), 16  
 chief technology officers (CTOs), 16  
 CIOs (chief information officers), 16  
 clients, dNFS, configuring, 217-218  
 Closson, Kevin, 30  
 cloud  
   CloudDBA, 290  
   vCAC (vCloud Automation Center), 338  
   vCloud Hybrid Service, 338  
 clusters, 157  
   EMC VPLEX stretch clusters, 6-7  
   isolated ESXi, 150-151  
   RAC (Real Application Clusters), 55  
     virtualizing, 51-52  
     virtual machines, 89-98  
   vMSC (vSphere Metro Storage Clusters), 338  
   VMware, 42  
 cold/offline backups, 270  
 Comparison of Storage Protocol  
   Performance in VMware vSphere 5.x, 72  
 compliance (IT)  
   achieving, 24  
   maintaining, 48

- configuration
  - vCenter SRM connections, 257
  - vSwitches, 125
- connections, vCenter SRM, breaking and configuring, 257
- consolidating databases, 48-51
- consolidating platforms, 17-19
- controllers
  - LSI, 222
  - pvSCSI, 222-224
- cost-effective disaster recovery, 19-21
- costs, applications, reducing, 41
- CPUs (central processing units)
  - affinity, 158
  - best practices, 65-68
  - overhead, 157
  - ready time, 160-164
  - versus vCPUs, 159-164

## D

---

- dashboard, Oracle Database
  - configuring metric graph widget, 235
  - creating, 233-235
  - finalizing, 236-238
- database administrators (DBAs). *See* DBAs (database administrators)
- Database Backup strategy matrix, 287
- database-as-a-service (DBaaS). *See* DBaaS (database-as-a-service)
- databases
  - architectures, 11-12
  - backup and recovery, 10
  - backup methods, 270-273
  - capacity planning, 8
  - consolidating, 48, 51
  - DBaaS (database-as-a-service), 22, 311-314, 338
  - DBAs (database administrators), 1, 8, 15
    - general tasks, 9-10
    - performance management role, 158-159
  - design, 8
  - designing on VMware, 41-42
    - HA (high availability), 44-47
    - scalability on demand, 42-44
  - enhancing availability, 19-21
  - HA (high availability), 10, 44-47, 84-86
  - HammerDB, 30
  - meeting SLA demands, 27-30
  - monitoring, 10
  - Oracle Database
    - virtualization, 291-302
      - dashboard, 233-238
      - server metrics, 230
  - performance studies, 36
  - performance tuning, 10
  - protecting against downtime, 98-100
  - protecting applications, Symantec AppHA, 86-87
  - protecting virtualized environment, vSphere HA, 84-86
  - provisioning rapid server environments, 21-22
  - recovery, cost-effective, 19-21
  - security, 9
  - sizing application vendor recommendations, 32-34
  - sizing workloads, 30-31
  - stress studies, 36
  - troubleshooting, 10
- Data Guard (DG). *See* DG (Data Guard)
- data pump/export/import backups, 272
- data warehouse (DWH) application, 15
- datacenters
  - consolidating platforms, 17-19
  - implementing dynamic resource management, 26-27
  - protecting, site recovery manager, 46-47
  - SDDC (software-defined datacenter), 4
  - virtualizing, 17
- datastores
  - adding extents to, 223
  - growing, 224
- DBaaS (database-as-a-service), 22, 311, 314, 338
  - allocating storage, 313
  - benefits, 312-313
  - choosing components, 313-314
- DBAs (database administrators), 1, 8, 15
  - general tasks, 9-10
  - performance management role, 158-159
- decision support system (DSS) kits, 30
- dedicated RAC interconnect networks, 178
- dependencies, application blueprints, adding to, 305-306

designing databases on VMware, 41  
     HA (high availability), 44-47  
     scalability on demand, 42-44

DG (Data Guard), 20, 248-249, 284-286  
     combining with VR (vSphere Replication), 250-251  
     repairing logical data block corruption, 249

Direct-NFS (dNFS). *See* dNFS (Direct-NFS)

disaster recovery, cost-effective, 19-21

disaster recovery (DR) plan, 337  
     configuring VMs in, 261  
     creating, 257-261  
     editing, 261  
     exporting, 264  
     removing, 261  
     running, 263-264  
     testing, 262

disk groups (ASM), virtual provisioning, 255-256

disk profiles, application blueprints, 308-309

Distributed Resource Scheduler (DRS), 19, 23, 43-44, 167

dNFS (Direct-NFS), 215-216, 270, 337  
     client configuration, 217-218  
     verification, 218-219

downtime, protecting databases against, 98-100

DRD (Distributed Resource Scheduler), 19

DR (disaster recovery) plan, 337  
     configuring VMs in, 261  
     creating, 257-261  
     editing, 261  
     exporting, 264  
     removing, 261  
     running, 263-264  
     testing, 262

dropped packets, monitoring, 177-181

DRS (Distributed Resource Scheduler), 23, 43-44, 167

DSS (decision support system) kits, 30

Dunington, Brent, 331

DWH (data warehouse) application, 15

dynamic expansion, VMFS, 223

dynamic resource management, datacenter, 26-27

---

## E

eager-zeroed thick disks, 225-226

EMC IT (Information Technology) case study, 321

EMC VPLEX stretch clusters, 6-7

EMC World, 331

Enterprise Manager, vCOPS, 228-238

ESX  
     hosts and VMs, SCSI queues, 191-219  
     implementing host best practices, 57-58  
     upgrading, 60-61

ESXi, 1  
     architecture, 38  
     hosts and VMs, SCSI queues, 191-219  
     hypervisor, 38-41  
     isolated cluster, 150-151  
     upgrading, 60-61  
     virtualization, 3

Esxtop, 109

events, 330-331  
     key trigger, 24-25

exporting  
     recovery history results, 264  
     recovery plans, 264

extended page tables, 19

extents, adding to datastores, 223

external services, application blueprints, 309-310

---

## F

fast recovery area (FRA), 255

file systems  
     dNFS (Direct-NFS), 270  
     VMFS (Virtual Machine File System), 282-284

files  
     managing, ASM, 73-74  
     partitions, aligning, 77-78

first write penalties, 225

Flashback, 272

flash storage arrays, 339

four Vfs, 1-2, 16, 37, 56, 83-84, 105, 142, 242, 267, 317, 333

four V2s approach, 289

four Vss (viability, value, versatility, and vision), 155-156  
 FRA (fast recovery area), 255

## G

Giles, Dominic, 107  
 GMP (Green Mountain Power) case study, 323  
 GOS (Guest Operating Systems), 270  
 Green Mountain Power (GMP) case study, 323  
 growing datastores, 224  
 Guest Operating Systems (GOS), 270

## H

HA (high availability), 83, 104, 337  
   databases, 10  
     protecting against downtime, 98-100  
     protecting virtualized environment, 84-86  
   designing databases, 44-47  
   enhancing databases, 19-21  
   protecting applications, 86-87  
   vMotion, 121-122  
   VMware, 121-122  
 HammerDB, 30  
 hardware-assisted Memory Management Unit (MMU), 61  
 Haverfield, Scott, 331  
 Health Insurance Portability and Accountability Act of 1996 (HIPAA), 48  
 high availability. *See* HA (high availability)  
 high-performance scheduler, 19  
 HIPAA (Health Insurance Portability and Accountability Act), 48  
 hosts  
   ESX, implementing best practices, 57-58  
   RAC nodes, transitioning between, 100-103  
   VMware, 42  
 hot-add feature (VMware), 42  
 hot/online backups, 271  
 HP Load Runner, 30

Huge Pages, 167-172  
 Huge Pages (Linux), 19  
 hypervisor (ESXi), 38-41

## I

IBM SAN Volume Controller (SVC), 5-6  
 Idaho Supreme Court (ISC) case study, 324-325  
 IgniteVM, 156  
 Independent Oracle User Groups (IOUG) Collaborate, 331  
 Indiana University case study, 318-319  
 information technology (IT). *See* IT (information technology)  
 in-guest software solutions, backing up data, 270  
 installing  
   OEM adapter, 232  
   RAC, 130-132  
 Iometer (Windows Server 2008), 198-203  
 Iperf network testing tool, 176  
 ISC (Idaho Supreme Court) case study, 324-325  
 isolated ESXi cluster, 150-151  
 isolation, 158  
 IT (Information Technologies), 16  
   achieving compliance, 24  
   compliance, maintaining, 48  
   organization hierarchy, 17  
   VMware, 327-330

## J-L

jumbo frame networks, 179-181  
 key trigger events, 24-25  
 Kim, Charles, 331  
 lab architecture, Oracle Database, 295-296  
 large pages, supporting, 64-65  
 large-scale order entry benchmark kit (Swingbench), 122-124  
 Large-Scale Order Entry Benchmark Kit with Swingbench, 106-108  
 latency, 157

- layout, storage, 126-129
- legacy application rehosting, 50
- licensing, 335
  - Oracle, 141-143
- licensing Oracle, 147-154
  - isolated ESXi cluster, 150-151
  - subcapacity approach, 153
  - subcluster approach, 151-152
- Linux Huge Pages, 168
- load testing networks, 176
- logical data block corruption, repairing, DG (Data Guard), 249
- logical volume manager, 224
- LSI controllers, 222
- LVMs (logical volume managers)
  - versus ASM (Automatic Storage Management), 225-228

---

## M

---

- managing files ASM (Automatic Storage Management), 73-74
- Mani, Kannan, 327
- Maritz, Paul, 7
- mega vMotion-RAC functional stress test, 135-138
- memory
  - allocation, 336
  - AMM (Automatic Memory Management), 168-169
  - ballooning, 171-172
  - oversubscription, 18
  - performance management, 164-167
    - Huge Pages, 167-172
    - NUMA (non-uniform memory access), 172-175
    - TPS (Transparent Page Sharing), 172
  - virtual, 4
  - virtual machine memory reservation, 18
- Memory Management Unit (MMU),
  - hardware-assisted, 61
- memory-related best practices,
  - implementing, 61-64
- metric graph widget, Oracle Database
  - dashboard, configuring, 235
- metric sampling, 155

- migration
  - Oracle Database, physical to virtual, 296-302
  - virtual machines, 117-119
- minimizing server sprawl, 27
- MMU (Memory Management Unit),
  - hardware-assisted, 61
- monitoring databases, 10
- multinode RAC, implementing, 90-91
- multipathing, 178

---

## N

---

- nested page tables, 19
- nested/extended page tables, 19
- NetApp
  - integrating with vSphere, 274
  - Snap Creator, 276
  - Snap Manager, 276
  - Snapshot, 274-275
    - backing up virtualized databases, 275-276
    - restoring virtualized databases, 277-279
  - storage, backing up virtualized databases, 277
- NetApp Insight/Foresight, 331
- Netbackup, 273
- netstat, dNFS verification, 219
- network bandwidth, vMotion, 177-178
- Network File Storage (NFS), 336
- Network Virtualization and Security platform (NSX), 4
- networking guidelines, 78-79
- networks
  - dedicated RAC interconnect, 178
  - jumbo frame, 179-181
  - load testing, 176
  - monitoring dropped packets, 177-181
  - storage, 178
  - VM, 178
- NFS (Network File Storage), 336
  - benchmark test, setup notes, 217
  - datastores, backing up virtualized databases, 277
  - dNFS (Direct-NFS), 215-216
    - client configuration, 217-218
    - verification, 218-219

- mounting directly to ESXi hosts, 213
- SCSI queues, 183
- storage, 185-187
  - access latency, 187-190
- NICs (network Interface cards), application blueprints, adding multiple, 308
- NIST Cloud definitions, 332
- NSX (Network Virtualization and Security platform), 4
- NUMA (non-uniform memory access), 172-175

## O

- Occam's Razor, 268
- OCFS (Oracle Clustered File System), 74
- OCI (Oracle Call Interface), 103
- OEM (Oracle Enterprise Manager) adapter, 228
  - installing, 232
- OFA (Optimal Flexible Architecture), 73
- offline/cold backups, 270
- OLTP (online transaction processing) loads, 106-107, 157
- One Node (RAC), implementing, 89-90
- online/hot backups, 271
- online transaction processing (OLTP) loads, 106-107, 157
- ONS (Oracle Notification Services), 103
- operating system processes, 59-60
- Optimal Flexible Architecture (OFA), 73
- Oracle
  - adoption rate, 339
  - licensing, 147-154
    - isolated ESXi cluster, 150-151
    - subcapacity approach, 153
    - subcluster approach, 151-152
  - obtaining support, 145
  - software support and licensing, 141-143
  - support
    - negotiating terms of, 147
    - obtaining, 146
- Oracle Call Interface (OCI), 103
- Oracle Clustered File System (OCFS), 74
- Oracle Database
  - virtualization, 291-306
  - migration, 296-302
- virtualization
  - business scenario, 294-295
  - facilitating deployments, 294
  - lab architecture, 295-296
- Oracle Database dashboard
  - configuring metric graph widget, 235
  - creating custom, 233-235
  - finalizing, 236-238
- Oracle Database server metrics, 230
- Oracle Enterprise Manager (OEM)
  - adapter, 228
  - reports, 155
- Oracle I/O Numbers (ORION). *See* ORION (Oracle I/O Numbers)
- Oracle Notification Services (ONS), 103
- Oracle Open World (OOW), 330
- Oracle RAC node vMotion test, 132-135
- Oracle server, creating workflow, 299
- organization hierarchy, IT, 17
- ORION (Oracle I/O Numbers), 203-208
  - configuring disk shares, 209
  - disk-to-disk, 204-206
  - dNFS (Direct-NFS), 215-216
  - RAID, 209-210
  - RDM (raw device mapping), 211-219
  - SIOC (Storage I/O Control), 206-208
  - storage path selection, 210
- Ostrow, Howard, 331
- overhead, 157
- oversubscription, memory, 18

## P

- packets, dropped, monitoring, 177-181
- page tables, nested/extended, 19
- paravirtualized SCSI adapters, 78
- partitions, alignment, 220-222
- Partner Exchange (PEX), 330
- PCI (Payment Card Industry) standards, 48
- Pearson, Brian, 331
- Performance Analyzer, 156
- performance management, 157, 239-240
  - BIOS settings, 58
  - DBA role, 158-159
  - memory, 164-167
    - AMM (Automatic Memory Management), 168-169

- Huge Pages, 167-172
- NUMA (non-uniform memory access), 172-175
- TPS (Transparent Page Sharing), 172
- networking, 175-176
  - dropped packets, 177-181
  - network load testing, 176
- processing power, 159-164
- storage, 181, 219-220
  - ASM (Automatic Storage Management), 225-228
  - block alignment, 220-222
  - LSI controllers, 222
  - LVMs (logical volume managers), 225-228
  - NFS (Network File System), 185-190
  - pvSCSI controllers, 222-224
  - SCSI queues, 182-185, 191-219
- terminology, 157-158
- vCOPS, 228-238
- vSphere, 79-81
- performance studies, 36
- performance tuning, 10, 155
- PGA (process global area), 165-167
- platforms, consolidating, 17-19
- process global area (PGA), 165-167
- processes, operating systems, 59-60
- protecting databases against downtime, 98-100
- provisioning, 289-291
  - application blueprints, 310
- pvSCSI controllers, 222-224

## Q-R

---

QoS (quality of service), applications, improving, 41

RAC (Real Application Clusters), 1, 55, 103-104, 139, 333-335

- architecture, 123-124
- deploying on vSphere, 92-98
- installation, 130-132
- mega vMotion-RAC functional stress test, 135-138
- multimode, implementing, 90-91
- nodes, transitioning between hosts, 100-103

- node vMotion test, 132-135
- One Node, implementing, 89-90
- virtual machines, 89-98
- virtualizing, 51-52
- workload characterization studies, 105, 121-129

RAC Deployment Guide, 327

RAID, ORION, 209-210

RAT (Real Application Testing), 30

raw device mapping (RDM). *See* RDM (raw device mapping)

RDBMS (relational database management system), 11, 37, 104, 243

RDM (raw device mapping), 282, 334

- detriments, 211
- ORION, 212-219
- RDM-P, 211-212
- RDM-V, 211-212
- versus VMFS (Virtual Machine File System), 76-77, 282-284

ready time, CPUs, 160-164

Real Application Clusters (RAC). *See* RAC (Real Application Clusters)

Real Application Testing (RAT), 30

RecoverPoint, 256

RecoverPoint with vCenter SRM, 256

recovery, 286, 338

- Avamar, 279-281
- DR (disaster recovery) plan
  - configuring virtual machines in, 261
  - creating, 257-264
  - editing, 261
  - exporting, 264
  - removing, 261
  - running, 263-264
  - testing, 262
- essential components, 268
- FRA (fast recovery area), 255
- NetApp, Snapshot, 277-279
- principles, 270-274
- Recovery Manager, 268
- SRM (Site Recovery Manager), 241-245
  - breaking connections, 257
  - configure connections, 257
  - exporting system logs, 257
  - protecting new virtual machines, 258-259
  - removing virtual machines, 260

- testing, 251-253
- unprotecting virtual machines, 260
- Recovery Manager (RMAN), 10, 268
- recovery point objective (RPO), 241
- recovery time objective (RTO), 241
- redo logs, 255
- relapsed time, 157
- relational database management system (RDBMS), 11, 37, 104, 243
- replication
  - application-based, 248
  - array-based, 258-264
  - combining DG and VR, 250-251
  - DG (Data Guard), 248-249
    - repairing logical data block corruption, 249
  - SRM (Site Recovery Manager), testing, 251-253
  - storage array-based, 247
  - vSphere, 253-255
  - VR (vSphere Replication), 245-246
- reports
  - AWR (Automatic Workload Repository), 155
  - Oracle Enterprise Manager, 155
- reprovisioning, 26-27
- resource pools, VMware, 42
- restoring. *See* recovery
- rEsxtop, display, 188
- RMAN (Recovery Manager), 10, 271, 279
- rolling view widget, Oracle Database
  - dashboard, configuring, 235
- RPO (recovery point objective), 241
- RTO (recovery time objective), 241

## S

- SAN Volume Controller (SVC), 5-6
- Sarbanes-Oxley (SOX) Act, 48
- scalability, 157
  - designing on demand, 42-44
- scale-out charts, multiple virtual machines, 116-117
- SCSI controllers, virtual, 76
- SCSI queues, 182-185
  - ESX/ESXi hosts and VMs, 191-219
  - Iometer, 198-203
- NFS, 183
- ORION (Oracle I/O Numbers), 203-208
  - storage benchmarking VMDK, 193-196
  - storage path throughput, 192-193
- SDDC (software-defined datacenter), 4, 7
- security, databases, 9
- server architecture, databases, 11-12
- server environments, databases, provisioning, 21-22
- servers
  - database, consolidating, 48-51
  - sizing vendor guidelines, 35
  - sprawl, minimizing, 27
- service level agreements (SLAs), 7, 19
- service-oriented architectures (SOAs), 26
- SGA (system global area), 12, 165-166
- SICO (Storage I/O Control), 206-208
- Silly Little Oracle Benchmark (SLOB), 30
- single-instance workload study, 106-121
- Site Recovery Manager (SRM). *See* SRM (Site Recovery Manager)
- SLAs (service level agreements), 7, 19
  - meeting demands, 27-30
- SLOB (Silly Little Oracle Benchmark), 30
- Snap Creator (NetApp), 276
- Snap Manager (NetApp), 276
- Snapshot (NetApp), 274-275
  - backing up virtualized databases, 275-276
  - restoring virtualized databases, 277-279
- snapshots, VMware, 228
- SOAs (service-oriented architectures), 26
- software support, VMware environments, 143-147
- software support and licensing, Oracle, 141-143
- software-defined datacenter (SDDC), 4, 7
- source environment, Oracle database lab
  - architecture, 296
- spindle busy time, 190
- SRAs (storage replication adapters), 21, 247
- SRM (Site Recovery Manager), 21, 241-245, 338
  - breaking connections, 257
  - configure connections, 257
  - exporting system logs, 257
  - protecting datacenters, 46-47
  - protecting new virtual machines, 258-259

- removing virtual machines, 260
- testing, 251-253
- unprotecting virtual machines, 260
- stakeholders, identifying key, 53-54
- storage
  - best practices, 68-76
  - categorizing virtualization technologies, 71
  - benchmark tools, 31, 193-196
  - DBaaS (database-as-a-service),
    - allocating, 313
  - layout, 114-115, 126-129
  - networks, 178
  - path throughput, SCSI queues, 192-193
  - SRAs (storage replication adapters), 21, 247
  - vendor backup solutions, 273-274
  - virtualization, 5
- storage array-based replication, vSphere, 253-255
- Storage I/O Control (SIOC), 206-208
- storage performance, 219-220
  - ASM, 225-228
  - block alignment, 220-222
  - LSI controllers, 222
  - LVMs, 225-228
  - monitoring, 181
    - NFS, 185-190
    - SCSI queues, 182-185
  - pvSCSI controllers, 222-224
  - SCSI queues, 191-219
- storage replication adapters (SRAs), 21, 247
- stress tests, 36
  - mega vMotion-RAC functional stress, 135-138
- stretch clusters, EMC VPLEX, 6-7
- studies
  - RAC workload characterization, 121-129
  - single-instance workload, 106-121
- subcapacity licensing, 153
- subcluster licensing, 151-152
- support, 335
  - Oracle, 145-146
- SVC (SAN Volume Controller), 5-6
- Swingbench large-scale order entry
  - benchmark kit, 30, 122-124

- switches, VDS (Virtual Distributed Switch), 336
- Symantec AppHA, protecting applications
  - with, 86-87
- system global area (SGA), 12, 165-166
- system logs, vCenter SRM, exporting, 257

---

## T

- TAF (Transparent Application Failover), 103
- target environment, Oracle database lab architecture, 296
- teaming, 178
- Technical Software Alliance Network (TSAnet), 335
- temp files, 255
- testing
  - recovery plans, 262
  - SRM vSphere replication, 251-253
- tests
  - stress, 36
    - mega vMotion-RAC functional stress, 135-138
    - Oracle RAC node vMotion, 132-135
    - single-instance workload study, 106-121
    - twenty-four hour workload, 131-132
- thin provisioning, VMDKs, 225
- throughput, 157
- timekeeping, virtual machines, 81
- TPS (Transparent Page Sharing), 172
- Transparent Application Failover (TAF), 103
- Transparent Page Sharing (TPS), 172
- troubleshooting databases, 10
- TSAnet (Technical Software Alliance Network), 335
- twenty-four hour workload test, 131-132

---

## U-V

- University of British Columbia (UBC) case study, 326
- upgrading ESX/ESXi and vSphere, 60-61

- VAAI (vSphere Storage APIs - Array Integration), 226-227
- vCAC (vCloud Automation Center), 21-22, 338
  - configuring application blueprints, 302-310
- vCenter
  - environment, Oracle database lab architecture, 296
  - Operations Manager Suite (vCops), 25
  - SRM (Site Recovery Manager), 241-245
    - breaking connections, 257
    - configure connections, 257
    - exporting system logs, 257
    - protecting new virtual machines, 258-259
    - removing virtual machines, 260
    - testing, 251-253
    - unprotecting virtual machines, 260
- vCloud Automation Center (vCAC), 21-22, 338
- vCloud Hybrid Service, 338
- vCOPS (vCenter Operations Manager), 25, 156, 228-238
  - Adapter for Hyperic Oracle Plug-in, 229
  - Database server metrics, 230
  - OEM (Oracle Enterprise Manager) adapter, 228
- vCPUs (virtual CPUs), 4
  - best practices, 65-68
  - hot-add, 119-120
  - ready time, 160-164
  - versus CPUs, 159-164
- vDBA (virtual DBA), 290
- VDPA (VMware Data Protection Advanced), 273, 282
- vDS (virtual distributed switch), 4, 336
- verification, dNFS, 218-219
- vFabric Application Director, 21
- vFabric Data Director, 26
- virtual CPUs (vCPUs). *See* vCPUs (virtual CPUs)
- virtual DBA (vDBA), 290
- virtual distributed switch (vDS), 4
- virtual machine disks (VMDKs). *See* VMDKs (virtual machine disks)
- Virtual Distributed Switch (VDS), 336
- virtualization, 1, 12
  - ESXi, 3
  - Oracle Database, 291-302
    - migration, 296-302
  - Oracle Database
    - business scenario, 294-295
    - facilitating deployments, 294
    - lab architecture, 295-296
  - SDDC (software-defined datacenter), 4, 7
  - storage, 5
  - vSphere, 3
- virtualized databases
  - backing up
    - NetApp Snapshot, 275-276
    - NetApp storage, 277
    - NFS datastores, 277
  - restoring, NetApp Snapshot, 277-279
- virtualized environment, databases, vSphere HA, 84-86
- Virtualizing SQL Server with VMware: Doing It Right, 173
- virtual machine disks (VMDKs). *See* VMDKs (virtual machine disks)
- Virtual Machine File System (VMFS). *See* VMFS (Virtual Machine File System)
- virtual machines, 4
  - live migration, 117-119
  - memory reservation, 18
  - protecting new, 259
  - protecting new, 258-259
  - RAC (Real Application Clusters), 89-98
    - architecture, 123-124
  - removing, 260
  - scale-out chart, 116-117
  - scale-out test, 113
  - single-instance workload study, single scale-up results, 111-112
  - storage layout, 114-115
  - timekeeping, 81
  - unprotecting, 260
- virtual memory, 4
- virtual network interface card (vNIC), 4
- Virtual Networking Concepts and Best Practices, 79
- virtual provisioning, ASM disk groups, 255-256

- virtual Real Application Clusters DBA (vRAC-DBA), 290
- virtual storage-area network (vSAN), 339
- VM networks, 178
- VMDKs (virtual machine disk), 4, 71, 184
  - thin provisioning, 225
  - storage benchmarking, 193-196
- VMFS (Virtual Machine File System), 43, 71, 282-334, 337
  - dynamic expansion, 223
  - versus RDM (Raw Device Mappings), 76-77, 282-284
- vMotion (VMware), 23, 43, 103-104, 117
  - HA (high availability), 121-122
  - mega vMotion-RAC functional stress test, 135-138
  - network bandwidth, 177-178
  - RAC node vMotion test, 132-135
  - transitioning RAC nodes between hosts, 100-103
- vMSC (vSphere Metro Storage Clusters), 338
- VMware
  - advising customers of, 149
  - clusters, 42
  - designing databases on, 41-47
  - DRS (Distributed Resource Scheduler), 43-44
  - HA (high availability), 121-122
  - hosts, 42
  - hot-add feature, 42
  - IT (Information Technologies), 327-330
  - obtaining VMware support, 146
  - resource pools, 42
  - snapshots, 228
  - solving deployment and management issues, 25-26
  - VAAI, 226
  - vMotion, 23, 43, 103-104, 117
    - HA (high availability), 121-122
    - mega vMotion-RAC functional stress test, 135-138
    - network bandwidth, 177-178
    - RAC node vMotion test, 132-135
    - transitioning RAC nodes between hosts, 100-103
  - VMware Data Protection Advanced (VDPA), 273, 281-282
  - VMware on VMware, 327
  - VMware Partner Exchange (PEX), 330
  - VMworld, 330
  - vNIC (virtual network interface card), 4
  - VR (vSphere Replication), 245-246
    - combining with DG (Data Guard), 250-251
  - vRAC-DBA (virtual Real Application Clusters DBA), 290
  - vSAN (virtual storage-area network), 339
  - vSphere, 333
    - best practices, 334
    - certification, 144-145
    - integrating with NetApp, 274
    - maximizing HA capabilities, 45
    - monitoring performance on, 79-81
    - storage array-based replication, 253-255
    - supporting large pages, 64-65
    - upgrading, 60-61
    - virtualization, 3
  - vSphere Metro Storage Clusters (vMSC), 338
  - vSphere Replication (VR), 245-246
  - vSphere Storage APIs - Array Integration (VAAI), 226-227
  - vSwitches, configuring, 125

## W-Z

- Welch, Dave, 331
- Williams, Chris, 331
- workload, 333
  - databases, sizing, 30-31
  - OLTP (online transaction processing) loads, 106-107
  - RAC workload characterization study, 121-129
  - single-instance workload study, 106-121