

vmware® PRESS



Virtualizing SQL Server with VMware®

Doing IT Right

Michael Corey
Jeff Szastak
Michael Webster

Foreword by **Duncan Epping**



Virtualizing SQL Server with VMware®

Doing IT Right


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 forty-two languages, including, but not limited to, Chinese (Simplified), Chinese (Traditional), French, German, Greek, Hindi, Japanese, Korean, Polish, Russian, and Spanish.

For more information about VMware Press, please visit vmwarepress.com.

vmware® PRESS



vmwarepress.com

Complete list of products • User Group Info • 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, eBooks, certification 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 certification exams.



Make sure to connect with us!
vmwarepress.com

vmware®

PEARSON
IT CERTIFICATION

Safari®
Books Online

This page intentionally left blank

Virtualizing SQL Server with VMware®

Doing IT Right

Michael Corey
Jeff Szastak
Michael Webster

vmware® PRESS

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

Virtualizing SQL Server with VMware®

Copyright © 2015 VMware, Inc.

Published by Pearson plc

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-13: 978-0-321-92775-0

ISBN-10: 0-321-92775-3

Library of Congress Control Number: 2014941961

Printed in the United States of America

First Printing August 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 CD 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 or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact international@pearsoned.com.

ASSOCIATE PUBLISHER

David Dusthimer

ACQUISITIONS EDITOR

Joan Murray

VMWARE PRESS

PROGRAM MANAGER

David Nelson

DEVELOPMENT EDITOR

Eleanor C. Bru

MANAGING EDITOR

Sandra Schroeder

PROJECT EDITOR

Mandie Frank

COPY EDITOR

Bart Reed

PROOFREADER

Sarah Kearns

INDEXER

Erika Millen

EDITORIAL ASSISTANT

Vanessa Evans

DESIGNER

Chuti Prasertsith

COMPOSITOR

Tricia Bronkella

This book is dedicated to my wife of 28 years, Juliann, who has supported me in every way possible, and my three children, John, Anmmarie, and Michael.

-Michael Corey

This book is dedicated to my wife, Heather, and my three children, Wyatt, Oliver, and Stella.

-Jeff Szastak

This book is dedicated to my wife, Susanne, and my four sons, Sebastian, Bradley, Benjamin, and Alexander, for their ongoing support. I also dedicate this book to the VMware community.

-Michael Webster

This page intentionally left blank

Contents

Foreword xvii

Preface xix

About the Authors xxiii

About the Technical Reviewer xxv

Acknowledgments xxvii

Reader Services xxix

1 Virtualization: The New World Order? 1

Virtualization: The New World Order 1

Virtualization Turns Servers into Pools of Resources 3

Living in the New World Order as a SQL Server DBA 3

A Typical Power Company 6

Summary 7

2 The Business Case for Virtualizing a Database 9

Challenge to Reduce Expenses 9

The Database Administrator (DBA) and Saving Money 10

Service Level Agreements (SLA) and the DBA 11

Avoiding the Good Intention BIOS Setting 12

DBAs' Top Reasons to Virtualize a Production Database 13

High Availability and Database Virtualization 14

Performance and Database Virtualization 16

Provisioning/DBaaS and Database Virtualization 17

Hardware Refresh and Database Virtualization 20

Is Your Database Too Big to Virtualize? 22

Summary 23

3 Architecting for Performance: The Right Hypervisor 25

What Is a Hypervisor? 25

Hypervisor Is Like an Operating System 26

What Is a Virtual Machine? 28

Paravirtualization 29

The Different Hypervisor Types 29

Type-1 Hypervisor 30

Type-2 Hypervisor 31

Paravirtual SCSI Driver (PVSCSI) and VMXNET3 31

Installation Guidelines for a Virtualized Database 32

It's About Me, No One Else But Me 33

Virtualized Database: It's About Us, All of Us 34

DBA Behavior in the Virtual World	34
Shared Environment Means Access to More If You Need It	35
Check It Before You Wreck It	36
Why Full Virtualization Matters	36
Living a DBA's Worst Nightmare	37
Physical World Is a One-to-One Relationship	38
One-to-One Relationship and Unused Capacity	38
One to Many: The Virtualized World	40
The Right Hypervisor	40
Summary	41
4 Virtualizing SQL Server: Doing IT Right	43
Doing IT Right	43
The Implementation Plan	44
Service-Level Agreements (SLAs), RPOs, and RTOs	45
Baselining the Existing vSphere Infrastructure	46
Baselining the Current Database Workload	48
Bird's-Eye View: Virtualization Implementation	50
How a Database Virtualization Implementation Is Different	51
Summary	55
5 Architecting for Performance: Design	57
Communication	58
Mutual Understanding	59
The Responsibility Domain	60
Center of Excellence	61
Deployment Design	63
SQL Workload Characterization	64
Putting It Together (or Not)	65
Reorganization	68
Tiered Database Offering	70
Physical Hardware	73
CPU	74
Memory	76
Virtualization Overhead	76
Swapping, Paging? What's the Difference?	78
Large Pages	79
NUMA	79
Hyper-Threading Technology	85
Memory Overcommitment	87

Reservations	87
SQL Server: Min/Max	90
SQL Server: Lock Pages in Memory	92
Storage	93
Obtain Storage-Specific Metrics	94
LSI Logic SAS or PVSCSI	94
Determine Adapter Count and Disk Layout	95
VMDK versus RDM	96
VMDK Provisioning Type	96
Thin Provisioning: vSphere, Array, or Both?	98
Data Stores and VMDKs	99
VMDK File Size	100
Networking	100
Virtual Network Adapter	100
Managing Traffic Types	101
Back Up the Network	103
Summary	104

6 Architecting for Performance: Storage 105

The Five Key Principles of Database Storage Design	106
Principle 1: Your database is just an extension of your storage	106
Principle 2: Performance is more than underlying storage devices	107
Principle 3: Size for performance before capacity	107
Principle 4: Virtualize, but without compromise	108
Principle 5: Keep it standardized and simple (KISS)	109
SQL Server Database and Guest OS Storage Design	109
SQL Server Database File Layout	110
Number of Database Files	110
Size of Database Files	114
Instant File Initialization	120
SQL Server File System Layout	122
SQL Server Buffer Pool Impact on Storage Performance	129
Updating Database Statistics	130
Data Compression and Column Storage	132
Database Availability Design Impacts on Storage Performance	135
Volume Managers and Storage Spaces	136
SQL Server Virtual Machine Storage Design	136
Virtual Machine Hardware Version	137
Choosing the Right Virtual Storage Controller	138
Choosing the Right Virtual Disk Device	143

SQL Virtual Machine Storage Layout	152
Expanding SQL Virtual Machine Storage	158
Jumbo VMDK Implications for SQL Server	159
vSphere Storage Design for Maximum SQL Performance	164
Number of Data Stores and Data Store Queues	165
Number of Virtual Disks per Data Store	170
Storage IO Control—Eliminating the Noisy Neighbor	173
vSphere Storage Policies and Storage DRS	177
vSphere Storage Multipathing	184
vSphere 5.5 Failover Clustering Enhancements	185
RAID Penalties and Economics	187
SQL Performance with Server-Side Flash Acceleration	198
VMware vSphere Flash Read Cache (vFRC)	199
Fusion-io ioTurbine	201
PernixData FVP	204
SQL Server on Hyperconverged Infrastructure	207
Summary	213

7 Architecting for Performance: Memory 217

Memory	218
Memory Trends and the Stack	218
Database Buffer Pool and Database Pages	219
Database Indexes	222
Host Memory and VM Memory	225
Mixed Workload Environment with Memory Reservations	226
Transparent Page Sharing	228
Internet Myth: Disable Memory TPS	229
Memory Ballooning	230
Why the Balloon Driver Must Run on Each Individual VM	232
Memory Reservation	232
Memory Reservation: VMware HA Strict Admission Control	233
Memory Reservations and the vswap File	233
SQL Server Max Server Memory	234
SQL Server Max Server Memory: Common Misperception	235
Formula for Configuring Max Server Memory	236
Large Pages	237
What Is a Large Page?	237
Large Pages Being Broken Down	238
Lock Pages in Memory	239
How to Lock Pages in Memory	241

Non-Uniform Memory Access (NUMA)	241
vNUMA	243
Sizing the Individual VMs	244
More VMs, More Database Instances	244
Thinking Differently in the Shared-Resource World	246
SQL Server 2014 In-Memory Built In	246
Summary	247

8 Architecting for Performance: Network 249

SQL Server and Guest OS Network Design	250
Choosing the Best Virtual Network Adapter	250
Virtual Network Adapter Tuning	252
Windows Failover Cluster Network Settings	254
Jumbo Frames	256
Configuring Jumbo Frames	259
Testing Jumbo Frames	262
VMware vSphere Network Design	264
Virtual Switches	265
Number of Physical Network Adapters	267
Network Teaming and Failover	270
Network I/O Control	274
Multi-NIC vMotion	276
Storage Network and Storage Protocol	279
Network Virtualization and Network Security	281
Summary	286

9 Architecting for Availability: Choosing the Right Solution 287

Determining Availability Requirements	287
Providing a Menu	288
SLAs, RPOs, and RTOs	290
Business Continuity vs. Disaster Recovery	291
Business Continuity	291
Disaster Recovery	291
Disaster Recovery as a Service	292
vSphere High Availability	294
Hypervisor Availability Features	294
vMotion	296
Distributed Resource Scheduler (DRS)	297
Storage vMotion	297
Storage DRS	297

Enhanced vMotion X-vMotion	298
vSphere HA	298
vSphere App HA	299
vSphere Data Protection	300
vSphere Replication	300
vCenter Site Recovery Manager	301
VMware vCloud Hybrid Service	302
Microsoft Windows and SQL Server High Availability	302
ACID	302
SQL Server AlwaysOn Failover Cluster Instance	304
SQL Server AlwaysOn Availability Groups	306
Putting Together Your High Availability Solution	308
Summary	310
10 How to Baseline Your Physical SQL Server System	311
What Is a Performance Baseline?	312
Difference Between Performance Baseline and Benchmarks	315
Using Your Baseline and Your Benchmark to Validate Performance	318
Why Should You Take a Performance Baseline?	319
When Should You Baseline Performance?	320
What System Components to Baseline	320
Existing Physical Database Infrastructure	321
Database Application Performance	323
Existing or Proposed vSphere Infrastructure	325
Comparing Baselines of Different Processor Types and Generations	328
Comparing Different System Processor Types	328
Comparing Similar System Processor Types Across Generations	330
Non-Production Workload Influences on Performance	331
Producing a Baseline Performance Report	332
Performance Traps to Watch Out For	333
Shared Core Infrastructure Between Production and Non-Production	333
Invalid Assumptions Leading to Invalid Conclusions	334
Lack of Background Noise	334
Failure to Considering Single Compute Unit Performance	335
Blended Peaks of Multiple Systems	335
vMotion Slot Sizes of Monster Database Virtual Machines	336
Summary	337

11	Configuring a Performance Test—From Beginning to End	339
	Introduction	339
	What We Used—Software	341
	What You Will Need—Computer Names and IP Addresses	341
	Additional Items for Consideration	342
	Getting the Lab Up and Running	342
	VMDK File Configuration	345
	VMDK File Configuration Inside Guest Operating System	352
	Memory Reservations	355
	Enabling Hot Add Memory and Hot Add CPU	356
	Affinity and Anti-Affinity Rules	358
	Validate the Network Connections	359
	Configuring Windows Failover Clustering	359
	Setting Up the Clusters	362
	Validate Cluster Network Configuration	368
	Changing Windows Failover Cluster Quorum Mode	369
	Installing SQL Server 2012	374
	Configuration of SQL Server 2012 AlwaysOn Availability Groups	387
	Configuring the Min/Max Setting for SQL Server	392
	Enabling Jumbo Frames	393
	Creating Multiple tempdb Files	394
	Creating a Test Database	396
	Creating the AlwaysOn Availability Group	399
	Installing and Configuring Dell DVD Store	406
	Running the Dell DVD Store Load Test	430
	Summary	436
	Appendix A Additional Resources	437
	Additional Documentation Sources	437
	User Groups	440
	VMUG: The VMware Users Group	440
	PASS: Professional Association of SQL Server	441
	VMware Community	442
	Facebook Groups	443
	Blogs	444
	Twitter: 140 Characters of Real-Time Action	445
	Index	447

This page intentionally left blank

Foreword

About 10 years ago, I started a new job. The company I started working for had a couple hundred physical servers at the time. When several new internal software development projects started, we needed to expand quickly and added dozens of new physical servers. Pretty soon we started hitting all the traditional datacenter problems, such as lack of floor space, high power consumption, and cooling constraints. We had to solve our problems, and during our search for a solution we were introduced to a new product called VMware ESX and Virtual Center. It didn't take long for us to see the potential and to start virtualizing a large portion of our estate.

During this exercise, we started receiving a lot of positive feedback on the performance of the virtualized servers. On top of that, our application owners loved the fact that we could deploy a new virtual machine in hours instead of waiting weeks for new hardware to arrive. I am not even talking about all the side benefits, such as VMotion (or vMotion, as we call it today) and VMware High Availability, which provided a whole new level of availability and enabled us to do maintenance without any downtime for our users.

After the typical honeymoon period, the question arose: What about our database servers? Could this provide the same benefits in terms of agility and availability while maintaining the same performance? After we virtualized the first database server, we quickly realized that just using VMware Converter and moving from physical to virtual was not sufficient, at least not for the databases we planned to virtualize.

To be honest, we did not know much about the database we were virtualizing. We didn't fully understand the CPU and memory requirements, nor did we understand the storage requirements. We knew something about the resource consumption, but how do you make a design that caters to those requirements? Perhaps even more importantly, where do you get the rest of the information needed to ensure success?

Looking back, I wish we'd had guidance in any shape or form that could have helped along our journey—guidance that would provide tips about how to gather requirements, how to design an environment based on these requirements, how to create a performance baseline, and what to look for when hitting performance bottlenecks.

That is why I am pleased Jeff Szastak, Michael Corey, and Michael Webster took the time to document the valuable lessons they have learned in the past few years about virtualizing tier 1 databases and released it through VMware Press in the form of this book you are about to read. Having gone through the exercise myself, and having made all the mistakes mentioned in the book, I think I am well qualified to urge you to soak in all this valuable knowledge to ensure success!

Duncan Epping

Principal Architect, VMware

Yellow-Bricks.com

This page intentionally left blank

Preface

As we traveled the globe presenting on how to virtualize the most demanding business-critical applications, such as SQL Server, Oracle, Microsoft Exchange, and SAP, it became very clear that there was a very real and unmet need from the attendees to learn how to virtualize these most demanding applications correctly.

This further hit home when we presented at the VMworld conferences in San Francisco and Barcelona. At each event, we were assigned a very large room that held over 1,800 people; within 48 hours of attendees being able to reserve a seat in the room, it was filled to capacity. We were then assigned a second large room that again filled up within 24 hours.

Recognizing that the information we had among the three of us could help save countless others grief, we decided to collaborate on this very practical book.

Target Audience

Our goal was to create in one book—a comprehensive resource that a solution architect, system administrator, storage administrator, or database administrator could use to guide them through the necessary steps to successfully virtualize a database. Many of the lessons learned in this book apply to any business-critical application being virtualized from SAP, E-Business Suite, Microsoft Exchange, or Oracle, with the specific focus of this book on Microsoft SQL Server. Although you don't have to be a database administrator to understand the contents of this book, it does help if you are technical and have a basic understanding of vSphere.

Approach Taken

Everything you need to succeed in virtualizing SQL Server can be found within the pages of this book. By design, we created the book to be used in one of two ways. If you are looking for a comprehensive roadmap to virtualize your mission-critical databases, then follow along in the book, chapter by chapter. If you are trying to deal with a particular resource that is constraining the performance of your database, then jump to Chapters 5 through 8.

At a high level, the book is organized as follows:

- Chapters 1 and 2 explain what virtualization is and the business case for it. If you are a database administrator or new to virtualization, you will find these chapters very helpful; they set the stage for why virtualizing your databases is “doing IT right.”

- Chapters 3 through 9 are the roadmap you can follow to successfully virtualize the most demanding of mission-critical databases. Each chapter focuses on a particular resource the database utilizes and how to optimize that resource to get the best possible performance for your database when it is virtualized. We purposely organized this section into distinct subject areas so that you can jump directly to a particular chapter of interest when you need to brush up. We expect that you will periodically return to Chapters 5 through 8 as you are fine-tuning the virtualized infrastructure for your mission-critical databases.
- The last two chapters walk you through how to baseline the existing SQL Server database so that you adequately determine the resource load it will put onto the virtualized infrastructure. In these chapters, we also provide detailed instructions on how to configure a stress test.

Here are the three major sections of the book with the associated chapters:

What Virtualization Is and Why You Should Do It

In this section, the reader will learn about the benefits of virtualization and why the world is moving towards 100% virtualization. The reader will learn the benefits of breaking the bond between hardware and software, and the benefits this brings to the datacenter and why virtualization is a better way to do IT.

Chapter 1: Virtualization: The New World Order?

Chapter 2: The Business Case for Virtualizing a Database

Optimizing Resources in a Virtualized Infrastructure

In Chapters 3-9, the reader will gain knowledge on how to properly architect and implement virtualized SQL Server. The reader will start off learning how to put together a SQL Server virtualization initiative, and then dive into an in-depth discussion on how to architect SQL Server on a vSphere platform. This section includes deep dives on storage, memory, networking, and high availability.

Chapter 3: Architecting for Performance: The Right Hypervisor

Chapter 4: Virtualizing SQL Server: Doing IT Right

Chapter 5: Architecting for Performance: Design

Chapter 6: Architecting for Performance: Storage

Chapter 7: Architecting for Performance: Memory

Chapter 8: Architecting for Performance: Network

Chapter 9: Architecting for Availability: Choosing the Right Solution

How to Baseline and Stress Test

The final two chapters walk the reader through the importance of setting up a baseline for their virtualized SQL Server implementation. Chapter 10 speaks to the why and the how of baselining, which is critical to successfully virtualizing SQL Server. In the final chapter, the reader will put all the knowledge presented in the previous chapters together and will be walked through a beginning-to-end configuration of SQL Server 2012 with AlwaysOn Availability Groups running on Windows Server 2012 on a vSphere 5.5 infrastructure.

Chapter 10: How to Baseline Your Physical SQL Server System

Chapter 11: Configuring a Performance Test—From Beginning to End

A database is one of the most resource-intensive applications you will ever virtualize, and it is our sincere intention that with this book as your guide, you now have a roadmap that will help you avoid the common mistakes people make—and more importantly, you will learn how to get optimal performance from your virtualized database.

We want to thank you for buying our book, and we hope after you read it that you feel we have achieved our goal of providing you with a comprehensive resource on how to do IT right. Feel free to reach out to us with any questions, suggestions, or feedback you have.

Michael Corey (@Michael_Corey) Michael.corey@ntirety.com

Jeff Szastak (@Szastak)

Michael Webster (@vcdxnz001)

This page intentionally left blank

About the Authors



Michael Corey (@Michael_Corey) is the President of Ntirety, a division of Hosting. Michael is an experienced entrepreneur and a recognized expert on relational databases, remote database administration, and data warehousing. Microsoft named Michael a SQL Server MVP, VMware named him a vExpert, and Oracle named him an Oracle Ace. Michael has presented at technical and business conferences from Brazil to Australia. Michael is a past president of the Independent Oracle Users Group; he helped found the Professional Association of SQL Server, is a current board member of the IOUG Cloud SIG, and is actively involved in numerous professional associations and industry user groups. Michael currently sits on the executive committee for the Massachusetts Robert H. Goddard Council for Science, Technology, Engineering, and Mathematics.



Jeff Szastak (@Szastak) is currently a Staff Systems Engineer for VMware. Jeff has been with VMware for over six years, holding various roles with VMware during his tenure. These roles have included being a TAM, Systems Engineer Specialist for Business-Critical Applications, Enterprise Healthcare Systems Engineer, and a CTO Ambassador. Jeff is a recognized expert for virtualizing databases and other high I/O applications on the vSphere platform. Jeff is a regular speaker at VMworld, VMware Partner Exchange, VMware User Groups, and has spoken at several SQL PASS events. Jeff holds a Master of Information Assurance degree as well as the distinguished CISSP certification. Jeff has over 13 “lucky” years in IT and is passionate about helping others find a better way to do IT.



Michael Webster (@vcdxnz001) is based in Auckland, New Zealand. He is a VMware Certified Design Expert (VCDX #66), author of long-whiteclouds.com (a top-15 virtualization blog), and a Top 10 VMworld Session Speaker for 2013. In addition, he is a Senior Solutions and Performance Engineer for Nutanix, vExpert, MCSE, and NPP. Michael specializes in solution architecture and performance engineering for Unix-to-VMware migrations as well as virtualizing business-critical applications such as SQL, Oracle, SAP, Exchange, Enterprise Java Systems, and monster VMs in software-defined data centers. Michael has more than 20 years experience in the IT industry and 10 years experience deploying VMware solutions in large-scale environments around the globe. He is regularly a presenter at VMware VMworld, VMware vForums, VMware User Groups, and other industry events. In addition to this book, Michael was technical reviewer of *VCDX Boot Camp* and *Virtualizing and Tuning Large-Scale Java Platforms*, both published by VMware Press.

This page intentionally left blank

About the Technical Reviewer

Mark Achtemichuk (VCDX #50) is currently a Senior Technical Marketing Architect, specializing in Performance, within the SDDC Marketing group at VMware. Certified as VCDX #50, Mark has a strong background in data center infrastructures and cloud architectures, experience implementing enterprise application environments, and a passion for solving problems. He has driven virtualization adoption and project success by methodically bridging business with technology. His current challenge is ensuring that performance is no longer a barrier, perceived or real, to virtualizing an organization's most critical applications on its journey to the software-defined data center.

This page intentionally left blank

Acknowledgments

We would like to thank the entire team at VMware Press for their support throughout this project and for helping us get this project across the line—especially Joan Murray for her constant support and encouragement. We would like to thank our editorial team. Thank you Ellie Bru and Mandie Frank for your attention to detail to make sure we put out a great book, and last but not least, we would especially like to thank our technical reviewer, Mark Achtemichuk (VCDX #50).

Michael Corey

Anyone who has ever written a book knows first hand what a tremendous undertaking it is and how stressful it can be on your family. It is for that reason I thank my wife of 28 years, Juliann. Over those many years, she has been incredible. I want to thank my children, Annmarie, Michael, and especially John, who this particular book was hardest on. John will know why if he reads this.

Jeff and Michael, my co-authors, are two of the smartest technologists I have ever had the opportunity to collaborate with. Thank you for making this book happen despite the many long hours it took you away from your families. Mark Achtemichuk, our technical reviewer, rocks! He helped take this book to a whole new level. To my friends at VMware—Don Sullivan, Kannan Mani, and Sudhir Balasubramanian—thank you for taking all my late-night emails and phone calls to discuss the inner workings of vSphere. To the publishing team at Pearson, what can I say? Thank you Joan Murray for believing and making this book possible.

Special thanks go to my Ntirety family—Jim Haas, Terrie White, and Andy Galbraith are all three incredible SQL Server technologists. And special thanks to people like David Klee and Thomas LaRock and to the entire SQL Server community. Every time I attend a SQLSaturday event, I always think how lucky I am to be party of such a special community of technologist who care a lot and are always willing to help.

Jeff Szastak

I would like to thank my loving wife, Heather, for her love, support, and patience during the writing of this book. I want to thank my children, Wyatt, Oliver, and Stella, for it is from you I draw inspiration. A huge thank-you to Hans Drolshagen for the use of his lab during the writing of this book! And thanks to my mentor, Scott Hill, who pushed me, challenged me, and believed in me. Thanks for giving a guy who couldn't even set a DHCP address a job in IT, Scott.

Finally, I would like to thank the VMware community. Look how far we have come. I remember the first time I saw a VMware presentation as a customer and thought, “If this software works half as well as that presentation says it does, this stuff will change the world.” And it has, because of you, the VMware community.

Michael Webster

I’d like to thank my wife, Susanne, and my four boys, Sebastian, Bradley, Benjamin, and Alexander, for providing constant love and support throughout this project and for putting up with all the long hours on weeknights and weekends that it required to complete this project. I would also like to acknowledge my co-authors, Michael and Jeff, for inviting me to write this book with them. I am extremely thankful for this opportunity, and it has been a fantastic collaborative process. Finally, I’d like to thank and acknowledge VMware for providing the constant inspiration for many blog articles and books and for creating a strong and vibrant community. Also, thanks go out to my sounding boards throughout this project: Kasim Hansia, VMware Strategic Architect and SAP expert, Cameron Gardiner, Microsoft Senior Program Manager Azure and SQL, and Josh Odgers (VCDX #90), Nutanix Senior Solutions and Performance Architect. Your ideas and support have added immeasurable value to this book and the IT community as a whole.

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 us directly 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, email address, and phone number. We will carefully review your comments and share them with the author and editors who worked on the book.

Email: 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/9780321927750 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

Architecting for Performance: Storage

All aspects of architecting your SQL Server Database for performance are important. Storage is more important than most when compared to the other members of the IT Food Group family we introduced in Chapter 5, “Architecting for Performance: Design,” which consists of Disk, CPU, Memory, and Network. Our experience has shown us, and data from VMware Support validates this belief, that more than 80% of performance problems in database environments, and especially virtualized environments, are directly related to storage. Understanding the storage architecture in a virtualized environment and getting your storage architecture right will have a major impact on your database performance and the success of your SQL Server virtualization project. Bear in mind as you work through your storage architecture and this chapter that virtualization is bound by the laws of physics—it won’t fix bad code or bad database queries. However, if you have bad code and bad queries, we will make them run as fast as possible.

TIP

Greater than 80% of all problems in a virtualized environment are caused by the storage in some way, shape, or form.

This chapter first covers the key aspects of storage architecture relevant to both physical and virtual environments as well as the differences you need to understand when architecting storage, specifically for virtualized SQL Server Databases. Many of the concepts we discuss will be valid for past versions of SQL Server and even the newest release, SQL Server 2014.

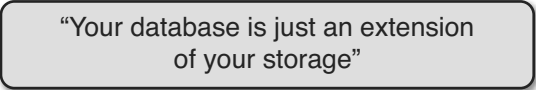
We provide guidance on what our experience has taught us are important database storage design principles. We present a top-down approach covering SQL Server Database and Guest OS Design, Virtual Machine Template Design, followed by VMware vSphere Hypervisor Storage Design and then down to the physical storage layers, including using server-side flash acceleration technology to increase performance and provide greater return on investment. We conclude the chapter by covering one of the biggest IT trends and its impact on SQL Server. Throughout this chapter, we give you architecture examples based on real-world projects that you can adapt for your purposes.

When designing your storage architecture for SQL Server, you need to clearly understand the requirements and have quantitative rather than subjective metrics. Our experience has taught us to make decisions based on fact and not gut feeling. You will need to benchmark and baseline your storage performance to clearly understand what is achievable from your design. Benchmarking and baselining performance are critical to your success, so we've dedicated an entire chapter (Chapter 10, "How to Baseline Your Physical SQL Server System") to those topics. In this chapter, we discuss some of the important storage system component performance aspects that will feed into your benchmarking and baselining activities.

The Five Key Principles of Database Storage Design

When architecting storage for SQL Server, it's important to understand a few important principles. These will help guide your design decisions and help you achieve acceptable performance both now and in the future. These principles are important because over the past decade, CPU performance has increased at a much faster pace than storage performance, even while capacity has exploded.

Principle 1: Your database is just an extension of your storage



"Your database is just an extension
of your storage"

Figure 6.1 Quote from Michael Webster, VMworld 2012

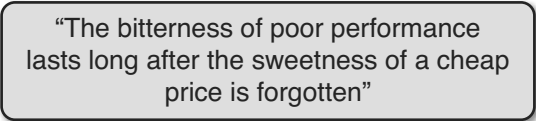
The first principle is highlighted in Figure 6.1: that your database is just an extension of your storage. A database is designed to efficiently and quickly organize, retrieve, and process large quantities of data to and from storage. So increasing the parallelism of access

to storage resources at low latency will be an important goal. Later in this chapter, we cover how to optimize the architecture of your database to maximize its storage performance and parallelism. When you understand this principle, it's easy to understand why getting your storage design and performance is so critical to the success of your SQL Server Database virtualization project.

Principle 2: Performance is more than underlying storage devices

The next key principle is that storage performance is more than just about underlying storage devices and spindles, although they are very important too. SQL Server storage performance is multidimensional and is tightly coupled with a number of different system components, such as the number of data files allocated to the database, the number of allocated vCPUs, and the amount of memory allocated to the database. This is why we like to use the term “IT Food Groups,” because it is so important to feed your database the right balance of these critical resources. This interplay between resources such as CPU, Memory, and Network and their impact on storage architecture and performance will be covered in subsequent sections of this chapter.

Principle 3: Size for performance before capacity



“The bitterness of poor performance lasts long after the sweetness of a cheap price is forgotten”

Figure 6.2 Quote from Michael Webster, VMworld 2013

Figure 6.2 is loosely based on the eighteenth-century quote “The bitterness of poor quality remains long after the sweetness of low price is forgotten,” by Benjamin Franklin. Both quotes are extremely relevant to SQL Server database and storage performance.

This brings us to the next key principle. In order to prevent poor performance from being a factor in your SQL Server virtualization project (refer to Figure 6.2), you should design storage for performance first (IOPS and latency), then capacity will take care of itself. Capacity is the easy part. We will show you later in this chapter how compromising on certain storage configurations on the surface can actually cost you a lot more by causing unusable capacity due to poor performance.

CAUTION

A lesson from the field: We were working with a customer, and they wanted to design and run a database on vSphere that could support sustained 20,000 IOPS. After we worked with the customer's vSphere, SAN, Network, and DBA teams, the customer decided to move forward with the project. The customer then called in a panic saying, "In our load test, we achieved 1,000 IOPS. We are 19,000 short of where we need to be." Trust me, this is a phone call you don't want to get. Playing the odds, we started with the disk subsystem. We quickly identified some issues. The main issue was the customer purchased for *capacity*, not *performance*. They had to reorder the right disk. Once the new (right) disk arrived and was configured, the customer exceeded the 20,000 IOPS requirement.

TIP

When it comes to storage devices, HDDs are cents per GB but dollars per IOP, whereas SSDs are cents per IOP and dollars per GB. SSDs should be considered cheap memory, rather than expensive disks, especially when it comes to enterprise SSDs and PCIe flash devices.

Principle 4: Virtualize, but without compromise

The next principle is that virtualizing business-critical SQL Server databases is all about reducing risk and not compromising on SLAs. Virtualize, but without compromise. There is no need to compromise on predictability of performance, quality of service, availability, manageability, or response times. Your storage architecture plays a big part in ensuring your SQL databases will perform as expected. As we said earlier, your database is just an extension of your storage. We will show you how to optimize your storage design for manageability without compromising its performance.


Believe it or not, as big of advocates as we are about virtualizing SQL Server, we have told customers in meetings that now is not the right time for this database to be virtualized. This has nothing to do with the capability of vSphere or virtualization, but more to do with the ability of the organization to properly operate critical SQL systems and virtualize them successfully, or because they are not able or willing to invest appropriately to make the project a success. If you aren't willing to take a methodical and careful approach to virtualization projects for business-critical applications, in a way that increases the chances of success, then it's not worth doing. Understand, document, and ensure requirements can

be met through good design and followed by testing and validation. It is worth doing, and it is worth “Doing It Right!”

Principle 5: Keep it standardized and simple (KISS)

This brings us to the final principle. Having a standardized and simplified design will allow your environment and databases to be more manageable as the numbers scale while maintaining acceptable performance (see Principle 4). If you have a small number of standardized templates that fit the majority of your database requirements and follow a building-block approach, this is very easy to scale and easy for your database administrators to manage. We’ll use the KISS principle (Keep It Standardized and Simple) throughout this chapter, even as we dive into the details. Once you’ve made a design decision, you should standardize on that decision across all your VM templates. Then when you build from those templates, you’ll know that the settings will always be applied.

SQL Server Database and Guest OS Storage Design



The starting point for any storage architecture for SQL Server Databases is actually with our last design principle: KISS (Keep It Standardized and Simple). But all of the principles apply. We will determine the smallest number of templates that are required to virtualize the majority (95%) of database systems, and anything that falls outside this will be handled as an exception.

Your first step is to analyze the inventory of the SQL Server Databases that will be virtualized as part of your project (refer to Chapter 4, “Virtualizing SQL Server 2012: Doing It Right”). From this inventory, you will now put each database and server into a group with similar-sized databases that have similar requirements. The storage requirements for all of these existing and new databases, based on their grouping, will be used to define the storage layouts and architecture for each of the SQL Server Databases, Guest OS, and VM template.

TIP

If you are virtualizing existing databases, you might consider using a tool such as VMware Capacity Planner, VMware Application Dependency Planner, Microsoft System Center, or Microsoft Assessment and Planning Toolkit to produce the inventory. VMware Capacity Planner and Application Dependency Planner are available from VMware Professional Services or your preferred VMware partner. When you're baselining a SQL Server database, a lot can happen in a minute. We recommend your sample period for CPU, Memory, and Disk be 15 seconds or less. We recommend you sample T-SQL every minute.

SQL Server Database File Layout

Database file layout provides an important component of database storage performance. If you have existing databases that will be virtualized, you or your DBAs will likely have already developed some practices around the number of database files, the size of database files, and the database file layout on the file system. If you don't have these practices already in place, here we provide you with some guidelines to start with that have proven successful.

Your SQL Server database has three primary types of files you need to consider when architecting your storage to ensure optimal performance: data files, transaction log files, and Temp DB files. Temp DB is a special system database used in certain key operations, and has a big performance impact on your overall system. The file extensions you'll see are .mdf (master data file), .ndf (for secondary data files), and .ldf for transaction log files. We will go over all of these different file types later in this chapter.

Number of Database Files

First, we need to determine the number of database files. There are two main drivers for the number of files you will specify. The first driver is the number of vCPUs allocated to the database, and the second is the total capacity required for the database now and in the future.

Two design principles come into play here: The parallelism of access to storage should be maximized by having multiple database files, and storage performance is more than just the underlying devices. In the case of data files and Temp DB files, they are related to the number of CPU cores allocated to your database. Table 6.1 provides recommendations from Microsoft and the authors in relation to file type.

NOTE

It is extremely unlikely you will ever reach the maximum storage capacity limits of a SQL Server 2012 database system. We will not be covering the maximums here. We recommend you refer to Microsoft (<http://technet.microsoft.com/en-us/library/ms143432.aspx>).

Table 6.1 Number of Data Files and Temp DB Files Per CPU

File Type	Microsoft Recommended Setting	Author Recommended Setting
Temp DB Data File	1 per CPU core	< 8 vCPU, 1 per vCPU > 8 vCPU, 8 total (increase number of files in increments of four at a time if required) Max 32
Database Data File	0.25 to 1.0 per file group, per CPU core	Min 1 per vCPU, max 32
Database Transaction Log File	1	1*
Temp DB Transaction Log File	1	1*

*If Temp DB and Transaction Log are deployed on local SSD or flash storage, especially when using AlwaysOn Availability Groups, then it is recommended to have an additional copy on SAN.

Microsoft recommends as a best practice that you should configure one Temp DB data file per CPU core and 0.25 to 1 data file (per file group) per CPU core. Based on our experience, our recommendation is slightly different.

If your database is allocated eight or fewer vCPUs as a starting point, we recommend you should configure at least one Temp DB file per vCPU. If your database is allocated more than eight vCPUs, we recommend you start with eight Temp DB files and increase by lots of four in the case of performance bottlenecks or capacity dictates.

TIP

Temp DB is very important because it's extensively utilized by OLTP databases during index reorg operations, sorts, and joins, as well as for OLAP, DSS, and batch operations, which often include large sorts and join activity.

We recommend in all cases you configure at least one data file (per file group) per vCPU. We recommend a maximum of 32 files for Temp DB or per file group for database files because you'll start to see diminishing performance returns with large numbers of database files over and above 16 files. Insufficient number of data files can lead to many writer processes queuing to update GAM pages. This is known as GAM page contention. The Global Allocation Map (GAM) tracks which extents have been allocated in each file. GAM contention would manifest in high PageLatch wait times. For extremely large databases into the many tens of TB, 32 files of each type should be sufficient.

Updates to GAM pages must be serialized to preserve consistency; therefore, the optimal way to scale and avoid GAM page contention is to design sufficient data files and ensure all data files are the same size and have the same amount of data. This ensures that GAM page updates are equally balanced across data files. Generally, 16 data files for tempdb and user databases is sufficient. For Very Large Database (VLDB) scenarios, up to 32 can be considered. See <http://blogs.msdn.com/b/sqlserverstorageengine/archive/2009/01/04/what-is-allocation-bottleneck.aspx>.

If you expect your database to grow significantly long term, we would recommend that you consider configuring more data files up front. The reason we specify at least one file per CPU is to increase the parallelism of access from CPU to data files, which will reduce any unnecessary data access bottlenecks and lower latency. This also allows for even data growth, which will reduce IO hotspots.

CAUTION

Having too few or too many Temp DB files can impact the overall performance of your database. Our guidance is conservative and aimed to meet the requirements for the majority of SQL systems. If you start to see performance problems such as higher than normal query response times or excessive database waits in *PAGELATCH_XX*, then you have contention in memory and may need to increase the number of Temp DB files further and/or implement trace flag 1118 (which we recommend), which prevents single page allocations. If you see waits in *PAGEIOLATCH_XX*, then the contention is at the IO subsystem level. Refer to <http://www.sqlskills.com/blogs/paul/a-sql-server-dba-myth-a-day-1230-Temp-DB-should-always-have-one-data-file-per-processor-core/> and Microsoft KB 328551 (<http://support.microsoft.com/kb/328551>).

TIP

The number of data files and Temp DB files is important enough that Microsoft has two spots in the Top 10 SQL Server Storage best practices highlighting the number of data files per CPU. Refer to <http://technet.microsoft.com/en-us/library/cc966534.aspx>.

NOTE

When you're determining the number of database files, a vCPU is logically analogous to a CPU core in a native physical deployment. However, in a native physical environment without virtualization, each CPU core may also have a hyper-thread. In a virtual environment, each vCPU is a single thread. There is no virtual equivalent of a hyper-thread.

Figure 6.3 shows an example of data files, Temp DB files, and transaction log files allocated to a SQL Server 2012 Database on a sample system with four vCPU and 32GB RAM.

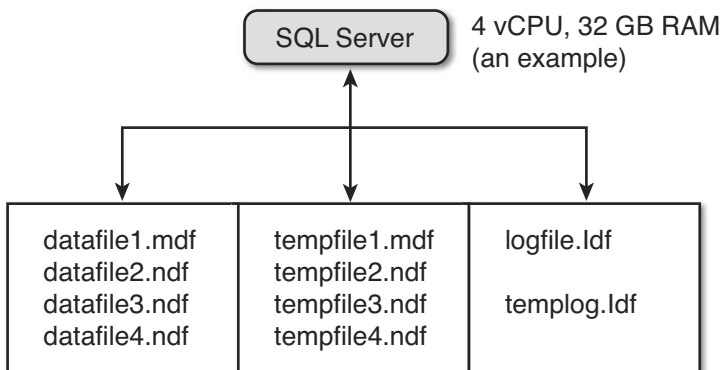


Figure 6.3 SQL Database data file allocation.

NOTE

As Figure 6.3 illustrates, there is only one transaction log file per database and per Temp DB. Log files are written to sequentially, so there is no benefit in having multiples of them, unless you exceed the maximum log file size (2TB) between backups. There is a benefit of having them on very fast and reliable storage, which will be covered later.

This page intentionally left blank

Index

Numbers

10Gb Ethernet NICs, 269

A

ABRTS/s counter, 326

ACID (Atomicity, Consistency,
Isolation, and Durability), 302-303

ActivePerl, 407

ActiveState ActivePerl, 407

adapter count, 95

adapters

- CAN (Converged Network Adapter),
276

- iSCSI, 276

- LSI Logic SAS, 95

- physical network adapters, 267-269

- PVSCSI, 95

- virtual network adapters, 100

 - choosing, 250-251

 - traffic types, 101-102

 - tuning, 252-254

- addresses (IP), 341-342

- Admission Control, 88

- affinity rules, 358

- AGs (Availability Groups), 306-308

- alignment of partitions, 128-129

- AlwaysOn Availability Groups

 - configuring, 387-391

 - creating, 399-405

- AlwaysOn Failover Cluster Instance,
125

- anti-affinity rules, 358

- Application Dependency Planner, 110

- AQLEN, 168

- arrays, 98-99

atomicity, 302-303

ATS (Atomic Test Set), 99

Auto Grow, 114-115

availability, 135

- ACID (Atomicity, Consistency, Isolation, and Durability), 302-303
- business continuity, 291
- determining availability requirements, 287-288
- disaster recovery, 291-294
- high availability, 14-16
- providing a menu of options, 288-289
- RPOs (recovery point objectives), 290
- RTOs (recovery time objectives), 290
- sample high availability chart, 308-309
- SLAs (service-level agreements), 290
- SQL Server AlwaysOn Failover Cluster Instance (FCI), 304-306
- SQL Server Availability Groups (AGs), 306-308

vSphere high availability

- DRS (Distributed Resource Scheduler), 297
- hypervisor availability features, 294-296
- Storage DRS, 297
- Storage vMotion, 297
- vCenter SRM (Site Recovery Manager), 301
- vCHS (vCloud Hybrid Service), 302
- vDP (vSphere Data Protection), 300

- vMotion, 296-297
- vSphere App HA, 299-300
- vSphere HA, 298-299
- vSphere Replication, 300-301
- X-vMotion, 298

Availability Groups (AGs), 306-308

Available Mbytes metrics, 321

Average Latch Wait Time(ms) metric, 324

Average Wait Time(ms) metric, 324

B

background noise, lack of, 334

backing up networks, 103

ballooning, 230-232

bandwidth, vMotion traffic, 276

baselines

- baseline performance reports, 332-333
- benchmarks, 315-316
 - developing, 317-318
 - industry-standard benchmarks, 316
 - validating performance with, 318
 - vendor benchmarks, 316-317
- common performance traps
 - blended peaks of multiple systems, 335
 - failure to consider SCU (Single Compute Unit) performance, 335
 - invalid assumptions, 334
 - lack of background noise, 334

- shared core infrastructure
 - between production and non-production, 333-334
 - vMotion slot sizes of monster database virtual machines, 336-337
 - comparing
 - different processor generations, 330-331
 - different processor types, 328-330
 - customer deployments, 71
 - database workload, 48-50
 - explained, 311-314
 - metrics
 - ESXTOP counters, 325-327
 - SQL Server baseline infrastructure metrics, 321-322
 - SQL Server Perfmon counters, 323-324
 - SQL Server Profiler counters, 324-325
 - non-production workload influences on performance, 331-332
 - reasons for, 319-320
 - validating performance with, 318
 - vSphere infrastructure, 46-48
 - when to record, 320
- Batch Requests/sec metric, 324
- Batch/ETL (Extract Transform Load) workloads, 64
- benchmarks, 315-316
- developing
 - benchmark model based on recorded production performance, 318
 - benchmark model based on system nonfunctional requirements, 317
 - industry-standard benchmarks, 316
 - validating performance with, 318
 - vendor benchmarks, 316-317
- BIOS settings, 12-13
- blended peaks of multiple systems, 335
- blocks, Pointer Block Eviction Process, 163-164
- blogs, 444
 - Thomas LaRock's blog, 445
 - vLaunchPad, 444-445
- breaking down large pages, 238-239
- buffer
 - Buffer Cache, 49
 - Buffer Cache Hit Ratio, 50, 323
 - Buffer Manager, 323
 - Buffer Pool, 129-130, 219-220
- built-in in-memory, 246-247
- business case for virtualization, 9
 - BIOS settings, 12-13
 - DBA (database administrator) advantages, 10-11
 - hardware refresh, 20-22
 - high availability, 14-16
 - large databases, 22-23
 - performance, 16-17
 - provisioning/DBaaS, 17-20
 - database tiering, 19-20
 - shared environments, 20
 - reduced expenses, 9-10
 - SLAs (service level agreements), 11-12

business continuity, 291
 business transparency, 73
 Bytes Total/sec metric, 322

C

cache

- Buffer Cache, 49
- Buffer Cache Hit Ratio, 50, 323
- CACHEUSED counter, 326
- Fusion-io ioTurbine, 201-203
- vRFC (vSphere Flash Read Cache), 199-201

CACHEUSED counter, 326

CAN (Converged Network Adapter), 276

capacity, one-to-one relationships and unused capacity, 38-40

Center of Excellence (CoE), 61-63

charge back, 73

“check it before you wreck it” rule, 36

choosing virtual network adapters, 250-251

cloud, vCHS (vCloud Hybrid Service), 302

Cluster Validation Wizard, 363-364

clusters

- failover cluster instance storage layout, 157

- vSphere 5.5 failover clustering environments, 185-186

- Windows Failover Clustering

- configuring, 359-368

- quorum mode, 369-374

- validating, 368

- WSFC (Windows Server Failover Clustering), 304

CMDS/s counter, 326

CoE (Center of Excellence), 61-63

Column Storage, 134-135

commands, SP_Configure, 246

communication, 58-59

- communication responsiveness, 253

- mutual understanding, 59-60

- responsibility domains, 60-61

comparing performance baselines, 328

- different processor generations, 330-331

- different processor types, 328-330

Complete page (SQL Server installation), 387

compression, 133

compromise, virtualization without, 108-109

computer names, requirements for performance testing, 341-342

configuration

- AlwaysOn Availability Groups, 387-391

- Hot-Add Memory and Hot-Add CPU, 356-358

- jumbo frames, 259-262, 393-394

- max/min memory, 392

- Max Server Memory, 236-237

- performance test labs, 342

- performance tests, 339-340

- affinity and anti-affinity rules, 358

- AlwaysOn Availability Groups configuration, 387-391

AlwaysOn Availability Groups
 creation, 399-405
computer name/IP address
 requirements, 341-342
Dell DVD Store installation and
 configuration, 406-430
Dell DVD Store load test,
 430-436
Hot-Add Memory and Hot-Add
 CPU, 356-358
jumbo frame configuration,
 393-394
max/min memory configuration,
 392
memory reservations, 355
multiple tempdb files, 394-395
network connection validation,
 359
performance test lab setup,
 342-345
software requirements, 341
SQL Server 2012 installation,
 374-387
test database creation, 396-398
VMDK file configuration,
 345-354
Windows Failover Clustering
 configuration, 359-368
Windows Failover Clustering
 quorum mode, 369-374
Windows Failover Clustering
 validation, 368
trace flags, 215
Configure Cluster Quorum Wizard,
 370-374

consistency, 302-303
consolidations, 53, 68
continuity (business), 291
controllers, virtual storage, 138-143
Converged Network Adapter (CAN),
 276
Cores per Socket, 83
counters
 ESXTOP, 325-327
 Perfmon, 50, 323-324
 Profiler, 324-325
CPUs, 74-76
 CPU Scheduler, 86
 hot-add CPUs, 3-4
CPU Scheduler, 86
Create Cluster Wizard, 366
CrossSubnetDelay, 255
CrossSubnetThreshold, 255
%CSTP counter, 326
customer deployment baselines, 71

D

DaaS (Database as a Service), 73
 and database virtualization, 17-20
Darwin, Charles, 1-3
database administrators. *See* DBAs
Database as a Service. *See* DaaS
database availability, 287
 ACID (Atomicity, Consistency,
 Isolation, and Durability), 302-303
 business continuity, 291
 determining availability
 requirements, 287-288

- disaster recovery, 291-294
- providing a menu of options, 288-289
- RPOs (recovery point objectives), 290
- RTOs (recovery time objectives), 290
- sample high availability chart, 308-309
- SLAs (service-level agreements), 290
- SQL Server AlwaysOn Availability Groups (AGs), 306-308
- SQL Server AlwaysOn Failover Cluster Instance (FCI), 304-306
- vSphere high availability
 - DRS (Distributed Resource Scheduler), 297
 - hypervisor availability features, 294-296
 - Storage DRS, 297
 - Storage vMotion, 297
 - vCenter SRM (Site Recovery Manager), 301
 - vCHS (vCloud Hybrid Service), 302
 - vDP (vSphere Data Protection), 300
 - vMotion, 296-297
 - vSphere App HA, 299-300
 - vSphere HA, 298-299
 - vSphere Replication, 300-301
 - X-vMotion, 298
- database availability design, 135
- database buffer pool, 219-220
- database consolidations, 53
- Database Engine Configuration page (SQL Server installation), 382
- database files
 - file system layout, 110-122
 - data files, 123-126
 - log files, 123-126
 - NTFS file system allocation unit size, 126-127
 - OS, application binaries, and page file, 122
 - partition alignment, 128-129
 - Temp DB files, 123-126
 - Instant File Initialization (IFI), 120-122
 - number of, 110-113
 - size of, 114-116
 - data files, 116
 - Temp DB files, 116
 - transaction log file sizing, 117-120
- database indexes, 222-225
- database installation guidelines, 32-36
- database instances, number of, 244-245
- database metrics, 324
- database pages
 - explained, 219-220
 - large pages
 - breaking down into default page size, 238-239
 - explained, 237-238
 - locking pages in memory, 239-241
 - paging, 220-221
 - swapping, 220-221
 - TPS (transparent page sharing), 228-229

- database statistics, updating, 130-132
 - with maintenance plan, 131-132
 - with trace flag 2371, 131
- database storage. *See* storage
- database tiering, 19-20
- database virtualization. *See* virtualization
- database workload, baselining, 48-50
- Data Compression, 133
- data files
 - file system layout, 123-126
 - sizing, 116
- data protection (vDP), 300
- data stores
 - number of, 165-169
 - virtual disks per data store, 170-173
- DAVG/cmd counter, 326
- DAVG/cmd metric, 48
- DBA Guide to Databases on VMware (white paper), 439
- DBAs (database administrators)
 - business case for virtualization, 10-11
 - BIOS settings, 12-13
 - hardware refresh, 20-22
 - high availability, 14-16
 - large databases, 22-23
 - performance, 16-17
 - provisioning/DBaaS, 17-20
 - SLAs (service level agreements), 11-12
 - SLAs (service level agreements), 11-12
- Decision Support System. *See* DSS workload
- default queue depth (QLogic HBA), 166
- Dell DVD Store, 327
 - installing and configuring, 406-430
 - load test, 430-436
- Dell DVD Store Custom Install Wizard, 408
- Denneman, Frank, 444
- deployment, 54, 63
- design
 - deployment, 54, 63
 - networks. *See* network design
 - server-side flash acceleration, 198-199
 - Fusion-io ioTurbine, 201-203
 - PernixData FVP, 204-206
 - vSphere Flash Read Cache (vFRC), 199-201
- SQL Server database and guest OS storage, 109
 - Buffer Pool, 129-130
 - Column Storage, 134-135
 - database availability, 135
 - database statistics, updating, 130-132
 - Data Compression, 133
 - file system layout, 110, 122-129
 - Instant File Initialization (IFI), 120-122
 - number of database files, 110-113
 - size of database files, 114-120
 - Storage Spaces, 136
 - Volume Managers, 136
- SQL Server on hyperconverged infrastructure, 207-213

- SQL Server virtual machine storage, 136
 - expanding, 158-159
 - Jumbo VMDKs, 159-164
 - layout, 152-157
 - virtual disk devices, 143-152
 - virtual storage controllers, 138-143
 - VM hardware version, 137
- storage design principles
 - database as extension of storage, 106
 - KISS principle (Keep It Standardized and Simple), 109
 - performance and underlying storage devices, 107
 - sizing for performance before capacity, 107-108
 - virtualization without compromise, 108-109
- vSphere storage, 164
 - multipathing, 184-185
 - number of data stores and data store queues, 165-169
 - number of virtual disks per data store, 170-173
 - RAID (Redundant Array of Independent Disks), 187-197
 - storage DRS, 177-183
 - Storage IO Control (SIOC), 173-177
 - storage policies, 177-183
 - vSphere 5.5 failover clustering environments, 185-186
 - determining availability requirements, 287-288
- developing benchmarks
 - benchmark model based on recorded production performance, 318
 - benchmark model based on system nonfunctional requirements, 317
- disaster recovery, 291-294
- discovery, 53
- Disk Management utility, 352
- Disk Space Requirements page (SQL Server installation), 382
- disks
 - disk layout, 95
 - Disk Management utility, 352
 - enterprise flash disks (EFDs), 195-197
 - RAID (Redundant Array of Independent Disks)
 - economics of RAID performance, 194-197
 - IO penalties, 189-194
 - randomness of IO pattern, 187-188
 - read/write bias, 188
 - virtual disk devices, 143
 - IO blender effect, 151-152
 - Raw Device Map (RDM), 149-151
 - Thick Eager Zero disks, 147-148
 - Thin versus Thick Lazy Zero disks, 144-146
- Distributed Resource Scheduler (DRS), 297
- distributed switches, 263
- distributed virtual switches, 100

documentation
 online resources, 437-440
 reading, 43-44

dozing, 13

DQLEN, 168

DRaaS, 293-294

drivers (PVSCSI), 31

%DRPPX counter, 327

%DRPTX counter, 327

DRS (Distributed Resource Scheduler), 297

DSS (Decision Support System)
 workload, 64

durability, 302-303

dynamic threshold for automatic
 statistics update, 131

E

E1000 adapters, 250-252

E1000E adapters, 250-252

economics of RAID performance,
 194-197

education, 60

EFDs (enterprise flash disks), 195-197

Effects of Min and Max Server Memory
 (article), 235

enabling. *See* configuration

encapsulation, 28

enterprise flash disks (EFDs), 195-197

Epping, Duncan, 444

Error Reporting page (SQL Server
 installation), 385

ESXi host swap file location, 78

ESXTOP counters, 325-327

ETL (Extract Transform Load), 64

expanding SQL virtual machine storage
 layout, 158-159

expenses, reducing, 9-10

F

Facebook groups, 443

Failover Cluster Instance (FCI), 98,
 304-306

Failover Cluster Manager, 362

Failover Clustering
 configuring, 359-368

Failover Cluster Instance (FCI), 98,
 304-306

Failover Cluster Manager, 362

failover clustering environments,
 185-186

network settings, 254-256

network teaming, 270-273

quorum mode, 369-374

storage layout, 157

validating, 368

FCI (Failover Cluster Instance), 98,
 304-306

Feature Selection page (SQL Server
 installation), 379

file system layout, 110-122
 data files, 123-126
 log files, 123-126
 NTFS file system allocation unit size,
 126-127

- OS, application binaries, and page file, 122
- partition alignment, 128-129
- Temp DB files, 123-126

files

- database files
 - Instant File Initialization (IFI), 120-122
 - number of, 110-113
 - size of, 114-120
- file system layout, 110-122
 - data files, 123-126
 - log files, 123-126
 - NTFS file system allocation unit size, 126-127
 - OS, application binaries, and page file, 122
 - partition alignment, 128-129
 - Temp DB files, 123-126
- multiple tempdb files, 394-395
- VLFs (Virtual Log Files), 118-120
- VMDK files, configuring
 - inside guest operating system, 352-354
 - on virtual machines, 345-351
- vswap, memory reservations and, 233-234
- vswp files, 88

flash, server-side flash acceleration, 198-199

- Fusion-io ioTurbine, 201-203
- PernixData FVP, 204-206
- vSphere Flash Read Cache (vFRC), 199-201

Flash Virtualization Platform (FVP), 204-206

frames

- jumbo frames, 256-259
 - configuring, 259-262, 393-394
 - testing, 262-264
- pause frames, 268

Free System Page Table Entries metric, 321

full virtualization, importance of, 36-38

Fusion-io ioTurbine, 201-203

Fusion-io ioTurbine Profiler, 203

FVP (Flash Virtualization Platform), 204-206

G

Gage, John, 281

GAVG/cmd for NFS Datastores counter, 326

General Statistics, 324

GPT (GUID Partition Table), 136

Gray, Jim, 302-303

groups

- AlwaysOn Availability Groups
 - configuring, 387-391
 - creating, 399-405

- database groups, 69

- user groups

- Facebook groups, 443

- PASS (Professional Association of SQL Server), 441-442

VMUG (VMware Users Group), 440

VMWare Community, 442-443

guest OS storage, 27, 109

 Buffer Pool, 129-130

 Column Storage, 134-135

 database availability, 135

 database statistics, updating, 130-132

 with maintenance plan, 131-132

 with trace flag 2371, 131

Data Compression, 133

file system layout, 110, 122

 data files, 123-126

 log files, 123-126

 NTFS file system allocation unit size, 126-127

 OS, application binaries, and page file, 122

 partition alignment, 128-129

 Temp DB files, 123-126

Instant File Initialization (IFI), 120-122

number of database files, 110-112

size of database files, 114-116

 data file sizing, 116

 Temp DB file sizing, 116

 transaction log file sizing, 117-120

Storage Spaces, 136

VMDK file configuration in, 352-354

Volume Managers, 136

GUID Partition Table (GPT), 136

H

HA (high availability)

 ACID (Atomicity, Consistency, Isolation, and Durability), 302-303

 DRS (Distributed Resource Scheduler), 297

 hypervisor availability features, 294-296

 sample high availability chart, 308-309

 SQL Server AlwaysOn Availability Groups (AGs), 306-308

 SQL Server AlwaysOn Failover Cluster Instance (FCI), 304-306

 Storage DRS, 297

 Storage vMotion, 297

 vCenter SRM (Site Recovery Manager), 301

 vCHS (vCloud Hybrid Service), 302

 vDP (vSphere Data Protection), 300

 vMotion, 296-297

 vSphere App HA, 299-300

 vSphere HA, 298-299

 vSphere Replication, 300-301

 X-vMotion, 298

hardware. *See* physical hardware

hardware independence, 28

hardware refresh and database virtualization, 20-22

Heap size, 160-162

heartbeat vNICs, 256

help. *See* resources

high availability. *See* HA (high availability)

high-level virtualization implementation plan, 50-51

phase 1: requirements gathering, 51-52

phase 2: discovery, 53

phase 2.1: database consolidations, 53

phase 3: infrastructure adjustments, 53

phase 4: validation and testing, 54

phase 5: migration and deployment, 54

phase 6: monitoring and management, 54

Hirt, Allan, 445

Hogan, Cormac, 445

host-local swap, 78

host memory, 225-226

host <servername> CPU% value, 434

Hot-Add CPU, 3-4, 356-358

Hot-Add Memory, 4-5, 356-358

HTT (Hyper-Threading Technology), 85-87

hyperconverged infrastructure, 207-213, 280

Hyper-Threading Technology (HTT), 85-87

Hyperic, 343-345

hypervisor, 25

availability features, 294-296

compared to OS, 26-27

explained, 25-27

importance of full virtualization, 36-38

one-to-many relationships, 40

one-to-one relationships and unused capacity, 38-40

paravirtualization, 29

PVSCSI (paravirtual SCSI driver), 31

Type-1 hypervisors, 30

Type-2 hypervisors, 31

virtualized database installation guidelines, 32-36

VMs (virtual machines), 28

VMware ESXi versions, 40-41

VMXNET3, 32

I

IFI (Instant File Initialization), 120-122

ILM (information life cycle management), 207

implementation plans

database workload baselines, 48-50

high-level plan, 50-51

items to consider, 44-45

phase 1: requirements gathering, 51-52

phase 2: discovery, 53

phase 2.1: database consolidations, 53

phase 3: infrastructure adjustments, 53

phase 4: validation and testing, 54

phase 5: migration and deployment, 54

phase 6: monitoring and management, 54

- RPOs (recovery point objectives), 45-46
 - RTOs (recovery time objectives), 45-46
 - SLAs (service-level agreements), 45-46
 - vSphere infrastructure baselines, 46-48
 - Independent Persistent (SQL FCI), 157
 - indexes (database), 222-225
 - industry-standard benchmarks, 316
 - information life cycle management (ILM), 207
 - Infrastructure Navigator, 343-344
 - initialization, Instant File Initialization (IFI), 120-122
 - installation. *See also* configuration
 - Dell DVD Store, 406-430
 - SQL Server 2012, 374-377, 384-387
 - Complete page, 387
 - Database Engine Configuration page, 382
 - Disk Space Requirements page, 382
 - Feature Selection page, 379
 - Installation Configuration Rules page, 385
 - Installation Rules page, 380
 - Instance Configuration page, 380
 - License Terms page, 377
 - preflight check, 375
 - Product Key page, 375
 - Ready to Install page, 385
 - Server Configuration page, 382
 - Setup Role page, 379
 - virtualized database installation guidelines, 32-36
 - Installation Configuration Rules page (SQL Server installation), 385
 - Installation Rules page (SQL Server installation), 380
 - Instance Configuration page (SQL Server installation), 380
 - Instant File Initialization (IFI), 120-122
 - invalid assumptions, 334
 - IO blender effect, 151-152
 - IOBlazer, 327
 - IOMeter, 142-143, 327
 - ioTurbine (Fusion-io), 201-203
 - ioTurbine Profiler (Fusion-io), 203
 - IP addresses, requirements for performance testing, 342
 - iSCSI
 - adapters, 276
 - port binding, 281
 - isolation, 28, 302-303
-
- ## J
- jumbo frames, 256-259
 - configuring, 259-262, 393-394
 - testing, 262-264
 - Jumbo VMDKs, 159
 - Pointer Block Eviction Process, 163-164
 - VMFS Heap size considerations, 160-162

K

KAVG/cmd counter, 326
Keep It Standardized and Simple (KISS), 109
KISS principle (Keep It Standardized and Simple), 109
Klee, David, 445

L

LACP, 273
Lam, William, 444
large databases and database virtualization, 22-23
large pages, 79

- breaking down into default page size, 238-239
- explained, 237-238
- locking pages in memory, 239-241

LaRock, Thomas, 445
latches, 324
layout

- file system layout, 110
 - data files, 123-126
 - log files, 123-126
- NTFS file system allocation unit size, 126-127
- OS, application binaries, and page file, 122
- partition alignment, 128-129
- Temp DB files, 123-126

virtual machine storage layout, 152-157

Leaf-Spine network architecture, 273
licenses

- License Terms page (SQL Server installation), 377
- VMware vCloud Suite licenses, 285

load test (Dell DVD Store), 430-436
locking pages in memory, 92, 239-241
locks, 324
log files

- file system layout, 123-126
- sizing, 117-120

Log Flush Wait Time, 324
Log Flush Waits/sec, 324
LogicalDisk(*): Avg Disk Sec/Read, 322
LogicalDisk(*): Avg. Disk Sec/Write, 322
LogicalDisk Disk Bytes/sec, 321
Logins/sec, 324
Logout/sec, 324
Lowe, Scott, 445
LSI Logic SAS, 95, 137-142
LUN queue depth, 167

M

maintenance plans, updating database statistics with, 131-132
Maintenance Plan Wizard, 420-422
management, 54
Max Server Memory, 234-236
MaxAddressableSpaceTB, 163-164
maximum memory, configuring, 392
maximum storage capacity limits, 111
MbRX/s counter, 327

- MbTx/s counter, 327
- MCTLSZ (MB) counter, 47-48, 326
- memory
 - Buffer Pool, 129-130
 - cache
 - Buffer Cache, 49
 - Buffer Cache Hit Ratio, 50, 323
 - CACHEUSED counter, 326
 - Fusion-io ioTurbine, 201-203
 - vRFC (vSphere Flash Read Cache), 199-201
 - host memory, 225-226
 - hot-add memory, 4-5
 - large pages
 - breaking down into default page size, 238-239
 - explained, 237-238
 - locking pages in memory, 239-241
 - memory ballooning, 230-232
 - Memory Grants Pending, 324
 - Memory Manager, 324
 - memory overcommitment, 87
 - memory reservation, 355
 - explained, 232-233
 - mixed workload environment with memory reservations, 226-228
 - VMware HA strict admission control, 233
 - vswap file, 233-234
 - memory trends and the stack, 218
 - database buffer pool, 219-220
 - database indexes, 222-225
 - database pages, 219-221
 - paging, 220-221
 - swapping, 220-221
 - min/max memory, configuring, 392
 - mixed workload environment with memory reservations, 226-228
 - NUMA (Non-uniform Memory Access)
 - explained, 241-243
 - vNUMA, 243
 - overview, 76, 217-218
 - RAM, 87
 - shared-resource world, 246
 - SQL Server 2014 in-memory, 246-247
 - SQL Server Max Server Memory, 234-236
 - SQL Server Min Server Memory, 235
 - TPS (transparent page sharing), 228
 - VMs (virtual machines), 225-226
 - number of, 244-245
 - sizing, 244
 - vRFC (vSphere Flash Read Cache), 199-201
 - xVelocity memory, 134
- Memory Grants Pending, 324
- Memory Manager, 324
- metrics
 - baseline metrics
 - ESXTOP counters, 325-327
 - SQL Server baseline infrastructure metrics, 321-322
 - SQL Server Perfmon counters, 323-324
 - SQL Server Profiler counters, 324-325

- Buffer Cache Hit Ratio, 50
- Cache Hit Ratio, 50
- DAVG/cmd, 48
- MCTLSZ, 47-48
- %MLMTD, 47-48
- %RDY, 47
- READs/s, 47-48
- storage-specific metrics, 94
- Microsoft Assessment and Planning Toolkit, 110
- “Microsoft Clustering on VMware vSphere: Guidelines for Supported Configurations,” 345
- Microsoft System Center, 110
- migration, 54
- Min Server Memory, 235
- minimum memory, configuring, 392
- mixed workload environment with memory reservations, 226-228
- MLAG (Multi-Link Aggregation Group), 273
- %MLMTD metric, 47-48
- models (benchmark workload)
 - based on recorded production performance, 318
 - based on system nonfunctional requirements, 317
- monitoring, 54
- Multi-Link Aggregation Group (MLAG), 273
- Multi-NIC vMotion, 276-278
- multipathing of storage paths, 184-185, 280
- multiple tempdb files, creating, 394-395

N

- National Institute of Standards and Technology (NIST), 291
- n_browse_overall value, 433
- NDFS (Nutanix Distributed File System), 207
- network connections, validating, 359
- network design, 264
 - Multi-NIC vMotion, 276-278
 - NIOC (Network IO Control), 101, 274-276
 - physical network adapters, 267-269
 - storage, 279-280
 - teaming and failover, 270-273
 - virtual switches, 265-267
- Network IO Control (NIOC), 101, 274-276
- network paths, verifying, 262
- network security, 103, 281-284
- network teaming, 270-273
- network virtualization, 281-284
- New Availability Group Wizard, 399-405
- NIOC (Network IO Control), 101, 274-276
- NIST (National Institute of Standards and Technology), 291
- N%L counter, 326
- n_newcust_overall value, 433
- non-production workload influences on performance, 331-332
- non-shared disks, 345
- non-uniform memory architecture.
See NUMA

n_overall value, 432
 n_purchase_from_start value, 434
 n_purchase_overall value, 433
 n_rollbacks_from_start value, 434
 n_rollbacks_overall value, 434
 NTFS allocation unit size, 126-127
 NUMA (non-uniform memory architecture), 79-85
 explained, 241-243
 NUMA Scheduler, 81
 vNUMA, 82, 243
 Wide NUMA, 81
 NUMA Scheduler, 81
 Nutanix Bible (Poitras), 210
 Nutanix Distributed File System (NDFS), 207
 Nutanix Virtual Computing Platform, 207-213

O

OLAP (Online Analytical Processing), 64
 OLTP (Online Transaction Processing), 64
 one-to-many relationships, 40
 one-to-one relationships and unused capacity, 38-40
 Online Analytical Processing (OLAP), 64
 Online Transaction Processing (OLTP), 64
 operating systems. *See* OSs
 opm value, 432

OSs (operating systems)
 application binaries, and page file, 122
 compared to hypervisor, 26-27
 guest operating systems, 27

P

PAGEIOLATCH, 124
 PAGELATCH_XX, 112
 Page Life Expectancy metric, 323
 pages
 explained, 219-220
 large pages
 breaking down into default page size, 238-239
 explained, 237-238
 locking pages in memory, 239-241
 paging, 78, 220-221
 swapping, 78, 220-221
 TPS (transparent page sharing), 228-229
 Pages/Sec metrics, 321
 paging, 78, 220-221
 Paging File(_Total): %Usage metric, 321
 parallelism of storage design, 106
 paravirtualization, 29
 Paravirtualized SCSI (PVSCSI), 31, 95, 137-141
 partition alignment, 128-129
 partitioning, 28
 patches, 68
 pause frames, 268
 penalties (RAID IO), 189-194

Perfmon counters, 50, 323-324

performance baselines, 311-315

baseline performance reports, 332-333

benchmarks, 315

developing, 317-318

industry-standard benchmarks,
316

validating performance with, 318

vendor benchmarks, 316-317

comparing, 319-327

different processor generations,
330-331

different processor types, 328-330

common performance traps

blended peaks of multiple
systems, 335

failure to consider SCU (Single
Compute Unit) performance,
335

invalid assumptions, 334

lack of background noise, 334

shared core infrastructure
between production and non-
production, 333-334

vMotion slot sizes of monster
database virtual machines,
336-337

and database virtualization, 16-17

metrics

ESXTOP counters, 325-327

SQL Server baseline
infrastructure metrics, 321-322

SQL Server Perfmon counters,
323-324

SQL Server Profiler counters,
324-325

non-production workload influences
on performance, 331-332

server-side flash acceleration, 198-199

Fusion-io ioTurbine, 201-203

PernixData FVP, 204-206

vSphere Flash Read Cache
(vFRC), 199-201

storage

KISS principle (Keep It
Standardized and Simple), 109

performance and underlying
storage devices, 107

sizing for performance before
capacity, 107-108

virtualization without
compromise, 108-109

validating performance with, 318

vSphere storage design, 164

multipathing, 184-185

number of data stores and data
store queues, 165-169

number of virtual disks per data
store, 170-173

RAID (Redundant Array of
Independent Disks), 187-197

storage DRS, 177-183

Storage IO Control (SIOC),
173-177

storage policies, 177-183

vSphere 5.5 failover clustering
environments, 185-186

when to record, 320

- performance tests, 339
 - affinity and anti-affinity rules, 358
 - AlwaysOn Availability Groups configuration, 387-391
 - AlwaysOn Availability Groups creation, 399-405
 - Dell DVD Store installation and configuration, 406-430
 - Dell DVD Store load test, 430-436
 - Hot-Add Memory and Hot-Add CPU, 356-358
 - jumbo frame configuration, 393-394
 - lab configuration, 342-345
 - max/min memory configuration, 392
 - memory reservations, 355
 - multiple tempdb files, 394-395
 - network connection validation, 359
 - reasons for performance testing, 339-340
 - requirements
 - computer names and IP addresses, 341-342
 - resources, 342
 - software, 341
- SQL Server 2012 installation, 374-377, 384-387
 - Complete page, 387
 - Database Engine Configuration page, 382
 - Disk Space Requirements page, 382
 - Error Reporting page, 385
 - Feature Selection page, 379
 - Installation Configuration Rules page, 385
 - Installation Rules page, 380
 - Instance Configuration page, 380
 - License Terms page, 377
 - preflight check, 375
 - Product Key page, 375
 - Ready to Install page, 385
 - Server Configuration page, 382
 - Setup Role page, 379
- test database creation, 396-398
- VMDK file configuration
 - inside guest operating system, 352-354
 - on virtual machines, 345-351
- Windows Failover Clustering configuration, 359-368
- Windows Failover Clustering quorum mode, 369-374
- Windows Failover Clustering validation, 368
- performance traps
 - blended peaks of multiple systems, 335
 - failure to consider SCU (Single Compute Unit) performance, 335
 - invalid assumptions, 334
 - lack of background noise, 334
 - shared core infrastructure between production and non-production, 333-334
 - vMotion slot sizes of monster database virtual machines, 336-337
- PernixData FVP, 204-206
- PernixData FVP Datasheet, 205

phases of virtualization implementation

- database consolidations, 53
- discovery, 53
- infrastructure adjustments, 53
- migration and deployment, 54
- monitoring and management, 54
- requirements gathering, 51-52
- validation and testing, 54

physical hardware, 73

- adapter count, 95
- CPUs, 74-76
- data stores, 99
- disk layout, 95
- hardware compatibility, 62
- HTT (Hyper-Threading Technology), 85-87
- large pages, 79
- LSI Logic SAS adapters, 95
- memory, 76
- memory overcommitment, 87
- NUMA (non-uniform memory architecture), 79-85
- PVSCSI adapters, 95
- reservations, 87-89
- SQL Server Lock Pages in Memory, 92
- SQL Server Min Server Memory/Max Server Memory, 90-91
- storage, 93
- storage-specific metrics, 94
- swapping files, 78
- Thin Provisioning, 98-99
- virtualization overhead, 76-77
- VMDKs, 99

- file size, 100

- provisioning types, 96-98
- versus RDM, 96

Physical Mode RDMs (pRDM), 159

physical network adapters, 267-269

PKTRX/s counter, 327

PKTTX/s counter, 327

plans. *See* implementation plans

Pointer Block Eviction Process, 163-164

Poitras, Steven, 210

policies (storage), 177-183

pool, database buffer, 219-220

port groups, 265

power management policies, 253

pRDM (Physical Mode RDMs), 159

processors baseline comparisons

- between different processor generations, 330

- between different processor types, 328-330

Processor(_Total):Privileged Time metric, 321

Processor(_Total)[metric] Processor Time metric, 321

Product Key page (SQL Server installation), 375

Product Updates page (SQL Server installation), 377

Professional Association of SQL Server (PASS)

- PASS—Virtualization Virtual Chapter, 441

- PASS SQLSaturday, 441-442

Profiler counters, 324-325
protocols, 279-280
provisioning
 and database virtualization, 17-20
 VMDK, 96-98
PVSCSI (Paravirtualized SCSI), 31, 95,
137-141

Q

QFULL SCSI sense code, 166
QLogic HBA, 166
QoS, 102-103
Query Plan Optimizer, 130
queues, data store, 165-169
quorum mode (Failover Clustering),
369-374

R

RAID (Redundant Array of
Independent Disks), 187
 economics of RAID performance,
 194-197
 IO penalties, 189-194
 randomness of IO pattern, 187-188
 read/write bias, 188
RAM, 87
randomness of IO pattern, 187-188
Raw Device Map (RDM), 149-151
RDM (Raw Device Map), 149-151
 versus VMDK, 96
%RDY counter, 47, 325

reading documentation, 43-44
READs/s counter, 47-48, 326
read/write bias, 188
Ready to Install page (SQL Server
installation), 385
recovery
 disaster recovery, 291-294
 recovery point objectives (RPOs),
 45-46, 290
 recovery time objectives (RTOs),
 45-46, 290
 vCenter SRM (Site Recovery
 Manager), 301
recovery point objectives (RPOs), 45-46,
290
recovery time objectives (RTOs), 45-46,
290
reducing expenses, 9-10
Redundant Array of Independent Disks.
 See RAID
relationships
 one-to-many relationships, 40
 one-to-one relationships, 38-40
reorganizing SQL workloads, 68-69
replication (vSphere), 300-301
reports (performance), 332-333
requirements gathering, 51-52
reservation (memory), 87-89, 355
 explained, 232-233
 mixed workload environment with
 memory reservations, 226-228
 VMware HA strict admission control,
 233
 vswap file, 233-234

RESETS/s counter, 327

resource pools, Network IO Control, 275

resources

- blogs
 - Thomas LaRockTs blog, 445
 - vLaunchPad, 444-445
- documentation and white papers, 437-440
- Twitter, 445-446
- user groups, 440
 - Facebook groups, 443
 - PASS (Professional Association of SQL Server), 441-442
 - VMUG (VMware Users Group), 440
 - VMWare Community, 442-443

responsibility domains, 60-61

rollback_rate value, 434

Route Based on Physical NIC Load, 271

RPOs (recovery point objectives), 45-46, 290

rt_browse_avg_msec value, 433

rt_login_avg_msec value, 433

rt_newcust_avg_msec value, 433

RTOs (recovery time objectives), 45-46, 290

rt_purchase_avg_msec value, 433

rt_tot_avg, 433

rt_tot_avg value, 433

rt_tot_lastn_max value, 432

rules, affinity/anti-affinity, 358

S

SameSubnetDelay, 255

SameSubnetThreshold, 255

sample high availability chart, 308-309

SAP

- benchmark examples between different processor generations, 330-331
- benchmark examples between different processor types, 328-330
- Data Compression with, 133

SCU (Single Compute Unit), 335

security, 281-284

Server Configuration page (SQL Server installation), 382

server-side flash acceleration, 198-199

- Fusion-io ioTurbine, 201-203
- PernixData FVP, 204-206
- vSphere Flash Read Cache (vFRC), 199-201

service-level agreements (SLAs), 11-12, 45-46, 100, 290

settings (BIOS), 12-13

“Setup for Failover Clustering and Microsoft Cluster Service” (white paper), 439

Setup Role page (SQL Server installation), 379

shared core infrastructure between production and non-production, 333-334

shared disks, 345

shared environments, 20, 35

shared-resource world, 246

- SIOC (Storage IO Control), 157, 173-177
- Site Recovery Manager (SRM), 16, 301
- sizing
 - Heap size, 160-162
 - database files, 114-116
 - data file sizing, 116
 - Temp DB file sizing, 116
 - transaction log files, 117-120
 - databases, 22-23
 - performance before capacity, 107-108
 - VMs (virtual machines), 244
- SLAs (service-level agreements), 11-12, 45-46, 100, 290
- software requirements for performance testing, 341
- SP_Configure command, 246
- SQL AlwaysOn Failover Cluster Instances, 157
- SQL Compilations/sec metric, 324
- SQL Re-Compilations/sec metric, 324
- SQL Server 2012 installation, 374-387
 - Complete page, 387
 - Database Engine Configuration page, 382
 - Disk Space Requirements page, 382
 - Feature Selection page, 379
 - Installation Configuration Rules page, 385
 - Installation Rules page, 380
 - Instance Configuration page, 380
 - License Terms page, 377
 - preflight check, 375
 - Product Key page, 375
 - Ready to Install page, 385
 - Server Configuration page, 382
 - Setup Role page, 379
- SQL Server 2014 & The Data Platform (data sheet), 247
- “SQL Server Best Practices” (white paper), 210
- SQL Server Max Server Memory, 90-91
- SQL Server Min Server Memory, 90-91
- “SQL Server on VMware—Availability and Recovery Options” (white paper), 439
- “SQL Server on VMware—Best Practices Guide” (white paper), 439
- SQL Server SysPrep, 71
- SQL Statistics, 324
- SQL workloads, 64-67
 - Batch/ETL workloads, 64
 - DSS workloads, 64
 - OLAP workloads, 64
 - OLTP workloads, 64
 - reorganization, 68-69
 - tiered database offering, 70-73
- SQLIOSim, 327
- SQLSaturday, 441-442
- SRM (Site Recovery Manager), 16, 301
- stack, memory trends and, 79, 218
 - database buffer pool, 219-220
 - database indexes, 222-225
 - database pages, 219-220
 - paging, 220-221
 - swapping, 220-221

storage

design principles

- database as extension of storage, 106

- KISS principle (Keep It Standardized and Simple), 109

- performance and underlying storage devices, 107

- sizing for performance before capacity, 107-108

- virtualization without compromise, 108-109

- overview, 93, 105-106

- server-side flash acceleration, 198-199

 - Fusion-io ioTurbine, 201-203

 - PernixData FVP, 204-206

 - vSphere Flash Read Cache (vFRC), 199-201

- SQL Server database and guest OS storage, 109

 - Buffer Pool, 129-130

 - Column Storage, 134-135

 - database availability, 135

 - database statistics, updating, 130-132

 - Data Compression, 133

 - file system layout, 110, 122-129

 - Instant File Initialization (IFI), 120-122

 - number of database files, 110-113

 - size of database files, 114-120

 - Storage Spaces, 136

 - Volume Managers, 136

- SQL Server on hyperconverged infrastructure, 207-213

- SQL Server virtual machine storage, 136

 - expanding, 158-159

 - Jumbo VMDKs, 159-164

 - layout, 152-157

 - virtual disk devices, 143-152

 - virtual storage controllers, 138-143

 - VM hardware version, 137

- Storage Acceleration, 96

- storage arrays, 98

- Storage DRS, 177-183, 297

- Storage IO Control (SIOC), 157, 173-176

- storage networks, 279-280

- storage policies, 177-183

- storage protocols, 279-280

- Storage Spaces, 136

- storage-specific metrics, 94vSphere storage design, 164

 - multipathing, 184-185

 - number of data stores and data store queues, 165-169

 - number of virtual disks per data store, 170-173

 - RAID (Redundant Array of Independent Disks), 187-197

 - storage DRS, 177-183

 - Storage IO Control (SIOC), 173-177

 - storage policies, 177-183

 - vSphere 5.5 failover clustering environments, 185-186

- Storage Acceleration, 96

- Storage DRS, 177-183, 297

Storage IO Control (SIOC), 157, 173-176
 Storage Spaces, 136
 Storage vMotion, 297
 Strawberry Perl, 406
 swap files, 78
 Swapin counter, 326
 Swapout counter, 326
 swapping, 78, 220-221
 switches
 virtual switches, 265-267
 distributed virtual switch, 100
 port groups, 265
 teaming methods, 270
 vSphere distributed switches, 263
 vSS, 265
 %SWPWT counter, 326
 %SYS counter, 325
 Szastak, Jeff, 311

T

Target Server Memory (KB) metric, 324
 TDE (Transparent Data Encryption), 122
 teams
 Center of Excellence (CoE), 61-63
 communication, 58
 mutual understanding, 59-60
 responsibility domains, 60-61
 Temp DB files
 file system layout, 123-126
 multiple tempdb files, creating, 394-395
 sizing, 116
 test databases, creating, 396-398
 testing, 54
 baseline. *See* performance baselines
 benchmarks, 315
 developing, 317-318
 industry-standard benchmarks, 316
 vendor benchmarks, 316-317
 jumbo frames, 262-264
 performance tests, 339
 affinity and anti-affinity rules, 358
 AlwaysOn Availability Groups
 configuration, 387-391
 AlwaysOn Availability Groups
 creation, 399-405
 Dell DVD Store installation and
 configuration, 406-430
 Dell DVD Store load test,
 430-436
 Hot-Add Memory and Hot-Add
 CPU, 356-358
 jumbo frame configuration,
 393-394
 lab configuration, 342-345
 max/min memory configuration,
 392
 memory reservations, 355
 multiple tempdb files, 394-395
 network connection validation,
 359
 reasons for performance testing,
 339-340
 requirements, 341-342

- SQL Server 2012 installation, 374-387
 - test database creation, 396-398
 - VMDK file configuration, 345-354
 - Windows Failover Clustering configuration, 359-368
 - Windows Failover Clustering quorum mode, 369-374
 - Windows Failover Clustering validation, 368
 - Thick Eager Zero disks, 147-148
 - Thick Lazy Zero disks, 144-146
 - Thick Provisioned Eager Zeroed, 97
 - Thick Provisioned Lazy Zeroed, 97
 - Thick Provisioned LUNs, 98
 - Thin disks, 144-146
 - Thin Provisioned LUNs, 98
 - Thin Provisioned VMDK, 97-99
 - Thin Provisioning, 98-99
 - tiered database offering, 70-73
 - tiers (database), 19-20
 - TLB (translation lookaside buffer), 79
 - TPC (Transaction Processing Performance Council), 316
 - TPS (transparent page sharing), 79, 228-229
 - trace flags
 - enabling, 215
 - list of, 215
 - trace flag 2371, 131
 - traffic types, 101-102
 - transaction log files, sizing, 117-120
 - Transaction Processing Performance Council (TPC), 316
 - translation lookaside buffer (TLB), 79
 - Transactions/sec metric, 324
 - Transparent Data Encryption (TDE), 122
 - transparent page sharing (TPS), 79, 228-229
 - troubleshooting common performance traps
 - blended peaks of multiple systems, 335
 - failure to consider SCU (Single Compute Unit) performance, 335
 - invalid assumptions, 334
 - lack of background noise, 334
 - shared core infrastructure between production and non-production, 333-334
 - vMotion slot sizes of monster database virtual machines, 336-337
 - tuning virtual network adapters, 252-254
 - Twitter, 445-446
 - Type-1 hypervisors, 30
 - Type-2 hypervisors, 31
- ## U
-
- Understanding Memory Resource Management in VMware vSphere 5.0 (study), 229-230
 - “Understanding VMware vSphere 5.1 Storage DRS” (white paper), 182
 - unused capacity and one-to-one relationships, 38-40

updating database statistics, 130-132
 with maintenance plan, 131-132
 with trace flag 2371, 131
%USED counter, 325
User Connections metric, 324
user groups, 440
 Facebook groups, 443
 PASS (Professional Association of SQL Server)
 PASS—Virtualization Virtual Chapter, 441
 PASS SQLSaturday, 441-442
 VMUG (VMware Users Group), 440
 VMWare Community, 442-443

V

VAAI (vStorage APIs for Array Integration), 96
Validate a Configuration Wizard, 364
validation, 54
 cluster network configuration, 368
 network connections, 359
 performance with baselines/ benchmarks, 318
vCenter Hyperic, 343-345
vCenter Infrastructure Navigator, 343-344
vCenter Operations Manager, 87
vCenter SRM (Site Recovery Manager), 301
vCHS (vCloud Hybrid Service), 302
vCloud Hybrid Service (vCHS), 302
vCPUs (virtual CPUs), 4
vDP (vSphere Data Protection), 266, 300
 Multi-NIC vMotion, 277
vDS (vSphere distributed switch), 262
vendor benchmarks, 316-317
verifying network paths, 262
Virtual Computing Platform (Nutanix), 207-213
virtual CPUs (vCPUs), 4
virtual disks, number per data store, 170-173
Virtual Log Files (VLFs), 118-120
virtuallyGhetto, 444
virtual machine storage, 136
 expanding, 158-159
 Jumbo VMDKs, 159
 Pointer Block Eviction Process, 163-164
 VMFS Heap size considerations, 160-162
layout, 152-157
number of VMs, 244-245
sizing, 244
virtual disk devices, 143
 IO blender effect, 151-152
 Raw Device Map (RDM), 149-151
 Thick Eager Zero disks, 147-148
 Thin versus Thick Lazy Zero disks, 144-146
virtual storage controllers, 138-143
 VM hardware version, 137
 VM memory, 225-226
Virtual Mode RDMs (vRDM), 159

- virtual network adapters, 100
 - choosing, 250-251
 - traffic types, 101-102
 - turning, 252-254
- virtual server access ports, 281
- virtual switches, 265-267
 - distributed virtual switch, 100
 - port groups, 265
 - teaming methods, 270
- virtualization, 93
 - advantages of, 3
 - hot-add CPUs, 3-4
 - hot-add memory, 4-5
 - business case for, 9
 - BIOS settings, 12-13
 - DBA (database administrator) advantages, 10-11
 - hardware refresh, 20-22
 - high availability, 14-16
 - large databases, 22-23
 - performance, 16-17
 - provisioning/DBaaS, 17-20
 - reduced expenses, 9-10
 - SLAs (service level agreements), 11-12
 - and compromise, 108-109
 - documentation, 43-44
 - explained, 1-2
 - hypervisor, 25
 - compared to OS, 26-27
 - explained, 25-27
 - importance of full virtualization, 36-38
 - one-to-many relationships, 40
 - one-to-one relationships and unused capacity, 38-40
 - paravirtualization, 29
 - PVSCSI (paravirtual SCSI driver), 31
 - Type-1 hypervisors, 30
 - Type-2 hypervisors, 31
 - virtualized database installation guidelines, 32-36
 - VMs (virtual machines), 28
 - VMware ESXi versions, 40-41
 - VMXNET3, 32
 - implementation plan
 - database workload baselines, 48-50
 - high-level plan, 50-51
 - items to consider, 44-45
 - phase 1: requirements gathering, 51-52
 - phase 2: discovery, 53
 - phase 2.1: database consolidations, 53
 - phase 3: infrastructure adjustments, 53
 - phase 4: validation and testing, 54
 - phase 5: migration and deployment, 54
 - phase 6: monitoring and management, 54
 - RPOs (recovery point objectives), 45-46
 - RTOs (recovery time objectives), 45-46

- SLAs (service-level agreements), 45-46
- vSphere infrastructure baselines, 46-48
- importance of full virtualization, 36-38
- overhead, 76-77
- performance baselines
 - baseline performance reports, 332-333
 - benchmarks, 315-318
 - common performance traps, 333-337
 - comparing, 328-331
 - explained, 311-315
 - metrics, 321-327
 - non-production workload influences on performance, 331-332
 - reasons for, 319-320
 - validating performance with, 318
 - when to record, 320
- power company example, 6
- world before database virtualization, 5-6
- “Virtualization Overview” (white paper), 21
- virtualized database installation guidelines, 32-36
- virtualized security zones, 283
- vLaunchPad, 444-445
- VLFs (Virtual Log Files), 118-120
- VMDKs, 99
 - files, configuring
 - file size, 100
 - inside guest operating system, 352-354
 - on virtual machines, 345-351
 - Jumbo VMDKs, 159-160
 - Pointer Block Eviction Process, 163-164
 - VMFS Heap size considerations, 160-162
 - provisioning types, 96-98
 - versus RDM, 96
 - virtual machine storage layout, 152-157
- VMFS heap size considerations, 160-162
- vMotion, 296-297
 - slot sizes of monster database virtual machines, 336-337
 - traffic, 276
- VMs (virtual machines). *See* virtual machine storage
- VMUG (VMware Users Group), 440
- %VMWait counter, 326
- VMware App Director, 70
- VMware Capacity Planner, 110
- VMWare Community, 442-443
- VMware ESXi versions, 40-41
- VMware HA strict admission control, 233
- VMware NSX, 283
- VMware PowerCLI, 253
- VMware Site Recovery Manager (SRM), 16

- VMware vCloud Suite licenses, 285
- VMware vSphere on Nutanix Best Practices (white paper), 210
- VMXNET 3, 100
- VMXNET3, 32, 251, 257
- vNIC, 256
- vNUMA, 82, 243
- Volume Managers, 136
- vRAM, 87
- vRDM (Virtual Mode RDMs), 159
- vRFC (vSphere Flash Read Cache), 199-201
- vSphere, 98-99
 - baselining, 46-48
 - ESXTOP counters, 325-327
 - failover clustering environments, 185-186
 - high availability
 - DRS (Distributed Resource Scheduler), 297
 - hypervisor availability features, 294-296
 - Storage DRS, 297
 - Storage vMotion, 297
 - vCenter SRM (Site Recovery Manager), 301
 - vCHS (vCloud Hybrid Service), 302
 - vDP (vSphere Data Protection), 300
 - vMotion, 296-297
 - vSphere App HA, 299-300
 - vSphere HA, 298-299
 - vSphere Replication, 300-301
 - X-vMotion, 298
 - storage design, 164
 - multipathing, 184-185
 - number of data stores and data store queues, 165-169
 - number of virtual disks per data store, 170-173
 - RAID (Redundant Array of Independent Disks), 187-197
 - storage DRS, 177-183
 - Storage IO Control (SIOC), 173-177
 - storage policies, 177-183
 - vSphere 5.5 failover clustering environments, 185-186
 - vDS (vSphere distributed switch), 262-263
 - vFRC (vSphere Flash Read Cache), 199-201
 - vNUMA, 243
 - vSphere App HA, 299-300
 - vSphere HA, 298-299
 - vSphere Hot Plug Memory, 91
 - vSphere Replication, 300-301
 - vSphere Web Client, 260
 - vSS (vSphere standard switch), 265, 271
 - “vSphere Monitoring & Performance” (white paper), 440
 - “vSphere Resource Management” (white paper), 440
 - “vSphere Storage” (white paper), 440
 - vSphere Web Client, 260
 - vSS (vSphere standard switch), 265, 271

vStorage APIs for Array Integration (VAAI), 96

vswap file, memory reservations and, 233-234

W

Webster, Michael, 106, 170

white papers

“DBA Guide to Databases on VMware,” 439

“Setup for Failover Clustering and Microsoft Cluster Service,” 439

“SQL Server Best Practices,” 210

“SQL Server on VMware—Availability and Recovery Options,” 439

“SQL Server on VMware—Best Practices Guide,” 439

“Understanding VMware vSphere 5.1 Storage DRS,” 182

“VMware vSphere on Nutanix Best Practices,” 210

“vSphere Monitoring & Performance,” 440

“vSphere Resource Management,” 440

“vSphere Storage,” 440

Wide NUMA, 81

Windows Failover Cluster Heartbeat settings, 255

Windows Failover Clustering
configuring, 359-368
network settings, 254-256

quorum mode, 369-374

validating, 368

Windows Server Failover Clustering (WSFC), 304

Windows vNIC properties page, 261

wizards

Cluster Validation Wizard, 363-364

Configure Cluster Quorum Wizard, 370-374

Create Cluster Wizard, 366

Dell DVD Store Custom Install Wizard, 408

Maintenance Plan Wizard, 420-422

New Availability Group Wizard, 399-405

Validate a Configuration Wizard, 364

Workload Driver Configuration Wizard, 409

Workload Driver Configuration Wizard, 409

workloads

based on recorded production performance, 318

based on system nonfunctional requirements, 317

baselining, 48-50

mixed workload environment with memory reservations, 226-228

SQL workloads, 64-67

reorganization, 68-69

tiered database offering, 70-73

worlds (VMs), 86

Writes/s counter, 326

WSFC (Windows Server Failover Clustering), 304

X-Y-Z

xVelocity memory, 134

X-vMotion, 298

Yellow-Bricks, 444