

THE EXPERT'S VOICE® IN SQL SERVER

Pro SQL Server 2012 Practices

*CAREER-BUILDING INSIGHTS INTO SQL
SERVER DATABASE ADMINISTRATION
AND PERFORMANCE OPTIMIZATION*

Bradley Ball, TJay Belt, Glenn Berry, Jes Borland, Carlos Bossy, Louis Davidson, Ben DeBow, Grant Fritchey, Jonathan Gardner, Jesper Johansen, Jeremy Lowell, Wendy Pastrick, Kellyn Pot'vin, Mladen Prajdić, Herve Roggero, Chris Shaw, Gail Shaw, and Jason Strate

Apress®

Pro SQL Server 2012 Practices



**Bradley Ball, TJay Belt, Glenn Berry,
Jes Borland, Carlos Bossy, Louis Davidson,
Ben DeBow, Grant Fritchey, Jonathan Gardner,
Jesper Johansen, Jeremy Lowell, Wendy Pastrick,
Kellyn Pot'vin, Mladen Prajdić, Herve Roggero,
Chris Shaw, Gail Shaw, Jason Strate**

Apress®

Pro SQL Server 2012 Practices

Copyright © 2012 by Bradley Ball, TJay Belt, Glenn Berry, Jes Borland, Carlos Bossy, Louis Davidson, Jeremy Lowell, Ben DeBow, Grant Fritchey, Wendy Pastrick, Kellyn Pot'vin, Jonathan Gardner, Jesper Johansen, Mladen Prajdić, Herve Roggero, Chris Shaw, Gail Shaw, Jason Strate

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

ISBN 978-1-4302-4770-8

ISBN 978-1-4302-4771-5 (eBook)

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

President and Publisher: Paul Manning

Lead Editor: Jonathan Gennick

Developmental Editor: Brian MacDonald

Technical Reviewer: Rudi Bruchez, Robin Dewson, David Dye, Edgar Lanting, Ken Simmons

Editorial Board: Steve Anglin, Ewan Buckingham, Gary Cornell, Louise Corrigan, Morgan Ertel, Jonathan Gennick, Jonathan Hassell, Robert Hutchinson, Michelle Lowman, James Markham, Matthew Moodie, Jeff Olson, Jeffrey Pepper, Douglas Pundick, Ben Renow-Clarke, Dominic Shakeshaft, Gwenan Spearing, Matt Wade, Tom Welsh

Coordinating Editor: Anamika Panchoo

Copy Editors: Roger LeBlanc, Vanessa Moore

Compositor: Bytheway Publishing Services

Indexer: SPi Global

Artist: SPi Global

Cover Designer: Anna Ishchenko

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com.

For information on translations, please e-mail rights@apress.com, or visit www.apress.com.

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Special Bulk Sales—eBook Licensing web page at www.apress.com/bulk-sales.

Any source code or other supplementary materials referenced by the author in this text is available to readers at www.apress.com. For detailed information about how to locate your book's source code, go to www.apress.com/source-code.

To my wife, Silva, and my children, Zachary, William, Serenity, and Chesney: Thank you for putting up with my long hours on the computer. I love you all No Matter What!

—Bradley Ball

To Ali, Adam, and Gaby, for all they taught me about writing, computer science, and devotion.

—Carlos Bossy

I dedicate my chapter on database design to the poor programmers having to suffer with malformed databases.

—Louis Davidson

I type this on the date 9/11/2012. I've never dedicated one of my books before because I never saw the point. But this time, I'd like to dedicate my contribution to this book to all the first responders and those who serve in the military. Thanks for everything you do to protect our civilization.

—Grant Fritchey

To my lovely wife, Birgitte; my wonderful children, Alexander, Andreas, Jakob, Mathilde; my father, Ib; my best friend, Lars; and not least, Morten.

—Jesper Johansen

I'd like to dedicate my chapter to my children, Sam, Cait, and Josh. Everything I do, I do for them—as a Mom, a caregiver, and most of all, as an example for them to take into this world.

—Kellyn Pot'Vin

I'd like to dedicate this book to my mom, Desanka, and sister, Tajana, who are awesome supporters of my endeavors.

—Mladen Prajdić

How can I dedicate my work to all those who have influenced me over my lifetime in a few short words? Only one word is needed, family.

To my wife, Gigi; my son, Mike; my mom, Ruth; our best friend, Rock; and brother, Anthony. To my Italian by association family Dick, Linda, Grandma (Barbara), and Grandpa (Don).

For my Marine brothers and sisters, and those who stand or stood next to us.

For my #SQLFamily: we can do it together.

For my Dad, Herb: you may be gone but you will never be forgotten.

—Chris Shaw

To Nikolai, Aspen, and Dysin who missed out on a few rounds of Dungeon Defenders while I worked on this book.

—Jason Strate

Contents at a Glance

■ About the Authors.....	xiv
■ About the Technical Reviewers.....	xix
■ Acknowledgments.....	xxi
■ Introduction	xxiii
■ Chapter 1: Be Your Developer's Best Friend	1
■ Chapter 2: Getting It Right: Designing the Database for Performance.....	17
■ Chapter 3: Hidden Performance Gotchas	43
■ Chapter 4: Dynamic Management Views.....	71
■ Chapter 5: From SQL Trace to Extended Events	101
■ Chapter 6: The Utility Database	135
■ Chapter 7: Indexing Outside the Bubble	161
■ Chapter 8: Release Management.....	197
■ Chapter 9: Compliance and Auditing.....	221
■ Chapter 10: Automating Administration.....	235
■ Chapter 11: The Fluid Dynamics of SQL Server Data Movement.....	271
■ Chapter 12: Windows Azure SQL Database for DBAs	293
■ Chapter 13: I/O: The Untold Story	313
■ Chapter 14: Page and Row Compression	335
■ Chapter 15: Selecting and Sizing the Server.....	361
■ Chapter 16: Backups and Restores Using Availability Groups	375
■ Chapter 17: Big Data for the SQL Server DBA	395
■ Chapter 18: Tuning for Peak Load.....	429
■ Index	465

Contents

■ About the Authors.....	xiv
■ About the Technical Reviewers.....	xix
■ Acknowledgments.....	xxi
■ Introduction	xxiii
■ Chapter 1: Be Your Developer's Best Friend	1
My Experience Working with SQL Server and Developers	1
Reconciling Different Viewpoints Within an Organization.....	2
Preparing to Work with Developers to Implement Changes to a System	3
Step 1: Map Your Environment	3
Step 2: Describe the New Environment.....	5
Step 3: Create a Clear Document	7
Step 4: Create System-Management Procedures	7
Step 5: Create Good Reporting	10
Ensuring Version Compatibility.....	11
Setting Limits	12
Logon Triggers.....	12
Policy-Based Management.....	15
Logging and Resource Control.....	15
Next Steps	15
■ Chapter 2: Getting It Right: Designing the Database for Performance.....	17
Requirements	18
Table Structure	20
A Really Quick Taste of History	20

Why a Normal Database Is Better Than an Extraordinary One	21
Physical Model Choices	33
Design Testing	40
Conclusion	41
■ Chapter 3: Hidden Performance Gotchas	43
Predicates	43
Residuals	59
Spills	65
Conclusion	70
■ Chapter 4: Dynamic Management Views	71
Understanding the Basics	71
Naming Convention	72
Groups of Related Views	72
Varbinary Hash Values	73
Common Performance-Tuning Queries	74
Retrieving Connection Information	74
Showing Currently Executing Requests	75
Locking Escalation	77
Finding Poor Performing SQL	78
Using the Power of DMV Performance Scripts	80
Divergence in Terminology	82
Optimizing Performance	82
Inspecting Performance Stats	85
Top Quantity Execution Counts	86
Physical Reads	87
Physical Performance Queries	87
Locating Missing Indexes	87
Partition Statistics	92
System Performance Tuning Queries	93
What You Need to Know About System Performance DMVs	93
Sessions and Percentage Complete	93
Conclusion	99
■ Chapter 5: From SQL Trace to Extended Events	101
SQL Trace	101
Trace rowset provider	103

Trace file provider	107
Event Notifications	110
Extended Events	114
Events	115
Predicates	115
Actions	116
Types and Maps	116
Targets	117
Sessions	118
Built in Health Session	121
Extended Events .NET provider	123
Extended Events UI	125
Conclusion	133
■ Chapter 6: The Utility Database	135
Start with Checklists	136
Daily Checklist Items	136
Longer-Term Checklist Items	137
Utility Database Layout	138
Data Storage	138
Using Schemas	140
Using Data Referential Integrity	140
Creating the Utility Database	140
Table Structure	141
Gathering Data	143
System Tables	143
Extended Stored Procedures	143
CLR	144
DMVs	144
Storage	144
Processors	146
Error Logs	148
Indexes	149
Stored Procedure Performance	151
Failed Jobs	152
Reporting Services	153
Mirroring	154

AlwaysOn	156
Managing Key Business Indicators	156
Using the Data	158
Automating the Data Collection	158
Scheduling the Data Collection	159
Conclusion	160
■ Chapter 7: Indexing Outside the Bubble	161
The Environment Bubble	162
Identifying Missing Indexes	162
Index Tuning a Workload	170
The Business Bubble	191
Index Business Usage	191
Data Integrity	193
Conclusion	195
■ Chapter 8: Release Management	197
My Release Management Process	197
A Change Is Requested	198
Release Process Overview	199
Considerations	199
Documents	207
Release Notes	208
Release Plan Template and Release Plans	212
Document Repository	219
Conclusion	219
■ Chapter 9: Compliance and Auditing	221
Compliance	221
Sarbanes-Oxley	221
Health Insurance Portability and Accountability Act	223
New Auditing Features in SQL Server 2012	224
Server-Level Auditing for the Standard Edition	225
Audit Log Failure Options	225
Maximum Rollover Files	225
User-Defined Auditing	225
Audit Filtering	225
Auditing	226

Server Audit.....	226
Server Audit Specification	228
Database Audit Specification.....	230
Query the Audit File	231
Pro Tip: Alert on Audit Events.....	232
Conclusion.....	234
■ Chapter 10: Automating Administration.....	235
Tools for Automation	235
Performance Monitor.....	235
Dynamic Management Views	237
SQL Server Agent	238
Maintenance Plans	252
SQL Server Integration Services	259
PowerShell	262
What to Automate	263
Monitoring	264
Backups and Restores.....	267
Database Integrity	269
Index Maintenance	269
Statistics Maintenance	270
Conclusion	270
■ Chapter 11: The Fluid Dynamics of SQL Server Data Movement.....	271
Why the Need for Replicating Data?.....	271
SQL Server Solutions.....	273
Replication.....	274
Log Shipping.....	278
Database Mirroring.....	280
AlwaysOn	282
Failover Clustering	284
Custom ETL Using SQL Server Integration Services.....	286
Bulk Copy Process.....	287
Choosing the Right Deployment	288
Keeping the Data Consistent	290
Conclusion.....	292

■ Chapter 12: Windows Azure SQL Database for DBAs	293
SQL Database Architecture	294
Infrastructure.....	294
Availability and Failover.....	295
Hardware	295
Differences with SQL Server	296
Database Components.....	296
Management Platform	297
Security	298
Other Important Information	299
Federations.....	300
Key Terms	300
T-SQL Changes for Federations	301
Federation Example.....	302
Limitations.....	303
Troubleshooting Performance Issues	304
DMVs Available	304
Execution Plans	305
Performance Dashboard.....	306
Related Services.....	308
Windows Azure SQL Reporting	308
Windows Azure SQL Data Sync.....	309
Import/Export Feature.....	310
Cost of SQL Database	311
Conclusion.....	312
■ Chapter 13: I/O: The Untold Story	313
The Basics	314
Monitoring	314
Considerations.....	315
Tactical	317
Code or Disk?	321
Times Have Changed.....	323
Getting to the Data	324
Addressing a Query	328
Environmental Considerations.....	331

Conclusion.....	334
■ Chapter 14: Page and Row Compression	335
Before You Get Started	336
Editions and Support	336
What to Compress and How to Compress It	337
Row Compression.....	338
Page Compression.....	341
What Do You Compress?	346
Fragmentation and Logged Operations	355
Conclusion.....	359
■ Chapter 15: Selecting and Sizing the Server	361
Understanding Your Workload	361
SQL Server 2012 Enterprise Edition Consideration Factors	362
Server Vendor Selection	364
Server Form Factor Selection.....	364
Server Processor Count Selection	366
Dell 12th Generation Server Comparison	366
Dell PowerEdge R320	366
Dell PowerEdge R420	367
Dell PowerEdge R520	367
Dell PowerEdge R620	367
Dell PowerEdge R720	368
Dell PowerEdge R720xd	368
Dell PowerEdge R820	368
Dell Comparison Recap	368
Processor Vendor Selection.....	369
Processor Model Selection.....	370
Memory Selection	372
Conclusion.....	373
■ Chapter 16: Backups and Restores Using Availability Groups	375
Setting Up an Availability Group	376
Configuring the Windows Server	376
SQL Server Availability Group	377
Enabling Backups on Availability Groups.....	383

Backup Location	383
Backup Priority	384
Automating Backups on Availability Groups	386
Maintenance Plans	386
T-SQL Scripts	388
Recovery on Availability Groups	391
Conclusion	392
■ Chapter 17: Big Data for the SQL Server DBA	395
Big Data Arrives with Hadoop	397
MapReduce: The Nucleus of Hadoop	398
Hardware	404
DBA As Data Architect	405
Big Data for Analytics	406
Using SQL Server with Hadoop	407
The DBA's Role	407
Big Data in Practice	408
Exporting from HDFS	415
Hive	416
Hive and Excel	419
JavaScript	420
Pig	423
Big Data for the Rest of Us	425
Business Intelligence	425
Big Data Sources	425
Big Data Business Cases	426
Big Data in the Microsoft Future	427
Conclusion	428
■ Chapter 18: Tuning for Peak Load	429
Define the Peak Load	429
Determine Where You Are Today	431
Perform the Assessment	433
Define the Data to Capture	436
Analyze the Data	446
Analyzing Application-Usage Data	446
Analyzing Perfmon Data	449

Analyzing Configuration Data	455
Analyzing SQL Performance Data.....	458
Devise a Plan.....	462
Conclusion.....	463
■ Index	465

About the Authors

■ **Bradley Ball** is a husband and a father. During the work day, he is also a SQL DBA with over 10 years of IT experience. Bradley spent 8 years working as a defense contractor for clients such as the U.S. Army and The Executive Office of the President of the United States. He is currently a Senior Consultant with Pragmatic Works, specializing in DBA and BI solutions. He has presented at many SQL Saturdays, for the MAGICPASS & OPASS SSUGs, SQL Rally 2011 and 2012, DevConnections 2012, and at the PASS Summit in 2011 and 2012. Bradley can be found blogging on <http://www.sqlballs.com>.



■ **TJay Belt** is a professional database administrator and has been developing his expertise over the past eight years. He has been working in IT for over 18 years. He enjoys delivering technical presentations and has presented for Microsoft, Professional Association for SQL Server, SQL Server Worldwide User Group, PASS SQL Saturday events, and local SQL Server user groups. TJay has published technical articles for SQL Server Central, Simple Talk, and SQL-Server-Performance. He is the founder and was the leader of the PASS chapter of the SQL Server user group in Utah County. TJay has volunteered as a chapter leader for PASS, a regional mentor for PASS chapters, and the PASS Chapter Committee. He also was the founder of the Auditing & Compliance Special Interest Group with PASS. TJay currently volunteers as a regional mentor for PASS chapters.



■ **Glenn Berry** works as a database architect at Avalara in Parker, Colorado (USA). He has been a SQL Server MVP since 2007 and has a whole collection of Microsoft certifications, including MCITP, MCDBA, MCSE, MCSA, MCAD, and MCTS, which proves that he likes to take tests. His expertise includes DMVs, high availability, hardware selection, full text search, and SQL Azure. He is also an adjunct faculty member at University College at the University of Denver, where has been teaching since 2000. He has completed the Master Teacher Program at University College at the University of Denver. He is heavily involved in the SQL Server community and is a frequent speaker at user groups, SQL Saturdays, and the PASS Community Summit. He is the author of the book *SQL Server Hardware* (Red Gate Books, 2011), and he wrote chapters for *SQL Server MVP Deep Dives* (Manning Publications, 2009) and *SQL Server MVP Deep Dives Volume 2* (Manning Publications, 2011). Glenn blogs regularly at [http://](http://sqlserverperformance.wordpress.com)

sqlserverperformance.wordpress.com. He is active on Twitter as @GlennAlanBerry.

■ **Jes Borland** is a consultant with Brent Ozar PLF and a Microsoft SQL Server MVP. She holds an AAS - Programmer/Analyst degree and loves SQL Server Reporting Services and day-to-day administration. She is an active member of PASS, a user-group founder, a Tech on Tap brewmaster, and a LessThanDot.com blogger. She's a frequent presenter at user groups, SQL Saturdays, and other community events. She is also an avid runner and novice chef.



■ **Carlos Bossy** (MCTS, MCP SQL Server 2008 BI, CBIP) is a consultant with 25 years of experience in software and database development. As the Principal of his company, Carlos focuses on the development of Business Intelligence solutions, including data warehouses, predictive analytics, integration, and reporting. He has worked with SQL Server for 12 years and is very enthusiastic about its powerful features. Carlos has developed data warehouses and BI solutions for a variety of industries and state agencies and is a frequent speaker at SQL Server events.



■ **Louis Davidson** has been in the IT industry for what is starting to seem like a really long time as a corporate database developer and architect. He has been a Microsoft MVP for 9 years and has written five books on database design for Apress, including *Pro SQL Server 2012 Relational Database Design and Implementation* (2012). Louis has been active speaking on the topic of database design and implementation at many conferences over the past ten years, including SQL PASS, SQL Rally, SQL Saturday events, and CA World, as well as the Devlink developer conference. Louis has worked for the Christian Broadcasting Network as a developer, DBA, and Data Architect supporting offices in Virginia Beach, Virginia, and Nashville, Tennessee for over 14 years. Despite what his website at drsql.org implies, Louis has a bachelor's degree from the University of Tennessee at Chattanooga in computer science.



■ **Ben DeBow** is a consultant who works with customers of all sizes on their mission-critical applications and SQL Server deployments. His versatility allows him to assist with everything from virtualization and consolidation to schema design, performance tuning, and disk architectures. He has tuned and deployed some of the largest and most complex SQL Server solutions in the world. An active member of the SQL Server community, Ben is also an MCDBA, MCSE, and MCITP. He is also the co-founder of SQLHA with Allan Hirt, and can be reached at <http://www.sqlha.com>.

■ **Grant Fritchey** works for Red Gate Software—a market-leading, software-tools vendor—as its product evangelist. He has done development of large-scale applications in languages such as VB, C#, and Java; and he has worked in SQL Server since version 6.0. In addition, he has worked in insurance, finance, and consulting, as well as for three failed dot coms. He is the author of *SQL Server Execution Plans* (Simple Talk Publishing, 2008) and *SQL Server 2012 Query Performance Tuning* (Apress, 2012). Grant is a co-author of *Beginning SQL Server 2012 Administration* (Apress, 2012) and has one chapter in *SQL Server MVP Deep Dives Volume 2* (Manning Publications, 2011).



■ **Jonathan Gardner** is a senior systems architect, PMP, and MCITP working with Geocent, L.L.C. in New Orleans, Louisiana (USA). He has over 10 years of experience with a wide range of IT solutions, including database administration, virtualization, SAN administration, and Windows Server environments. He writes from his website <http://jonathanagardner.com>.



■ **Jesper Johansen** has been a SQL Server DBA for more than two decades. He is one of the heavy guys in the Danish SQL Server field.

Jesper joined Miracle Five years ago and has since developed efficient techniques for helping customers save time, money, and energy by consolidating their SQL Server installations.

Apart from consolidation, Jesper enjoys the occasional Backup/Restore game, helping customers realize that what they thought was lost can still be dug out from the belly of SQL Server.

As if the above wasn't enough to keep Jesper busy, he is also known to take on DB2 tasks from time to time. He is a board member of the Danish SQL Server User Groupe.

■ **Jeremy Lowell** is the founder of Data Realized and has been helping DBAs and clients with SQL Server for over a decade. Jeremy has a track record of solving complex issues for small and large enterprises at any point in the life cycle of a business. From startups to Fortune 100 enterprises, Jeremy has been a leader in implementing best practices and helping business leaders realize their data.

■ **Wendy Pastrick** (@wendy_dance) is a self-taught database administrator in Chicago, Illinois (USA) since 2000, replicating data across environments using various SQL methods. You can find her at SQL PASS (Professional Association for SQL Server) events, such as SQL Saturday or the annual Summit, either volunteering or presenting, or perhaps singing at an after-event gathering. With a degree in Psychology and previous technical training experience, she has taken the podium across the country at events talking about SQL Server, personal networking, and support for Women In Technology.

■ **Kellyn Pot'vin** has more than 15 years of experience in information technology and more than 12 years of experience with Oracle and SQL Server. She specializes in database performance-tuning and the glories of simplifying database management. In recent years, she has been working with large SSD databases and Exadata, taking advantage of the newest technology and performance-tuning features. Kellyn is known for her technical blog, <http://dbakevlar.com>, and contributions to numerous technical sites. She is on the board of directors for RMOUG (Rocky Mountain Oracle User Group) as the group's Training Days Director, a yearly conference of almost 1000 attendees. She has performed numerous presentations in the US and abroad at both Oracle and SQLPass User Group conferences. Kellyn lives with her partner, Tim, and her three children in Broomfield, Colorado (USA). Currently, she is honing her technical skills at Enkitec (<http://www.enkitec.com>) and can be followed on twitter under the tweet handle @DBAkevlar.



■ **Mladen Prajdić** is a SQL Server MVP from Slovenia. He started programming in 1999 in Visual C++. Since 2002, he has been actively developing different types of applications in .Net (C#) and SQL Server, ranging from standard line-of-business to image-processing applications.

He graduated from the College of Electrical Engineering at the University of Ljubljana, majoring in Medical Cybernetics. He's a regular speaker at various conferences and user-group meetings. He blogs at <http://weblogs.sqlteam.com/mladenp> and has authored various articles about SQL Server. He really likes to optimize slow SQL statements, analyze performance, and find unconventional solutions to difficult SQL Server problems.

In his free time, among other things, he also develops a very popular add-in for SQL Server Management Studio, called SSMS Tools Pack.



■ **Herve Roggero**, Windows Azure MVP, is the founder of Blue Syntax Consulting, a company specializing in cloud-computing products and services. Herve's experience includes software development, architecture, database administration, and senior management, with both global corporations and startup companies. Herve holds multiple certifications, including an MCDBA, MCSE, and MCSA. He also holds a Master's degree in Business Administration from Indiana University. Herve is the co-author of *PRO SQL Azure* (Apress, 2010). For more information, visit <http://www.bluesyntax.net>.



■ **Chris Shaw** started his database career in 1993. He began by working with Lotus during his time in the Marine Corps. Following his time with the Marines, Chris continued working with databases for companies such as Wells Fargo, Pulte Mortgage, and Yellow Pages Inc. He later consulted with insurance companies, including Anthem Blue Shield and Admini Quest. Chris has been writing and speaking about SQL Server for over 15 years at events such as SQL Connections, PASS, and SSWUG Ultimate conferences. Chris was the Conference Director for SSWUG in 2008. Chris received the Microsoft MVP award in 2009, 2010, and 2011, and he just received his fourth MVP award in April 2012. He is the founding member of the Colorado Springs SQL Server User Group and presently is the Co-President. Chris has many featured articles on the SSWUG website, and he was a contributing

author for *SQL Server 2005 Bible* (Wiley, 2006) and *SQL Server MVP Deep Dives, Volume 2* (Manning Publications, 2011). Blogs, information, questions, quizzes, and interviews can be found at Chris' blog site at <http://chrisshaw.wordpress.com>. He can also be reached on twitter at @SQLShaw or by email at Chris@sqlshaw.com



■ **Gail Shaw** is a senior consultant with Xpertease, based in Johannesburg, South Africa. She specializes in database performance tuning and database recovery, with a particular focus on topics such as indexing strategies, execution plans, and writing T-SQL code that performs well and scales gracefully.

Gail is a Microsoft Certified Master for SQL Server 2008 and a SQL Server MVP. She is a frequent poster on the SQL Server Central forums, writes articles for both SQLServerCentral and Simple-Talk, and blogs on all things relating to database performance at <http://sqlinthewild.co.za>. She has spoken at TechEd Africa, at a number of the 24 Hours of PASS web events and, on multiple occasions, at the PASS Community Summit.

Gail is an Aikido Shodan (1st degree black belt) and a keen archer. In her copious spare time, she is pursuing a Master's degree in Computer Science at the University of South Africa.



■ **Jason Strate** is a database architect at Pragmatic Works, with more than 15 years of experience. He has been a recipient of Microsoft's "Most Valuable Professional" designation for SQL Server since July 2009. His experience includes the design and implementation of both OLTP and OLAP solutions, as well as the assessment and implementation of SQL Server environments for best practices, performance, and high-availability solutions.

Jason is an active member of the SQL Server community. He currently serves as regional mentor for the North Central PASS region, helping community members and chapters connect with each other. In the community, he presents on SQL Server and related topics at local, regional, and national events, including SQL Saturdays and the PASS Summit. He also blogs at <http://www.jasonstrate.com> and <http://www.sqlkaraoke.com>, and he can be contacted on twitter as @stratesql.

About the Technical Reviewers

■ **Rudi Bruchez** is an independent consultant and trainer based in Paris, France. He has 15 years' experience with SQL Server. He worked as a DBA for CNET Channel, a subsidiary of CNET.com, at the MSC (Mediterranean Shipping Company) headquarters in Geneva, and at Promovacances, an online travel company in Paris. Since 2006, he has provided consulting and audits as well as SQL Server training sessions. As SQL Server is evolving into a more complex solution, he is trying to make sure that developers and administrators keep mastering the fundamentals of the relational database and the SQL language. He co-authored one of the best-selling books about the SQL language in French, and in 2008, he published the only book in French about SQL Server optimization. He can be contacted at <http://www.babaluga.com/>.



■ **Robin Dewson** has been hooked on programming ever since he bought his first computer in 1980, a Sinclair ZX80. His first major application of his own was a Visual FoxPro program that could be used to run a Fantasy League system. At this point, he met up with a great help in his PC development life, Jon Silver, at Step One Technologies. In return for training, he helped Jon with some other Visual FoxPro applications. From there, realizing that the marketplace for Visual FoxPro was limited, he decided to learn Visual Basic and SQL Server.

Starting out with SQL Server 6.5, Robin soon moved to SQL Server 7 and Visual Basic 5 and became involved in developing several applications for clients both in the UK and in the United States. From there, he has moved through each version to SQL Server 2012. Robin now specializes in C# and VB .NET. Robin can be contacted at robin@fat-belly.com or at <http://www.fat-belly.com>.

■ **David Dye** is a Microsoft SQL Server MVP, instructor, and author specializing in relational database management systems, business intelligence systems, reporting solutions, and Microsoft SharePoint. For the past nine years, David's expertise has been focused on Microsoft SQL Server development and administration. His work earned him recognition as a Microsoft MVP in 2009 and 2010. He has been a moderator for the Microsoft Developer Network for SQL Server forums and was the Innovator of the Year runner-up in 2009 by *SQL Server Magazine*. He was also featured in the Training Associates Technical Trainer Spotlight in April 2011. David currently serves as a technical reviewer and co-author with Apress Publishing in the SQL Server 2012 series, and as an author with Packt Publishing.

■ **Edgar Lanting** is a certified database specialist from the Netherlands with over 18 years of experience in IT. He is currently working as a database consultant for Ciber, where he assists various companies with the support and management of their database environments running on Microsoft SQL Server and Oracle software.

In his spare time, he enjoys hikes and walks with his wife and dog, is a big animal and nature photography enthusiast, and enjoys an occasional snowboarding trip when wintertime has arrived.

■ **Ken Simmons** is a database administrator and developer specializing in Microsoft SQL Server. He is an author for multiple SQL Server websites and books, including *Pro SQL Server 2008 Administration* (Apress, 2009), *Pro SQL Server 2008 Mirroring* (Apress, 2009), *Pro SQL Server 2008 Policy-Based Management* (Apress, 2010), and *Pro SQL Server 2012 Administration* (Apress, 2012). He currently holds certifications for MCP, MCAD, MCDBA, MCTS for SQL Server 2005, and MCITP for SQL Server 2008.

Acknowledgments

At some point early in 2012, my friend Chris Shaw innocuously asked me if I would be interested in writing. Chris and I were preparing for a presentation for a SQL event. I really thought that the question was simple. I had some articles at various SQL Server websites under my belt and had even just finished a series of three articles. It turned into much, much more.

As I was struggling with my early versions of my chapter, Allen Hirt went above and beyond the call of duty, as he volunteered to edit my document in its early stages. As an established author, he was profoundly helpful. I learned so much from him.

As other authors were working on their chapters, several of them spurred me on as they tweeted about their tasks via Twitter, reminding me of my own tasks. Special thanks go to my fellow authors Gail, Mladen, Chris, Grant, Jes, and Jonathan for their examples and pushes.

For years, I have tried to be like my heroes in the #SQLFamily who have gone before me with writing books, articles, blogs, and more. Thanks to Steve, Tom, Grant, Tim, and Brad. Thanks to the editors of the various websites and other resources who kept telling me to write. Thanks to each website that has published an article of mine and allowed me to practice this craft.

Thanks to each coworker, DBA, developer, project manager, and so on who helped me perfect this process of Release Management to the point that I could put all these actions and lessons to paper to describe a process I have grown to love and hate.

Thanks to all the fine folks at Apress who helped me along the way of creating this chapter. It was an experience I appreciate, and I am grateful for their expertise and assistance.

Super special thanks go out to the one person who had no idea how much he was helping me by kick-starting all my extracurricular activities. He happened to say the right thing at the right time that got me going. I will forever be grateful to Pat Wright for pushing me out of my comfort zone.

–TJay Belt

Thanks to Jonathan Gennick for giving me yet another project to do right after I finished my book *Pro SQL Server 2012 Relational Database Design and Implementation*.

Thanks to all of my MVP friends, who provide me with a plethora of new ideas, excellent examples, and horror stories that remind me that reality is often worse than any example that you can think of.

Thanks to the PASS Community for all of the opportunities to speak and get better at my craft over the past few years. SQL Saturday events have increased the number of speaking opportunities by orders of magnitudes, and the more people I get to speak to, the more I learn what they want to hear about.

–Louis Davidson

First, thanks to my wife, Danielle, for continuing to support me on this journey and listening to all of my technical conversations. Thanks to Jeremy Lowell for thinking of me when this opportunity came up and to Allan Hirt for reviewing the content and being a great sounding board. Lastly, thanks to everyone at Apress for putting together a great group of authors to share their industry expertise with the community. The strength of the SQL Server community starts with people giving back and sharing.

–Ben DeBow

Several of the great people you see on the cover of this book were already committed to this project when I was asked to take part. Looking at that list of amazing people, realizing that I could be a small part of it, I jumped at the chance. Thanks, guys, for all the learning and motivation you supply to our extended SQL family. All the fine folks at Apress have done another excellent job. Special thanks to Jonathan Gennick of Apress. I really appreciate every chance he has given me. I work for a great company at Red Gate, and I appreciate all that it does for me. Thanks Gareth, Ben, Neil, and Simon for all you've done to make it a fantastic place to work for, learn, and grow. Finally, thanks to my long-suffering family to whom I said I'd never write another book until next time. Well, next time came rather quickly.

–Grant Fritchey

Thanks to Jonathan Gennick for giving me the opportunity to write a chapter in this book, and to Anamika—what a job :-)

Special thanks to Morten for giving good input and helping with translation. Thanks to my wife for keeping me on my toes, and to my best friend, who woke me up from being a dusty DBA.

And thank you to all the wonderful people at Miracle and the Danish SQL Server user group that makes life exciting.

–Jesper Johansen

My chapter was made possible by the best clients any consultant can hope for. They continually present me with unique and challenging opportunities. All of the folks at Apress who helped me in writing my first chapter, ever—their professionalism is outstanding. Feedback from Richard Rodriguez, Chris Shaw, and Ben DeBow were instrumental in making this chapter relevant. Of course, I'd be remiss to not include my wonderful wife, who endured proofreading this a few times, and my six-year-old daughter, who makes it her mission to put a smile on my face. Thank you!

–Jeremy Lowell

Special thanks to Chris Shaw (@SQLShaw) for inspiring me to write. I'm sure he doesn't realize it, but he has had a significant impact on my interest in writing and presenting. Thank you, Chris, for the opportunities you have opened up for me, my friend. I would also like to thank Allan Hirt (@SQLHA) for helping me pick up the pieces and for providing guidance on this path to authorship. He has been a friend, mentor, and teacher, and his time and experience has taught me many things. Finally, having the support of my family in my endeavors in the SQL community is a blessing I am honored to benefit from.

–Wendy Pastrick

I want to thank Tim Gorman for all his loving support as I embraced my first challenge as an author of published work. His emotional support, as someone who understands what I hope to accomplish and as my partner, is more than anyone could hope for in this life.

–Kellyn Pot'vin

I don't want to let this opportunity go without saying thank you to my #SQLFamily for sharing ideas and opinions.

–Chris Shaw

Introduction

Pro SQL Server 2012 Practices is an anthology of writings on practices and techniques you should become familiar with on the way to becoming a top-notch database administrator (DBA) or database programmer. The authors have deep insight, and they want to share. The book covers a very wide array of topics, such as what makes for a well-performing design, how to choose database server hardware, indexing, tuning, tools to help you build your database, techniques for managing replication, compliance, backups, high availability, and more.

Working as a SQL Server professional is daunting and exhilarating at the same time. It's daunting due to the immense amount of “mechanics” that one must learn simply to run and program against a SQL Server database at a basic level. Exhilaration comes at first from mastering the mechanical aspects of creating a database, creating tables, adding users, storing data, retrieving data, and so forth. Then, in the long run, you get past the basic tasks and activities and begin to gain deep insight into SQL Server and how it operates. One day, you realize you actually do know a few things. You want to share. And that's how this book came about.

Could you come up with these practices on your own? Yes, given enough time and experience you would likely learn everything for yourself, from first principles. The authors' goal in this book, however, is to speed you on your way—not so that you can cheat the natural order of things, but so that you can continue from where they left off and go on to develop even greater insights, and then share what you have learned with those who come after.

Audience

Pro SQL Server 2012 Practices is written for working SQL Server professionals. The book is primarily aimed at DBAs, but it is also of interest to developers who are serious about building their reputation and expertise around Microsoft SQL Server. The book targets intermediate users, presuming a baseline of experience with the underlying product. Its goal is to provide the sort of deep expertise and guidance that lets you stand on the shoulders of giants—so to speak—as you build your own career and reputation in the field.

Authors

The author team on this book is a group of dedicated professionals who are passionate about using SQL Server, using it well, and sharing their expertise with those around them. Each of the authors has chosen a chapter in the belief that what you learn will make a significant and positive contribution to your growth as a SQL Server professional.

Topics covered in this book, listed by chapter, include the following:

1. Be Your Developer's Best Friend, by Jesper Johansen
2. Getting It Right: Designing for Database Performance, by Louis Davidson
3. Hidden Performance Gotchas, by Gail Shaw
4. Dynamic Management Views, by Kellyn Pot'vin
5. From SQL Trace to Extended Events, by Mladen Prajdić
6. The Utility Database, by Chris Shaw
7. Indexing Outside the Bubble, by Jason Strate
8. Release Management, by TJay Belt
9. Compliance and Auditing, by Jonathan Gardner
10. Automating Administration, by Jes Borland
11. The Fluid Dynamics of SQL Server Data Movement, by Wendy Pastrick
12. Windows Azure SQL Database for DBAs, by Herve Roggero
13. I/O: The Untold Story, by Jeremy Lowell
14. Page and Row Compression, by Bradley Ball
15. Selecting and Sizing the Server, by Glenn Berry
16. Backups and Restores Using Availability Groups, by Grant Fritchey
17. Big Data for the SQL Server DBA, by Carlos Bossy
18. Tuning for Peak Load, by Ben DeBow

Example Code

Downloadable example code is available for some chapters in the book. You'll find the code from the book's catalog page on the Apress.com website:

<http://www.apress.com/9781430247708>

Click on the Source Code/Downloads tab. You will see a link from which to download the example code archive.

CHAPTER 1



Be Your Developer's Best Friend

By Jesper Johansen

When asked to write a chapter for this book, I was in doubt about which subject to choose. Should I write about common problems I normally see? About how Microsoft SQL Server's easy installation is at the same time a strength and a weakness? About how to fix the 10 most general problems with SQL Server installation by querying a couple of dynamic management views (DMVs)? Or how about the *Yet Another Performance Profiling* (YAPP) method—a well-known performance method in the Oracle community that is just as usable in SQL Server with the implementation of Extended Events and DMVs that will show you what you are waiting for?

No. What really makes me tick is becoming friends with the developers by creating a good SQL Server environment and fostering an understanding of one another's differences. Just think what can be accomplished if both sides can live peacefully together instead of fighting every opposite opinion, digging the trenches even deeper and wider. Through fostering good relationships, I have seen environments move from decentralized development systems and standalone production databases to central solutions with easy access for developers, and calm uninterrupted nights for me. However, this does mean that you have to give up some sovereignty over your systems by relinquishing some of your admin power.

The main problem is focus. While the DBA thinks of space issues, data modeling, and the stability of everyday operations, the developers think of making things smarter, sexier, and shiny. To make the relationship work, we have to move through the entire palette of the system—standardization, accessibility, logging, information flow, performance information—all while ensuring that systems are stable and that developers know that they are not alone, and that DBAs still exist and decide things.

My Experience Working with SQL Server and Developers

After finishing my engineering studies, I started as a developer. I worked on CRM systems in DataEase under DOS and OS/2, and that combination gave me plenty of insight into the issues developers have with the DBA function. DataEase was the Microsoft Access of that time, and it had connectors to all the major databases (Oracle, SQL Server, and DB2). But most of the time, DBAs would not allow dynamic access to production data. Their resistance led to friction with the developers.

By coincidence, I ended up as a Microsoft Visual Basic programmer in a company developing and running systems for all the Danish municipalities. I was placed among the DB2/MVS DBAs, and I was by far the youngest (and only) GUI person (OS/2 and Windows). While I coded Visual Basic 3 applications, those DBAs were taking care of decentralized connections, such as ODBC on DB/2 MVS. These were the

days before having TCP/IP on the mainframe, so we're talking Systems Network Architecture (SNA) and IBM Communications Manager.

One day, my boss gave me responsibility for a new product called SQL Server. Why? Because I was the only one working with Windows.

My biggest problem was how to approach the current environment within the company. How many SQL Server databases did we already have? Which development groups were using it? Those were just some of the questions I had to grapple with.

I had to start from scratch. So I asked my DB/2 colleagues for help. After all, they had been working in these kind of environments for the last 15 years, handling systems with 15,000 concurrent users, 50,000 different programs, thousands of individual databases, and lots of data on every Danish citizen, such as taxes, pension funds, and other personal information. I wanted the benefit of their experience.

What I learned was that data modeling is a must. You need to have a naming standard for servers, for database objects, for procedures—for everything, actually. Starting the battle for naming standards and consistency took me on a six-year-long journey with developers, until most developers actually started to live safely. They came to understand that my requirements gave them more stable and consistent environments to develop on, made them more efficient, and got the job done faster for all.

Reconciling Different Viewpoints Within an Organization

The everyday battles between DBAs and developers mostly concern routine administrative tasks. Limitations on space allocations and limits to changes in production are perceived by developers as inhibiting innovation and stopping them from making a difference. They often see the DBA as someone who purposely limits them from doing their job. On the other hand, the admin group thinks that developers rarely plan ahead longer than the next drop-down box or the next screen, and that they never think in terms of the time period over which the software they build must run, which is often five to ten years or even longer.

The consequences of these differences are that developers create their own secret systems, move budget money out of reach of the DBA team, and generally do everything in their power to limit the admins in setting up the imaginary borders they believe are being set up. For example, often I would hear the sentence, "If you take away that privilege from me, I can no longer boot the machine at will." The problem with that thinking is that well-configured SQL Server systems need no more restarts than any other type of systems.

So how do we get out of this evil spiral, and what are the benefits of doing so? Dialog is the way out, and the benefits are a controlled production environment, clearly defined ownership of databases, consistent environments patched correctly, lower cost of maintenance, possible license savings, and almost certainly fewer calls at 4:00 in the morning interrupting your sleep.

Remember, all change starts with one's self, and it is far easier to change yourself than to change others. So get a hold of a piece of paper, divide it into a plus and a minus side, and start listing the good and bad things in your environment. For instance, it could be a plus that some developers have sysadmin privileges because they fix some things for themselves, but it could also be a minus because they meddle with things they are not supposed to meddle with and create objects without the knowledge of the DBA.

What you'll get from this chapter is my experience and how I managed to foster a productive and good relationship with developers. I'll provide a couple of code examples to help you on the way to success, or to just inspire you. My approach is not the only way to achieve good developer relations, but it has proven effective in the environments in which I've worked.

Preparing to Work with Developers to Implement Changes to a System

To make progress, you have to prepare for it. Implementing change will not work if you make demands of the developers without preparing. The battle will be hard, but it will be worth fighting for, because in the end you'll be eating cake with the developers while talking about the bad-old days with their unstable systems, anarchy, and crashes without backups.

Bring some good suggestions to the table. Do not approach developers without having anything to offer to make *their* lives easier. Think of yourself as a salesperson of stability and consistency—not even developers will disagree with those goals. As in any good marriage, however, the needs of both parties must be aired and acknowledged.

Put yourself in their place as well. Try to understand their work. You'll find that most of their requests are actually not that bad. For example, a common request is to be able to duplicate the production environment in the development environment on the weekend to test new ideas in their software. Would you rather spend your weekend doing that work for them? Isn't it preferable to facilitate having them do the work on their own so that you can be home with your family?

Listen to your developers. Ask them what they see as the biggest hurdles put in their way by the admin group. Your goal, ultimately, should be to create an environment that is good for the business. That means making everybody as happy as possible, easing the bureaucracy, and ensuring stable access for all.

A well-prepared environment can also lead to server consolidation, which in turn leads to saving power, simplifying patching, and ultimately less administrative effort. The money saved from having well-prepared environments can then be used for better things, such as buying Enterprise Edition or enabling Always On availability groups to provide an environment more resilient to failure.

By now, you are beginning to think that I am all talk. How can you get this process of working closely with developers started? The answer depends on how your business is structured. Following is a list of steps. Don't start by doing everything at once. Keep it simple. Build on each success as you go along. Remember that the goal is to create a stable environment for all:

1. Make a map of your existing environment.
2. Create a description of what your new environment should look like.
3. Document your description in written form so that it is clear and convenient to pass on to management and vendors.
4. Create system-management procedures for most changes.
5. Create system reports to report on all those changes.

These are a good series of steps to follow. Don't be too rigid, though. Sometimes you will need to divert from this sequence to make your business work. Adapt and do what is right for your business.

Step 1: Map Your Environment

If you have never experienced discovering an instance or server that has existed without your knowledge and without your knowing who owns it or uses it, you are a member of the 0.1 percent of SQL Server DBA in the world that has it easy. Indeed, not only is it common to run across unknown servers and instances, sometimes you'll find one and not even know what applications, if any, are running on it. Thus, Step 1 is to begin mapping your infrastructure so that you can know what you currently have.

Start by finding all the servers in your domain. Several free tools are available on the Internet to help you do this. Or maybe you already have the information in your Configuration Management Database (CMD) but you have never created reports on that data.

Try executing the following in a command prompt window:

```
SQLCMD -L
```

This command will list all the available servers on your network that are visible. You can get much more detailed information using tools such as SQLPING, Microsoft Map, Quest Discovery Wizard, or other similar products. A benefit of these products is that they often provide information like version numbers or patch levels.

Once you find your servers, you need to find out whether they are actually still in use. Most likely, you will have servers that were used only during an upgrade, but no one thought to shut them down once the upgrade was complete. One place where I have seen this go horribly wrong was in an organization that forgot the old server was still running, so it no longer got patched. Along came the SLAMMER virus, and down went the internal network. Another project I was on, involved consolidating about 200 servers. We found we could actually just switch off 25 percent of them because they were not being used.

Following is a piece of code to help you capture information about logins so that you can begin to identify who or what applications are using a given instance. The code is simple, using the sysprocesses view available on most versions of SQL Server. Why not use audit trace? Because audit trace takes up a lot of space. You need only unique logins, and viewing logs of all login attempts from audit trace is not easy on the eyes.

First, create the following small table in the msdb database. I use msdb because it is available in all versions of SQL Server. The table will record unique logins.

```
CREATE TABLE msdb.dbo.user_access_log
( id          int IDENTITY(1,1) NOT NULL,
  dbname      nvarchar(128)    NULL,
  dbuser      nvarchar(128)    NULL,
  hostname    nchar(128)       NOT NULL,
  program_name nchar(128)       NOT NULL,
  nt_domain   nchar(128)       NOT NULL,
  nt_username nchar(128)       NOT NULL,
  net_address nchar(12)        NOT NULL,
  logdate     datetime         NOT NULL
CONSTRAINT DF_user_access_log_logdate DEFAULT (getdate()),
CONSTRAINT PK_user_access_log PRIMARY KEY CLUSTERED (id ASC) )
```

Then run the following code to sample logins every 15 seconds. If you need smaller or larger granularity, you can easily just change the WAITFOR part of the code. You can even make the code into a job that automatically starts when the SQL Server Agent starts.

```
WHILE 1=1
BEGIN

    WAITFOR DELAY '00:00:15';

    INSERT INTO msdb.dbo.user_access_log
    ( dbname,
      dbuser,
      hostname,
      program_name,
      nt_domain,
```

```

        nt_username,
        net_address )
SELECT distinct
    DB_NAME(dbid) as dbname,
    SUSER_SNAME(sid) as dbuser,
    hostname,
    program_name,
    nt_domain,
    nt_username,
    net_address
FROM master.dbo.sysprocesses a
WHERE spid>50
    AND NOT EXISTS( SELECT 1
        FROM msdb.dbo.user_access_log b
        WHERE b.dbname = db_name(a.dbid)
            AND NULLIF(b.dbuser,SUSER_SNAME(a.sid)) IS NULL
            AND b.hostname = a.hostname
            AND b.program_name = a.program_name
            AND b.nt_domain = a.nt_domain
            AND b.nt_username = a.nt_username
            AND b.net_address = a.net_address )
END

```

When you begin this process of capturing and reviewing logins, you should create a small team consisting of a couple of DBAs and a couple of the more well-liked developers and system owners. The reason to include others, of course, is to create ambassadors who can explain the new setup to other developers and system owners. Being told something by your peers makes it that much harder to resist the changes or even refuse them. And these people also have a lot of knowledge about the business, how the different systems interact, and what requirements most developers have. They can tell you what would be a show-stopper, and catching those in the beginning of the process is important.

Step 2: Describe the New Environment

The next step is to describe your new environment or the framework in which you plan to operate. What should this description contain? Make sure to address at least the following items:

- **Version of SQL Server** The fewer SQL Server versions you have to support, the easier it is to keep systems up and running. You will have fewer requirements to run different versions of Windows Server, and the newer the version you keep, the easier it is to get support from Microsoft or the community at large. I have seen several shops running versions spanning the gamut from 7.0, 2000, 2005, 2008, 2008 R2 through to 2012. Why not just choose 2012? Having the latest version as the standard is also a good selling point to developers, because most of the exciting new features will then be available to them. You might not be able to get down to just one version, but work hard to minimize the number you are stuck with supporting.
- **Feature support** Get started studying all the different features. Describe how your environment is implemented, how it is used, what features are accepted, and which features are not. Take into account whether features require privileges such as sysadmin, access to the command shell, and so forth. The important thing in this process is to understand the advantages and disadvantages of every feature, and to