Pro SQL Database for Windows Azure

SQL Server in the Cloud
Second Edition

Scott Klein
Herve Roggero

Apress®

Pro SQL Database for Windows Azure: SQL Server in the Cloud

Copyright © 2012 by Scott Klein and Herve Roggero

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 978-1-4302-4395-3

ISBN 978-1-4302-4396-0 (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 images of the Android Robot (01 / Android Robot) are reproduced from work created and shared by Google and used according to terms described in the Creative Commons 3.0 Attribution License. Android and all Android- and Google-based marks are trademarks or registered trademarks of Google, Inc., in the U.S. and other countries. Apress Media, L.L.C. is not affiliated with Google, Inc., and this book was written without endorsement from Google, Inc.

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

Technical Reviewer: Thomas LaRock

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: Kevin Shea

Copy Editor: James Compton

Compositor: SPi Global

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/9781430246831. For detailed information about how to locate your book's source code, go to www.apress.com/source-code.

*To my wonderful wife, children, and family. Without them,
I'd be nothing.*-Scott To my dear wife, Frederique.

—Herve

Contents at a Glance

Foreword	xvii
About the Authors.....	xix
About the Technical Reviewer	xxi
Acknowledgments	xxiii
Introduction	xxv
■ Chapter 1: Getting Started with SQL Database	1
■ Chapter 2: Design Considerations	23
■ Chapter 3: Security.....	45
■ Chapter 4: Data Migration and Backup Strategies	67
■ Chapter 5: Programming with SQL Database	99
■ Chapter 6: SQL Reporting	125
■ Chapter 7: SQL Data Sync.....	143
■ Chapter 8: Windows Azure and ASP.NET	165
■ Chapter 9: Designing for High Performance	183
■ Chapter 10: Federations	207
■ Chapter 11: Performance Tuning	219
■ Chapter 12: Windows Azure Mobile Services.....	241
■ Appendix A: SQL Database Management Portal	257
■ Appendix B: Windows Azure SQL Database Quick Reference.....	275
Index.....	283

Contents

Foreword	xvii
About the Authors.....	xix
About the Technical Reviewer	xxi
Acknowledgments	xxiii
Introduction	xxv
■ Chapter 1: Getting Started with SQL Database	1
Introduction to Cloud Computing.....	1
Who Is Doing What in the Cloud?.....	2
Typical Cloud Services.....	2
Discovering the Microsoft Azure Platform.....	3
Why Microsoft Azure?.....	3
About Geographic Locations.....	4
Storing Data in Azure.....	5
SQL Database Primer	6
Registering for Azure	6
Creating a SQL Database Instance	7
Configuring the Firewall	10
Connecting with SQL Server Management Studio.....	11
Creating Logins and Users.....	15
Assigning Access Rights.....	18
Understanding Billing for SQL Database.....	18

Limitations in SQL Database	19
Security	20
Backups.....	20
Objects.....	21
Miscellaneous.....	21
Drivers and Protocols	22
Summary.....	22
■ Chapter 2: Design Considerations	23
Design Factors	23
Offsite Storage.....	23
High Availability	23
Performance	24
Data Synchronization.....	27
Direct vs. Serviced Connections.....	27
Pricing	28
Security	29
Review of Design Factors	29
Design Patterns.....	29
Direct Connection	29
Smart Branching.....	30
Transparent Branching	31
Sharding	31
Offloading	35
Aggregation	36
Mirroring.....	37
Combining Patterns.....	37
Transparent Branching + RWS	37
Cascading Aggregation.....	38
Sample Design: Application SLA Monitoring	39
Pre-Azure Application Architecture.....	39
Azure Implementation	40

Other Considerations	41
Blob Data Stores	42
Edge Data Caching	42
Data Encryption	43
SaaS Applications and Federations	43
Summary	43
■ Chapter 3: Security	45
Overview	45
Confidentiality	45
Integrity	46
Availability	46
Securing Your Data	48
Encryption	48
Hashing	50
Certificates	55
Access Control	59
Authentication (AUTHN)	59
Authorization (AUTHZ)	60
SQL Database Firewall	65
Internal Firewalls	65
Compliance	65
Summary	66
■ Chapter 4: Data Migration and Backup Strategies	67
Migrating Databases and Data to SQL Azure	67
The Import/Export Service	68
Generate and Publish Scripts Wizard	83
The bcp Utility	92

SQL Azure Backup Strategies.....	95
Copying a Database.....	95
Backing Up Using the Import/Export Features.....	97
Third-Party Backup Products.....	97
Summary.....	98
■ Chapter 5: Programming with SQL Database	99
Application Deployment Factors.....	99
On-Premises Application	100
Azure-Hosted Application	100
Which to Choose?	101
Connecting to SQL Database.....	101
ADO.NET.....	102
ODBC.....	107
sqlcmd.....	109
WCF Data Services	114
Creating a Data Service	114
Connecting the Service to the Model.....	115
Creating the Client Application	117
Best Practices	119
Transient Fault Handling Application Block	120
Using the Transient Fault Handling Application Block	121
Summary.....	123
■ Chapter 6: SQL Reporting	125
SQL Reporting Overview.....	125
Architecture	126
Feature Comparison	127
Provisioning Your SQL Reporting Server	128
Creating a Report	130
Creating the SQL Database Data Source	131
Creating the Report Design.....	136

Deploying the Report.....	137
Security	139
Roles.....	139
Using the Management Portal	141
Pricing	142
Summary.....	142
■ Chapter 7: SQL Data Sync.....	143
Understanding SQL Data Sync	143
Why the Need?	143
The Basic Scenario	144
Common Data Sync Scenarios	145
Architecture	145
Configuring Synchronization	147
Provision a SQL Data Sync Server	147
Creating a Sync Group	149
Defining the Hub and Member Databases	150
Selecting Tables to be Synchronized	156
Deploying the Sync Group	159
Debugging and the Log Viewer.....	160
Looking at the Synchronized Data.....	161
Editing Data and Resynchronizing.....	162
Data Sync Limitations	163
Data Sync Best Practices	163
Design Considerations.....	163
Initial Synchronization	164
Security	164
Sync Schedule.....	164
Summary.....	164

■ Chapter 8: Windows Azure and ASP.NET	165
Creating a Cloud Service	165
Creating a Windows Azure Project	168
Configuring Your Development Environment	168
Creating Your First Visual Studio Cloud Project	169
Connecting a GridView to SQL Database	173
Deployment and Configuration Files.....	176
Deploying an ASP.NET Application in Windows Azure.....	179
Summary.....	181
■ Chapter 9: Designing for High Performance	183
General Performance Concepts.....	183
Chatty vs. Chunky	183
Lazy Loading.....	183
Caching.....	184
Asynchronous User Interface.....	184
Parallel Processing	185
Shards	185
Coding Strategies Summary	185
Building a Shard	186
Designing the Shard Library Object.....	187
Managing Database Connections	188
Reading Using the Shard	190
Caching.....	193
Updating and Deleting Records in the Shard.....	194
Adding Records to the Shard	196
Managing a Shard	198
Managing Exceptions	198
Managing Performance	199
Working with Partial Shards	202
Managing Transaction Consistency	203
Managing Foreign Key Constraints.....	203

Designing a Multitenant System	204
Creating Vertical Partition Shards	206
Big Data with Hadoop for Windows Azure	206
Summary	206
■ Chapter 10: Federations	207
Introducing Federations	207
Federations versus Sharding	207
Why Use Federations?	208
Federations Overview	208
Creating More Federations	213
Managing Federations.....	214
Advanced Considerations.....	216
Limitations.....	216
Sharding Library	217
Summary	218
■ Chapter 11: Performance Tuning.....	219
What's Different with SQL Database	219
Methods and Tools.....	219
Coding Implications	220
Tuning Techniques.....	220
Dynamic Management Views	220
Connection Pooling.....	223
Execution Plans with SSMS.....	224
Execution Plans with the Management Portal.....	228
Query Performance with the Management Portal	230
Indexing.....	231
Indexed Views.....	235
Stored Procedures	236

Provider Statistics	238
Application Design	239
Summary	240
■ Chapter 12: Windows Azure Mobile Services	241
Getting Started	241
Data	248
Browse.....	248
Columns.....	249
Permissions	249
Script	251
Advanced Settings.....	251
Push.....	251
Identity.....	252
Scale.....	253
Summary	255
■ Appendix A: SQL Database Management Portal	257
Launching the Management Portal	257
Administration Features	259
Run T-SQL Statements.....	259
View Execution Plans.....	261
Monitor Current Query Performance	262
Design Features	263
Designing Tables.....	263
Designing Views	270
Designing Stored Procedures	271
Summary	273

■ **Appendix B: Windows Azure SQL Database Quick Reference.....275**

 Supported T-SQL Statements275

 Partially Supported T-SQL Statements276

 Unsupported T-SQL Statements280

 Supported Data Types280

Index.....283

Foreword

My journey from box software engineer to cloud service engineer began over four years ago when I moved into the SQL cloud team at Microsoft. I would like to say it was my brilliant foresight that led me to make the jump, but the truth is it was a fortuitous accident resulting from a group reorganization. I couldn't be more thankful for the opportunity, as working on Windows SQL Azure has been the most interesting and exciting project of my career. The journey to build a relational database as a service was not a straightforward one. We tried several different incarnations and had to back out from a couple of dead ends before we landed on the service you see today.

We also went through several iterations on the value proposition. What are we building? What is it good for? How does it help customers? There were a lot of naysayers who said a SQL relational database had no place in the cloud. Some thought cloud was just a fad and not really any different from what people already did. We believed differently. We knew you could have the benefits of a relational database system (a well understood model and programming APIs, and an existing rich tooling ecosystem) and combine them with the best of a cloud service (high availability, easy and fast provisioning, drastically reduced management, and a pay as you go model).

Clearly it is resonating with customers. We are seeing adoption from all sorts of usage scenarios. Some are well known scenarios of simple web sites, departmental enterprise applications. Others are net new scenarios like the one we call *data hub*, which is about sharing between islands of data, such as a company and its outlet sites, a company and its vendors/suppliers or between multiple companies that are collaborating. Our most common scenario is a mixture of old and new, the SaaS Line of Business application. This includes cloud-only companies as well as existing box software vendors who are moving their software to the cloud. Moving to an SaaS is a huge advantage for these companies. It allows them to streamline the sales cycle, focus their resources on their core competency, and easily extend their global market reach.

However, running a cloud service is not the same as building box software. While the business logic of the applications is pretty straightforward to get up and running, especially if you have existing code from an on-prem application; keeping it running well in the cloud is not easy. Security is something that goes beyond compliant algorithms and best practices for your customer. You are on the hook to prevent, or least identify and stop, malicious attacks on your system. You have to make sure you isolate one tenant from another. Availability and performance are things customers take for granted and they are your responsibility. Where there is a problem, you get the call at that wakes you out of bed. Troubleshooting is much harder. Root-causing an issue on a single SQL database can be difficult. How do you do it when you have 1000 SQL databases? There are also model shifts to absorb, like embracing failure and building compensation into your system. One of the hardest ones SQL Database customers struggle with is the notion that they have to scale out instead of scaling up. This is hard because scale up was the preferred pattern in the box world and scaling out is a lot harder. There is also a lot of new stuff to deal with, like multi-tenancy. How do I host 1000 customers? Which parts say isolated and which parts are shared? How do I do billing? How do I measure COGS (and make sure they are low enough!).

These are not trivial problems. In order to tackle them efficiently you need to understand the capabilities of the platform you are running on. Which of the problems can it can help you to solve and for which ones are you on your own? Authors Scott Klein and Herve Roggero have done a great job walking you through the ins and outs of Windows Azure SQL Database. From Programming Model to Migration Strategies, Security to Scale out. Whether you are focused on using the latest high-level interfaces like WCF and OData or diving into the details of performance, tuning the book covers everything you need to know about SQL Database.

Rick Negrin
Lead Program Manager
Windows Azure SQL Database

About the Authors

Scott Klein is a Windows Azure Technical Evangelist for Microsoft focusing on Windows Azure SQL Database (formally known as SQL Azure) and related cloud-ready data services. He started his career in SQL Server with version 4.2 and has worked with every version since then during his 20+ year career. Prior to joining Microsoft he was a SQL Server MVP and then became one of the first four SQL Azure MVPs. Scott is the author or coauthor of several books for both Wrox and Apress, including this one. Scott can be found talking about Windows Azure SQL Database wherever he can get people to listen to him. You can reach Scott at SQLScott@live.com and read some of his musings at <http://www.scotttklein.com>.

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 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 Database for Windows Azure* from Apress. For more information, visit www.bluesyntax.net.

About the Technical Reviewer

Thomas LaRock is a seasoned IT professional with over a decade of technical and management experience. Currently serving as a Technical Evangelist for Confio Software, Thomas has progressed through several roles in his career, including programmer, analyst, and DBA. Thomas holds an MS degree in Mathematics from Washington State University and is a member of the Usability Professionals' Association. Thomas currently serves on the Board of Directors for the Professional Association for SQL Server (PASS), and is also a SQL Server MVP. Thomas can also be found blogging at <http://thomaslarock.com> and is the author of *DBA Survivor: Become a Rock Star DBA* (<http://dbasurvivor.com>).

Acknowledgments

This book exists because of the diligence and patience of a handful of individuals to whom we are extremely grateful. First, to Jonathan Gennick at Apress for letting us do a second edition and being a wonderful sounding board for ideas and thoughts regarding this book. Second, to Kevin Shea of Apress and Chris Nelson for keeping us on track, reviewing our work, correcting our grammar, and making the writing of this book quite a delightful process.

A very special thanks to Tom LaRock, the technical editor, for his meticulous and detailed work ensuring that the our examples actually work and that the message we are sharing is clear and thorough for the readers. If you ever have an opportunity to meet Tom LaRock, be sure to thank him, and take the opportunity to talk SQL Server with him. You will be glad you did.

We can't thank the fine individuals at Microsoft enough for their insight and feedback on many of these chapters. So, a huge thank you to Rick Negin, Mike Morrison, Ariel Netz, and Barclay Hill. It is a pleasure to associate with individuals who have a passion in what they do. We also thank Rick Negrin again for contributing the Foreword for this book.

Nothing in this life is worth doing without the love and support of family. Thus, Scott would like to profoundly express thanks to his wife, Lynelle, and his children for their endless love, support, understanding, and LOADS of patience. Professionally, Scott would also like to thank a few of his co-workers for their thoughts, insight, and advice. Cory Fowler, Brady Gaster, and Wenming Yi, thank you!

Introduction

Windows Azure SQL Database, formally known as SQL Azure, appeared on the scene about five years ago. At the time, little was known about it, but Microsoft was beginning to talk quite a bit about the Azure platform. Most people thought that SQL Azure was another NoSQL offering, while in reality it was, and is, nothing of the sort. At that time, the largest database it could handle was 1GB, and no one was really taking it seriously. Since that time, Windows Azure SQL Database has grown into an enterprise-ready PaaS (Platform as a Service) offering based on the proven SQL Server technology.

Cloud computing is not hype anymore. Today, cloud-based solutions are becoming the norm rather than an afterthought or sitting on the fringe. The benefits of the Windows Azure cloud platform, including Windows Azure SQL Database, allow businesses to rapidly create and scale solutions with low acquisition costs, yet provide high availability and interoperability. SQL developers and DBAs can use existing skills and knowledge to extend their on-premises solutions and quicken cloud development time. This book covers the fundamental Windows Azure SQL Database concepts, practices, and approaches your valuable data needs as it prepares for the journey to the cloud and Windows Azure SQL Database.

Because of the rapid pace at which Windows Azure SQL Database is updated, some of the services discussed in this book are in Preview as we write and may change somewhat by the time you read this. However, we have tried our best to bring you the most up-to-date information. Updated information can be found at our blogs, the Windows Azure blog (<http://blogs.msdn.com/b/windowsazure/>), and the all-important Windows Azure home page (<http://www.windowsazure.com/>) where you can find features, pricing, developer information, and much more.

Our hope is that after reading this book you have a better understanding of, and appreciation for, Windows Azure SQL Database. Whether you are just getting started with it or are a “seasoned veteran,” each chapter contains scenarios and information that we hope you will find helpful and beneficial as you design and build Windows Azure projects. There also is a plethora of source code that is used in chapters where examples are given.

Who This Book Is For

Pro SQL Database for Windows Azure, 2nd Edition is aimed at developers and database administrators who want instant access to a fully-capable SQL Server database environment without the pain of sorting out and managing the physical infrastructure.

How This Book Is Structured

Pro SQL Database for Windows Azure is designed to take you from knowing almost nothing at all about SQL Database to being able to configure and deploy it for use by production applications. The book does assume a knowledge of databases in general, and of SQL in general. From that base, the book takes you from the point of getting started through performance tuning and the use of other Azure Data Services.

Chapters in the book are as follows:

Chapter 1, *Getting Started with SQL Database*, helps you create your first database in the cloud.

Chapter 2, *Design Considerations*, discusses design issues you should think about when creating an application to run against a database that is cloud-based instead of one hosted in your own data center.

Chapter 3, *Security*, covers the all-important issue of securing your data in a scenario in which it is accessed across the public Internet.

Chapter 4, *Data Migration and Backup Strategies*, helps you move data efficiently into and out of SQL Database. It also covers backup strategies to protect yourself in the event of your database being lost or damaged.

Chapter 5, *Programming with SQL Database*, covers the differences between using SQL Database from on-premises applications and from Azure-hosted applications.

Chapter 6, *SQL Reporting*, shows how to create cloud-based reports.

Chapter 7, *SQL Data Sync*, Covers replication between multiple SQL Database instances, and between SQL Database in the cloud and SQL Server within your data center.

Chapter 8, *Windows Azure and ASP.NET*, provides an example and guidance for building ASP.NET applications backed by Windows Azure and SQL Database.

Chapter 9, *Designing for High Performance*, covers topics and features such as sharding, lazy loading, caching, and others that are important in building highly performant applications.

Chapter 10, *Federations*, discusses how to achieve greater scalability and performance from the database tier through the horizontal partitioning of data.

Chapter 11, *Performance Tuning*, provides techniques to use in troubleshooting applications when they aren't executing fast enough to meet business requirements.

Chapter 12, *Windows Azure Mobile Services*, covers how to easily connect a scalable cloud-based data backend to your client and mobile applications.

Also in the book are two appendices. Appendix A provides a walk-through of the SQL Database Management Portal. Appendix B is a quick-reference to T-SQL syntax supported by SQL Database.

Conventions

Throughout the book, we've kept a consistent style for presenting SQL and results. Where a piece of code, a SQL reserved word or fragment of SQL is presented in the text, it is presented in fixed-width Courier font, such as this (working) example:

```
select * from dual;
```

Where we discuss the syntax and options of SQL commands, we've used a conversational style so you can quickly reach an understanding of a command or technique. This means we haven't duplicated large syntax diagrams that better suit a reference manual.

Downloading the Code

The code for the examples shown in this book is available on the Apress web site, www.apress.com. A link can be found on the book's information page under the Source Code/Downloads tab, located below the Related Titles section of the page.

Contacting the Authors

Should you have any questions or comments—or even spot a mistake you think we should know about—you can contact the authors at the following addresses: scottkl@microsoft.com (Scott Klein), and hroggero@BlueSyntax.onmicrosoft.com (Herve Roggero).