



Professional Expertise Distilled

# SQL Server 2014 Development Essentials

Design, implement, and deliver a successful database solution  
with Microsoft SQL Server 2014

Basit A. Masood-Al-Farooq

**[PACKT]** enterprise   
PUBLISHING professional expertise distilled

---

# SQL Server 2014 Development Essentials

Design, implement, and deliver a successful database solution with Microsoft SQL Server 2014

**Basit A. Masood-Al-Farooq**

**[PACKT]** enterprise   
PUBLISHING professional expertise distilled  
BIRMINGHAM - MUMBAI

---

# SQL Server 2014 Development Essentials

Copyright © 2014 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: July 2014

Production reference: 1180714

Published by Packt Publishing Ltd.  
Livery Place  
35 Livery Street  
Birmingham B3 2PB, UK.

ISBN 978-1-78217-255-0

[www.packtpub.com](http://www.packtpub.com)

Cover image by Angela Robledo ([angel.spo@gmail.com](mailto:angel.spo@gmail.com))

---

# Credits

**Author**

Basit A. Masood-Al-Farooq

**Project Coordinator**

Sageer Parkar

**Reviewers**

Basavaraj Biradar

Brenner Grudka Lira

David Loo

Richard Louie

José (Cheo) Redondo

**Proofreaders**

Simran Bhogal

Ameesha Green

**Indexers**

Rekha Nair

Tejal Soni

**Acquisition Editor**

Neha Nagwekar

**Graphics**

Ronak Dhruv

**Content Development Editor**

Neil Alexander

**Production Coordinators**

Kyle Albuquerque

Saiprasad Kadam

Conidon Miranda

**Technical Editor**

Pankaj Kadam

**Copy Editors**

Insiya Morbiwala

Sayanee Mukherjee

Aditya Nair

**Cover Work**

Kyle Albuquerque

---

# About the Author

**Basit A. Masood-Al-Farooq** is an internationally known Lead SQL DBA, trainer, and technical author with twelve years' experience of the Microsoft technology stack. He is an accomplished development and production SQL Server DBA with a proven record of delivering major projects on time and within budget. He is an expert at evaluating the clients' needs against the capabilities of the SQL Server product set, with the objective of minimizing costs and maximizing functions by making innovative use of advance capabilities. Basit has authored numerous SQL Server technical articles on various SQL Server topics for different SQL Server community sites, which include [SQLMag.com](http://SQLMag.com), [MSSQLTips.com](http://MSSQLTips.com), [SQLServerCentral.com](http://SQLServerCentral.com), [SSWUG.org](http://SSWUG.org), [SQL-Server-Performance.com](http://SQL-Server-Performance.com), and [SearchSQLServer.com](http://SearchSQLServer.com).

He has also developed and implemented many successful database infrastructures, data warehouses, and business intelligence projects. He holds a Master's degree in Computer Science from London Metropolitan University and industry-standard certifications from Microsoft, Sun, Cisco, Brainbench, ProSoft, and APM, which include MCITP Database Administrator 2008, MCITP Database Administrator 2005, MCDBA SQL Server 2000 and MCTS .NET Framework 2.0 Web Applications. He also has a good understanding of ITIL principles.

He can be reached via Twitter (@BasitAali), his blog (<http://basitaalishan.com>), or via LinkedIn (<http://uk.linkedin.com/in/basitfarooq>).

He was a technical reviewer for *SQL Server 2012 Reporting Services Blueprints*, *Marlon Ribunal and Mickey Stuewe*, Packt Publishing and *Reporting with Microsoft SQL Server 2012*, *James Serra and Bill Anton*, Packt Publishing.

---

# Acknowledgments

First and foremost, I would like to praise and thank Allah SWT, the compassionate, the almighty, the most merciful, who has granted me countless blessings, knowledge, and opportunities. Without the will of Allah SWT, none of this would be possible.

I would like to thank my parents for getting me started on my journey, giving me the opportunity for a great education, allowing me to realize my own potential, and giving me the freedom to choose my career path. Thanks Dad (Masood Ahmad Nisar) and Mom (Saeeda Perveen); you both have always supported me and encouraged me in everything I have ever done. You both have worked extremely hard all your life to give me the life I wanted.

I would like to thank my caring, loving, and supportive wife, Aniq, for all the encouragement and support. Despite the long days, sleepless nights, and long and exhausting marathons of writing, a few words of love and encouragement from you always successfully wipe away all my fatigue. Thank you for supporting me, coping with my hectic work schedule, and taking care of our kids and household activities, and giving me time so that I can pursue and concentrate on this book-writing project. I would also like to thank my two sons, Saifan and Rayyan, and would like to dedicate this book to them because they always make me smile and understood on those weekend mornings when I was writing this book instead of playing with them. I hope someday you both will read my book and understand why I spent so much time in front of my laptop.

I would also like to thank Packt Publishing for giving me the opportunity to write this book. Last but not least, I would like to thank Neha Nagwekar, the Acquisition Editor; Neil Alexander, the Content Development Editor; Pankaj Kadam, the Technical Editor; and all the technical reviewers, for their help, advice, and constructive comments, because without their feedback this book would not have been possible.

---

# About the Reviewers

**Basavaraj Biradar** holds a Master's degree in Computer Applications with gold medals from Gulbarga University, India. Besides these, he has a Microsoft Certified Professional certification in SQL Server 2000. Basavaraj has a rich experience of more than 13 years in designing and developing databases for complex, large, online systems. Currently, Basavaraj is working as Senior Technical Lead in a major IT security company.

Basavaraj writes blogs regularly at <http://sqlhints.com> on SQL Server technologies and his articles are quite popular in the industry. Basavaraj speaks about SQL Server in Microsoft User Group meetings and in his company as well. You may contact Basavaraj through his e-mail ID, [basav@sqlhints.com](mailto:basav@sqlhints.com).

**Brenner Grudka Lira** has been a data analyst and DBA at Neurotech since 2012. He has a Bachelor's degree in Computer Science and a postgraduate degree in Project Management, both from the Catholic University of Pernambuco in Recife, Brazil.

He also has experience in building and modeling data warehouses and has knowledge in SQL Server and MySQL database management. Today, he is dedicated to the study of project management and database tuning.

He has reviewed *Microsoft SQL Server 2012 Integration Services: An Expert Cookbook*, *Reza Rad and Pedro Perfeito*; *Oracle BI Publisher 11g: A Practical Guide to Enterprise Reporting*, *Daniela Bozdoc*; and *Getting Started with SQL Server 2014 Administration*, *Gethyn Ellis*, all by Packt Publishing.

---

**David Loo** is a senior software professional with over 25 years' experience in both software development and people management. He is respected for his ability to focus teams on service excellence and for designing and implementing practical process improvements and solutions. He is always on the lookout for ways to contribute his knowledge and experience of software development, team-building, and development best practices.

He has reviewed *Getting Started with SQL Server 2012 Cube Development*, Simon Lidberg and *Getting Started with SQL Server 2014 Administration*, Gethyn Ellis, both by Packt Publishing.

**Richard Louie** is a Senior Business Intelligence Developer with over 20 years' experience in software development and project management. He has extensive hands-on experience in Oracle and Microsoft SQL for ETL, SSIS, SSRS, SSAS, and VB.Net. Richard is a graduate of the University of California, Irvine in Information and Computer Science, and is ASQ Green Belt Certified.

He has reviewed *Getting Started with SQL Server 2012 Cube Development*, Simon Lidberg and *Getting Started with SQL Server 2014 Administration*, Gethyn Ellis, both by Packt Publishing.

**José (Cheo) Redondo** is a consultant, educator, mentor, and evangelist of technology databases and a SQL Server MVP in Latin America, specializing in enterprise databases and business intelligence solutions since 1998. Since that time, he has been providing consulting services and specialized education through academic and business-user groups. He gives conferences to the PASS Community (SQLSaturday events) in Latin America and the US in Spanish, and has been leading SQL PASS Venezuela for several years now. You can follow Cheo on Twitter at @redondoj or contact him through his blog, El Blog de Cheo Redondo in Spanish (<http://redondoj.wordpress.com/>), or his e-mail, [redondoj@gmail.com](mailto:redondoj@gmail.com).



---

# www.PacktPub.com

## Support files, eBooks, discount offers, and more

You might want to visit [www.PacktPub.com](http://www.PacktPub.com) for support files and downloads related to your book.

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at [www.PacktPub.com](http://www.PacktPub.com) and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at [service@packtpub.com](mailto:service@packtpub.com) for more details.

At [www.PacktPub.com](http://www.PacktPub.com), you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.



<http://PacktLib.PacktPub.com>

Do you need instant solutions to your IT questions? PacktLib is Packt's online digital book library. Here, you can access, read and search across Packt's entire library of books.

## Why subscribe?

- Fully searchable across every book published by Packt
- Copy and paste, print and bookmark content
- On demand and accessible via web browser

## Free access for Packt account holders

If you have an account with Packt at [www.PacktPub.com](http://www.PacktPub.com), you can use this to access PacktLib today and view nine entirely free books. Simply use your login credentials for immediate access.

## Instant updates on new Packt books

Get notified! Find out when new books are published by following [@PacktEnterprise](https://twitter.com/PacktEnterprise) on Twitter, or the *Packt Enterprise* Facebook page.

---

# Table of Contents

<b>Preface</b>	<b>1</b>
<b>Chapter 1: Microsoft SQL Server Database Design Principles</b>	<b>7</b>
<b>Database design</b>	<b>8</b>
The requirement collection and analysis phase	8
The conceptual design phase	9
The logical design phase	9
The physical design phase	10
The implementation and loading phase	10
The testing and evaluation phase	10
The database design life cycle recap	10
<b>Table design</b>	<b>11</b>
Tables	11
Entities	12
Attributes	12
<b>Relationships</b>	<b>12</b>
A one-to-one relationship	12
A one-to-many relationship	13
A many-to-many relationship	13
<b>Data integrity</b>	<b>14</b>
<b>The basics of data normalization</b>	<b>14</b>
The normal forms	15
The first normal form (1NF)	15
The second normal form (2NF)	15
The third normal form (3NF)	16
Denormalization	16
<b>The SQL Server database architecture</b>	<b>16</b>
Pages	17
Extents	18

The transaction log file architecture	19
The operation and workings of a transaction log	20
Filegroups	21
<b>The importance of choosing the appropriate data type</b>	<b>21</b>
SQL Server 2014 system data types	22
Alias data types	23
Creating and dropping alias data types with SSMS 2014	23
Creating and dropping alias data types using the Transact-SQL DDL statement	23
CLR user-defined types	24
<b>Summary</b>	<b>24</b>
<b>Chapter 2: Understanding DDL and DCL Statements in SQL Server</b>	<b>25</b>
<b>Understanding the DDL, DCL, and DML language elements</b>	<b>26</b>
Data Definition Language (DDL) statements	26
Data Manipulation Language (DML) statements	26
Data Control Language (DCL) statements	26
<b>Understanding the purpose of SQL Server 2014 system databases</b>	<b>27</b>
SQL Server 2014 system databases	27
The master database	27
The model database	28
The msdb database	28
The tempdb database	28
The resource database	29
The distribution database	29
<b>An overview of database recovery models</b>	<b>29</b>
The simple recovery model	29
The bulk-logged recovery model	29
Full recovery	30
<b>Creating and modifying databases</b>	<b>30</b>
Create, modify, and drop databases with T-SQL DDL statements	30
Creating a database with T-SQL DDL statements	30
Example 1 – creating a database based on a model database	32
Example 2 – creating a database that explicitly specifies the database data and the transaction log file's filespecs properties	32
Example 3 – creating a database on multiple filegroups	33
Modifying a database with T-SQL DDL statements	33
Example – adding a secondary data file to an existing database	34
Dropping a database with T-SQL DDL statements	35
Create, modify, and drop databases with SSMS 2014	35
Creating a database with SSMS 2014	35
Modifying a database with SSMS 2014	37
Dropping a database with SSMS 2014	38

---

<b>Creating and managing database schemas</b>	<b>39</b>
Managing schemas using T-SQL DDL statements	40
Managing schemas using SSMS 2014	40
<b>Creating and managing tables</b>	<b>41</b>
Creating and modifying tables	42
Creating and modifying tables with T-SQL DDL statements	42
Creating a table with T-SQL DDL statements	42
Modifying a table with T-SQL DDL statements	44
Dropping a table with T-SQL DDL statements	44
Creating and modifying tables with SSMS 2014	45
Creating a table with SSMS 2014	45
Modifying a table with SSMS 2014	46
Deleting a table with SSMS 2014	46
<b>Grant, deny, and revoke permissions to securables</b>	<b>46</b>
Grant, deny, and revoke permissions to securables with T-SQL DCL statements	46
Granting permissions to securables with T-SQL DCL statements	47
Denying permissions to securables with T-SQL DCL statements	47
Revoking permissions to securables with T-SQL DCL statements	48
Managing permissions using SSMS 2014	48
<b>Summary</b>	<b>48</b>
<b>Chapter 3: Data Retrieval Using Transact-SQL Statements</b>	<b>49</b>
<b>Understanding Transact-SQL SELECT, FROM, and WHERE clauses</b>	<b>50</b>
The SELECT statement	50
The FROM clause	51
The WHERE clause	51
<b>Using T-SQL functions in the query</b>	<b>52</b>
Aggregate functions	52
Configuration functions	53
Cursor functions	53
Date and time functions	53
Mathematical functions	54
Metadata functions	54
Rowset functions	54
Security functions	54
String functions	55
System statistical functions	55
<b>Multiple table queries using UNION, EXCEPT, INTERSECT, and JOINS</b>	<b>55</b>
The UNION operator	56
The EXCEPT operator	57
The INTERSECT operator	58

The JOIN operator	59
Using INNER JOIN	59
Using outer joins	60
<b>Subqueries</b>	<b>61</b>
Examples of subqueries	62
<b>Common Table Expressions</b>	<b>63</b>
<b>Organizing and grouping data</b>	<b>64</b>
The ORDER BY clause	64
The GROUP BY clause	65
The HAVING clause	65
The TOP clause	66
The DISTINCT clause	66
Pivoting and unpivoting data	66
<b>Using the Transact-SQL analytic window functions</b>	<b>68</b>
Ranking functions	69
PERCENT RANK	71
CUME_DIST	72
PERCENTILE_CONT and PERCENTILE_DISC	73
LEAD and LAG	74
FIRST_VALUE and LAST_VALUE	76
<b>Summary</b>	<b>77</b>
<b>Chapter 4: Data Modification with SQL Server Transact-SQL Statements</b>	<b>79</b>
<b>    Inserting data into SQL Server database tables</b>	<b>80</b>
The INSERT examples	82
Example 1 – insert a single row into a SQL Server database table	82
Example 2 – INSERT with the SELECT statement	84
Example 3 – INSERT with the EXEC statement	84
Example 4 – explicitly inserting data into the IDENTITY column	85
<b>    Updating data in SQL Server database tables</b>	<b>86</b>
The UPDATE statement examples	87
Example 1 – updating a single row	87
Example 2 – updating multiple rows	87
<b>    Deleting data from SQL Server database tables</b>	<b>88</b>
The DELETE statement examples	89
Example 1 – deleting a single row	89
Example 2 – deleting all rows	89
<b>    Using the MERGE statement</b>	<b>89</b>
The MERGE statement examples	91
<b>    The TRUNCATE TABLE statement</b>	<b>94</b>
<b>    The SELECT INTO statement</b>	<b>94</b>
<b>    Summary</b>	<b>95</b>

<b>Chapter 5: Understanding Advanced Database Programming</b>	<b>97</b>
<b>Objects and Error Handling</b>	<b>97</b>
<b>Creating and using variables</b>	<b>98</b>
Creating a local variable	98
Creating the cursor variable	99
Creating the table variable	99
<b>Control-of-flow keywords</b>	<b>100</b>
BEGIN...END keywords	100
The IF...ELSE expression	101
A CASE statement	102
WHILE, BREAK, and CONTINUE statements	102
RETURN, GOTO, and WAITFOR statements	103
<b>Creating and using views</b>	<b>104</b>
Creating views with Transact-SQL and SSMS 2014	104
Creating, altering, and dropping views with Transact-SQL DDL statements	104
Creating, altering, and dropping views with SSMS 2014	107
Indexed views	109
<b>Creating and using stored procedures</b>	<b>111</b>
Creating a stored procedure	113
Modifying a stored procedure	117
Dropping a stored procedure	118
Viewing stored procedures	119
Executing stored procedures	120
<b>Creating and using user-defined functions</b>	<b>120</b>
Creating user-defined functions	121
Creating a user-defined scalar function	121
Creating a user-defined table-valued function	124
Modifying user-defined functions	128
Using a user-defined table-valued function	129
Dropping user-defined functions	129
Viewing user-defined functions	130
<b>Creating and using triggers</b>	<b>131</b>
Nested triggers	131
Recursive triggers	132
DML triggers	132
Inserted and deleted logical tables	133
Creating DML triggers	133
Modifying a DML trigger	135
Dropping a DML trigger	135
Data Definition Language (DDL) triggers	135
The EVENTDATA function	135
Creating a DDL trigger	135
Modifying a DDL trigger	136

Dropping a DDL trigger	137
Disabling and enabling triggers	137
Viewing triggers	137
<b>Handling Transact-SQL errors</b>	<b>138</b>
An example of TRY...CATCH	139
An example of TRY...CATCH with THROW	140
An example of TRY...CATCH with RAISERROR	141
<b>Summary</b>	<b>141</b>
<b>Chapter 6: Performance Basics</b>	<b>143</b>
<b>Components of SQL Server Database Engine</b>	<b>143</b>
The SQL Server Relational Engine architecture	144
Parsing and binding	145
Query optimization	145
Query execution and plan caching	147
Query plan aging	148
The improved design in SQL Server 2014 for the cardinality estimation	148
Optimizing SQL Server for ad hoc workloads	148
Manually clearing the plan cache	149
<b>The SQL Server 2014 in-memory OLTP engine</b>	<b>149</b>
The limitations of memory-optimized tables	150
<b>Indexes</b>	<b>151</b>
The cost associated with indexes	151
How SQL Server uses indexes	151
Access without an index	152
Access with an index	152
The structure of indexes	152
<b>Index types</b>	<b>153</b>
Clustered indexes	153
Nonclustered indexes	154
Single-column indexes	154
Composite indexes	155
Covering indexes	157
Unique indexes	160
Spatial indexes	160
Partitioned indexes	161
Filtered indexes	161
Full-text indexes	162
XML indexes	163
Memory-optimized indexes	163
Columnstore indexes	164
<b>Guidelines for designing and optimizing indexes</b>	<b>168</b>
Avoid overindexing tables	168
Create a clustered index before creating nonclustered indexes when using clustered indexes	168
Index columns used in foreign keys	168

---

Index columns frequently used in joins	169
Use composite indexes and covering indexes to give the query optimizer greater flexibility	169
Limit key columns to columns with a high level of selectability	169
Pad indexes and specify the fill factor to reduce page splits	169
Rebuild indexes based on the fragmentation level	170
<b>Query optimization statistics</b>	<b>170</b>
Database-wide statistics options in SQL Server to automatically create and update statistics	171
Manually create and update statistics	171
Determine the date when the statistics were last updated	172
Using the DBCC SHOW_STATISTICS command	172
Using the sys.stats catalog view with the STATS_DATE() function	173
<b>The fundamentals of transactions</b>	<b>173</b>
Transaction modes	174
Implementing transactions	174
BEGIN TRANSACTION	174
COMMIT TRANSACTION	174
ROLLBACK TRANSACTION	175
SAVE TRANSACTION	175
An overview of locking	175
Basic locks	176
Optimistic and pessimistic locking	176
Transaction isolation	177
<b>SQL Server 2014 tools for monitoring and troubleshooting SQL Server performance</b>	<b>178</b>
Activity Monitor	178
The SQLServer:Locks performance object	178
Dynamic Management Views	179
SQL Server Profiler	179
The sp_who and sp_who2 system stored procedures	180
SQL Server Extended Events	180
<b>Summary</b>	<b>180</b>
<b>Index</b>	<b>181</b>

---





---

# Preface

Microsoft SQL Server is an enterprise database server that is the cornerstone of modern business applications and is in the center of the business processes of many leading organizations. The latest release of Microsoft SQL Server, SQL Server 2014, has many new features. These new features of SQL Server 2014 let you design, build, and deploy high-performance OLTP applications. Especially, the new in-memory technology of SQL Server 2014 helps you to design and implement high-performance OLTP applications. According to Microsoft, in some situations, implementing the new SQL Server 2014 in-memory technology for existing OLTP applications can improve the performance of these applications by 10 times. This book will provide you with all the skills you need to successfully design, build, and deploy databases using SQL Server 2014. Starting from the beginning, this book gives you an insight into the key stages of the SQL Server database process, provides you with an in-depth knowledge of the SQL Server database architecture, and shares tips to help you design the new database.

By sequentially working through the steps in each chapter, you will gain hands-on experience in designing, creating, and deploying SQL Server databases and objects. You will learn how to use SQL Server 2014 Management Studio and the advanced Transact-SQL queries to retrieve data from the SQL Server database. You will also learn how to add, modify, and delete data stored within a database. You will use Transact-SQL statements to create and manage advanced database objects that include scalar and table-valued functions, views, stored procedures, and triggers. Finally, you will learn about how the SQL Server 2014 relation engine works, how indexes and statistics improve query performance, and the new SQL Server 2014 in-memory technologies.

## What this book covers

*Chapter 1, Microsoft SQL Server Database Design Principles*, explains the database design process and the architecture and working of the SQL Server 2014 Storage Engine. This chapter covers the database development life cycle in detail, including the normalization and denormalization process, benefits of choosing appropriate data types, and the functioning of the SQL Server 2014 Storage Engine.

*Chapter 2, Understanding DDL and DCL Statements in SQL Server*, introduces the reader to the SQL Server 2014 Transact-SQL language elements and SQL Server 2014 Management Studio (SSMS 2014). This chapter explains Transact-SQL DDL, DCL, and DML language elements in detail, and how you can use them to create, manage, and secure SQL Server 2014 databases, schemas, and tables. This chapter also shows you how you can use SQL Server Management Studio to create and manage SQL Server 2014 databases, schemas, and tables. Finally, this chapter covers the purpose of SQL Server 2014 system databases and highlights the advantages and disadvantages of database recovery models.

*Chapter 3, Data Retrieval Using Transact-SQL Statements*, demonstrates how to query data from tables, how to write multiple table queries, and how to group, organize, and pivot result set data. This chapter explores the basic form of the SELECT statement and how it can be used to query data from tables. This chapter also highlights the different categories of built-in T-SQL functions and how you can use them in your SELECT statements. This chapter also explains different techniques that you can use to combine data from multiple tables, how to organize the data, and how to generate the summary data by grouping or pivoting it. Finally, this chapter covers the purpose of the CTE and SQL Server 2014 windowing functions and how to use them to quickly solve complex analytical tasks.

*Chapter 4, Data Modification with SQL Server Transact-SQL Statements*, illustrates how to add, modify, and delete data in tables using Transact-SQL DML statements. This chapter covers how to add data to a table using the INSERT statement, how to delete the data using the DELETE statement, and how to update existing data using the UPDATE statement. This chapter also covers the SELECT...INTO, MERGE, and TRUNCATE TABLE statements, and it highlights the key new enhancements of these statements in SQL Server 2014.

*Chapter 5, Understanding Advanced Database Programming Objects and Error Handling*, covers reusable programming objects that includes views, stored procedures (normal and natively compiled), functions and triggers (based on either DDL or DML), SQL Server 2014 control-of-flow statements, and structured error handling blocks. This chapter shows you how you can declare and use variables and how you can use control-of-flow statements to control your program execution. Next, it explains the purpose of views, stored procedures, user-defined functions, and triggers, and highlights the guidelines and restrictions to design each of these programmable objects. This chapter also illustrates how to handle errors that occur within Transact-SQL batches and programmable objects using the TRY...CATCH construct.

*Chapter 6, Performance Basics*, explains performance-related features of SQL Server 2014. This chapter first explains the architecture of the SQL Server Relational Engine. Then, it introduces the architecture of the SQL Server 2014 in-memory technology. Next, it covers all SQL Server 2014 index types and how they can be used to achieve optimal query performance while reducing the overall response time. Then, it explores the architectural differences of B-tree, Bw-tree, and xVelocity columnstore indexes. Finally, it explains core performance topics such as SQL Server query optimization statistics, SQL Server transactions and locks, and tools that come with SQL Server 2014 Database Engine, which you can use to monitor and troubleshoot its Database Engine performance.

## What you need for this book

The following are the software prerequisites to run the samples in the book:

- Windows 7.0 SP1 or later
- SQL Server 2014 Developer edition
- SQL Server Management Studio
- AdventureWorks2012 sample database, which is available for download from the CodePlex site at <http://msftdbprodsamples.codeplex.com/downloads/get/478214>

## Who this book is for

If you want to learn how to design, implement, and deliver successful database solutions with SQL Server 2014, this is the book for you.

## Conventions


In this book, you will find a number of styles of text that distinguish between different kinds of information. Here are some examples of these styles, and an explanation of their meaning.

Code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles are shown as follows: "Each row in the `Customer` table represents an individual customer."

A block of code is set as follows:

```
[default]
DECLARE @Table2 TABLE (
  COL1 [int],
  COL2 [varchar] (30),
  COL3 [datetime],
  INDEX [ixc_col3] CLUSTERED (col3)
    WITH (FILLFACTOR=80),
  INDEX [ixnc_col1_col2] NONCLUSTERED (col1, col2)
    WITH (FILLFACTOR=80)
);
```

**New terms** and **important words** are shown in bold. Words that you see on the screen, in menus or dialog boxes for example, appear in the text like this: "Right-click on **User-Defined Data Types** and choose **New User-Defined Data Type**."

 Warnings or important notes appear in a box like this.

 Tips and tricks appear like this.

## Reader feedback

Feedback from our readers is always welcome. Let us know what you think about this book – what you liked or may have disliked. Reader feedback is important for us to develop titles that you really get the most out of.

To send us general feedback, simply send an e-mail to [feedback@packtpub.com](mailto:feedback@packtpub.com), and mention the book title via the subject of your message.

If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, see our author guide on [www.packtpub.com/authors](http://www.packtpub.com/authors).

## Customer support

Now that you are the proud owner of a Packt book, we have a number of things to help you to get the most from your purchase.

## Downloading the example code

You can download the example code files for all Packt books you have purchased from your account at <http://www.packtpub.com>. If you purchased this book elsewhere, you can visit <http://www.packtpub.com/support> and register to have the files e-mailed directly to you.

## Errata

Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you find a mistake in one of our books – maybe a mistake in the text or the code – we would be grateful if you would report this to us. By doing so, you can save other readers from frustration and help us improve subsequent versions of this book. If you find any errata, please report them by visiting <http://www.packtpub.com/submit-errata>, selecting your book, clicking on the **errata submission form** link, and entering the details of your errata. Once your errata are verified, your submission will be accepted and the errata will be uploaded on our website, or added to any list of existing errata, under the Errata section of that title. Any existing errata can be viewed by selecting your title from <http://www.packtpub.com/support>.

## **Piracy**

Piracy of copyright material on the Internet is an ongoing problem across all media. At Packt, we take the protection of our copyright and licenses very seriously. If you come across any illegal copies of our works, in any form, on the Internet, please provide us with the location address or website name immediately so that we can pursue a remedy.

Please contact us at [copyright@packtpub.com](mailto:copyright@packtpub.com) with a link to the suspected pirated material.

We appreciate your help in protecting our authors, and our ability to bring you valuable content.

## **Questions**

You can contact us at [questions@packtpub.com](mailto:questions@packtpub.com) if you are having a problem with any aspect of the book, and we will do our best to address it.

---

# 1

## Microsoft SQL Server Database Design Principles

Database design is one of the most important tasks in the **Systems Development Life Cycle (SDLC)**, also referred to as **Application Development Life Cycle (ADLC)**. That's because databases are essential for all businesses, and good design is crucial to any business-critical, high-performance application. Poor database design results in wasted time during the development process and often leads to unusual databases that are unfit for use.

We'll be covering the following topics in this chapter:

- The database design process and considerations
- The table design process, which includes identifying entities and attributes, creating a relationship between entities, and ensuring data integrity
- The basics of data normalization
- The SQL Server database architecture
- The importance of choosing the appropriate data type



- [read Seve: Golf's Flawed Genius: The Updated Definitive Biography](#)
- [download online Night Frost for free](#)
- [download James Beard's & Hors D'oeuvre And Canapes](#)
- [download Secret Beaches Of The Salish Sea: The Southern Gulf Islands for free](#)
- [read Skeptics and True Believers pdf](#)
  
- <http://reseauplatoparis.com/library/What-to-Solve---Problems-and-Suggestions-for-Young-Mathematicians--Reprint-Edition---Oxford-Science-Publication>
- <http://interactmg.com/ebooks/The-Eating-of-the-Gods--An-Interpretation-of-Greek-Tragedy.pdf>
- <http://www.khoi.dk/?books/Dealing-with-the-Tough-Stuff--How-to-Achieve-Results-from-Key-Conversations--2nd-Edition-.pdf>
- <http://www.gateaerospaceforum.com/?library/Language--The-Cultural-Tool.pdf>
- <http://www.khoi.dk/?books/Skeptics-and-True-Believers.pdf>