

Training Kit (Exam 70-461): Querying Microsoft SQL Server 2012

Itzik Ben-Gan, Dejan Sarka, Ron Talmage

ISBN: 978-0-7356-6605-4

First printing: December, 2012

To ensure the ongoing accuracy of this book and its companion content, we've reviewed and confirmed the errors listed below. If you find a new error, we hope you'll report it to us on our website: www.microsoftpressstore.com/contact-us/errata.

Page	Location	Description	Date corrected
428, 429, 445-449	Code section in step 2	<p>Total of 14 related fixes:</p> <p>Page 428 Line 33: Reads: VALUES(101, N'Test2: Bad categoryid', 1, 1, 18.00, 0); Should read: VALUES(101, N'Test2: New productid', 1, 1, 18.00, 0);</p> <p>Page 429 Line 10: Reads: VALUES(101, N'Test2: Bad categoryid', 1, 1, 18.00, 0); Should read: VALUES(101, N'Test2: New productid', 1, 1, 18.00, 0);</p> <p>Page 445, line 15 Reads: VALUES(1, N'Test1: Ok categoryid', 1, 1, 18.00, 0); Should read: VALUES(1, N'Test1: Duplicate productid', 1, 1, 18.00, 0);</p> <p>Page 445, line 41: Reads: VALUES(1, N'Test1: Ok categoryid', 1, 1, 18.00, 0); Should read: VALUES(1, N'Test1: Duplicate productid', 1, 1, 18.00, 0);</p> <p>Page 446, Line 9 Reads: VALUES(101, N'Test2: Bad categoryid', 1, 1, 18.00, 0); Should read: VALUES(101, N'Test2: New productid', 1, 1, 18.00, 0);</p> <p>Page 446, Line 37: Reads: VALUES(1, N'Test1: Ok categoryid', 1, 1, 18.00, 0); Should read: VALUES(1, N'Test1: Duplicate productid', 1, 1, 18.00, 0);</p>	4/12/2013

Page	Location	Description	Date corrected
		<p>Page 447 Line 23: Reads: VALUES(1, N'Test1: Ok categoryid', 1, 1, 18.00, 0); Should read: VALUES(1, N'Test1: Duplicate productid', 1, 1, 18.00, 0);</p> <p>Page 447, Line 33 Reads: VALUES(101, N'Test2: Bad categoryid', 1, 1, 18.00, 0); Should read: VALUES(101, N'Test2: New productid', 1, 1, 18.00, 0);</p> <p>Page 448, Line 17: Reads: VALUES(1, N'Test1: Ok categoryid', 1, 1, 18.00, 0); Should read: VALUES(1, N'Test1: Duplicate productid', 1, 1, 18.00, 0);</p> <p>Page 448, Line 19: Reads: VALUES(101, N'Test2: Bad categoryid', 1, 1, 18.00, 0); Should read: VALUES(101, N'Test2: New productid', 1, 1, 18.00, 0);</p> <p>Page 449 Line 2: Reads: VALUES(1, N'Test1: Ok categoryid', 1, 1, 18.00, 0); Should read: VALUES(1, N'Test1: Duplicate productid', 1, 1, 18.00, 0);</p> <p>Page 449, Line 5: Reads: VALUES(101, N'Test2: Bad categoryid', 1, 1, 18.00, 0); Should read: VALUES(101, N'Test2: New productid', 1, 1, 18.00, 0);</p> <p>Page 449 Line 25: Reads: VALUES(1, N'Test1: Ok categoryid', 1, 1, 18.00, 0); Should read: VALUES(1, N'Test1: Duplicate productid', 1, 1, 18.00, 0);</p> <p>Page 449, Line 28: Reads: VALUES(101, N'Test2: Bad categoryid', 1, 1, 18.00, 0); Should read: VALUES(101, N'Test2: New productid', 1, 1, 18.00, 0);</p>	

Page	Location	Description	Date corrected
Practice test question on CD	question ID 461P_3.4_06 (explanation section)	<p>Reads: You use the FLOOR function to return the smallest integer less than or equal to a specific numeric value.</p> <p>Should read: You use the FLOOR function to return the largest integer less than or equal to a specific numeric value.</p>	11/12/2014
xxvi	First bulleted item	The components listed within the SQL Server 2012 Setup Feature Selection list should be all second-level bullets.	11/12/2014
17	Paragraph 3, Group Rows Based on the GROUP BY Clause	<p>Reads: YEAR(orderdate).</p> <p>Should read: YEAR(hiredate).</p>	4/12/2013
19	Second paragraph, last sentence	<p>Reads: In this case, the SELECT list returns for each country and order year group ...</p> <p>Should read: In this case, the SELECT list returns for each country and hire year group ...</p>	11/12/2014
20	First paragraph of 6, Handle Presentation Ordering	<p>Reads: The query indicates that the result rows should be ordered first by country (in ascending order by default), and then by numemployees, descending, yielding the following output.</p> <p>Should read: The query indicates that the result rows should be ordered first by country and then by yearhired, descending, yielding the following output.</p>	4/12/2013
28	First sentence	<p>Reads: One of most typical mistakes that T-SQL developers make...</p> <p>Should read: One of the most typical mistakes that T-SQL developers make...</p>	4/12/2013
34	Second Paragraph	<p>Reads: The rules say that the first character must be a letter in the range A through Z (lower or uppercase), underscore (_), at sign (@), or number sign (#).</p> <p>Should read: The rules say that the first character must be a letter defined by the Unicode Standard 3.2 (a-z, A-Z, and letters from other Unicode languages), underscore (_), at sign (@), or number sign (#).</p>	8/23/2013

Page	Location	Description	Date corrected
44	Sixth paragraph	<p>Reads: The GETUTCDATE function returns the current date and time in UTC terms as a DATETIME type, and SYSUTCDATE does the same, only returning the result as the more precise DATETIME2 type.</p> <p>Should read: The GETUTCDATE function returns the current date and time in UTC terms as a DATETIME type, and SYSUTCDATETIME does the same, only returning the result as the more precise DATETIME2 type.</p>	4/12/2013
47	Code block, second line	<p>Reads: CONCAT_NULL_YIELDS_NULL_INPUT</p> <p>Should read: CONCAT_NULL_YIELDS_NULL</p>	4/12/2013
48	Third paragraph	<p>Reads: As an example, the expression PATINDEX('%[0-9]%', 'abcd123efgh') looks for the first occurrence of a digit (a character in the range 0–9) in the second input, returning the position 6 in this case.</p> <p>Should read: position 5, a-d are 4 positions so the digit 1 is on the fifth position.</p>	4/12/2013
49	Last sentence within String formatting section	<p>"CASE Expression and Related Functions" currently has: Reads: FORMAT(1759, '000000000')</p> <p>Should read: FORMAT(1759, '000000000')</p>	4/12/2013
51	Last paragraph	<p>Reads: Observe that the type of the COALESCE expression is determined by the returned element, whereas...</p> <p>Should read: Observe that the type of the COALESCE expression is the type of the input expression with the highest precedence, whereas...</p>	11/12/2014
51	First paragraph	<p>After the first sentence, which reads: The COALESCE function accepts a list of expressions as input and returns the first that is not NULL, or NULL if all are NULLs.</p> <p>The following sentence should be added: If all inputs are the untyped NULL constant, SQL Server generates an error.</p>	12/13/2013
53	Quick Check Answer #2	<p>Reads: NEWSEQUENTIAL ID (note the space)</p> <p>Should read: NEWSEQUENTIALID (without the space)</p>	4/12/2013

Page	Location	Description	Date corrected
63	Third paragraph (right after the table)	<p>Reads: However, when NULLs are possible in the data, things get trickier. Consider the customer location columns country, region, and city in the Sales.Customers table.</p> <p>Should read: However, when NULLs are possible in the data, things get trickier. Consider the employee location columns country, region, and city in the HR.Employees table.</p>	11/12/2014
79	First paragraph, second line	<p>Reads: When DISTINCT is used, duplicates are removed; then the result rows don't necessarily map to source rows in a one-to-one manner, but rather than one-to-many.</p> <p>Should read: When DISTINCT is used, duplicates are removed; then the result rows don't necessarily map to source rows in a one-to-one manner, but rather one-to-many.</p>	11/12/2014
82	Exercise 2.1, sentence before query	<p>Reads: Add shipperid DESC to the ORDER BY clause, as follows.</p> <p>Should read: Add shippeddate DESC to the ORDER BY clause, as follows.</p>	4/12/2013
85	Last paragraph, first sentence	<p>Reads: The PERCENT option puts a ceiling on the resulting number of rows if it's not whole.</p> <p>Should read: The PERCENT option rounds up the number of rows if it's not whole.</p>	11/12/2014
87	Last SQL Statement on the page	<p>Reads: SELECT TOP (3) WITH TIES orderid..... ORDER BY orderdate DESC, orderid DESC;</p> <p>Should read: SELECT TOP (3) orderid, orderdate, custid, empid FROM Sales.Orders ORDER BY orderdate DESC, orderid DESC;</p>	4/12/2013
89	Last paragraph	<p>Reads: OFFSET (@pagesize - 1) * @pagesize ROWS FETCH NEXT @pagesize ROWS ONLY;</p> <p>Should read: OFFSET (@pagenum - 1) * @pagesize ROWS FETCH NEXT @pagesize ROWS ONLY;</p>	4/12/2013
103	Last query	<p>Reads: S.N</p> <p>Should read: S.n</p>	11/12/2014

Page	Location	Description	Date corrected
114	Fourth paragraph	Reads: ... based on the JOIN keyword for cross and outer joins. Should read: ... based on the JOIN keyword for cross and inner joins.	4/12/2013
116	Last sentence on page	Reads: by using parentheses or by repositioning the ON clauses Should read: by using parentheses and by repositioning the ON clauses	4/12/2013
128	Third sentence	Reads: As an example, consider the recursive CTE from the section about CTEs that returned the management chain leading to employee 9. Should read: As an example, consider the recursive CTE from the section about CTEs that returned the management chain leading all the way up to the CEO for a specified employee.	8/23/2013
133	under the column "unitprice", second row of results	Reads: 10.00X Should read: 10.00	4/12/2013
140	Third paragraph from bottom	Reads: Finally, set operators have precedence: INTERSECT precedes UNION and EXCEPT, and UNION and EXCEPT are considered equal. Should read: Finally, set operators have precedence: INTERSECT precedes UNION and EXCEPT, and UNION and EXCEPT are evaluated from left to right based on their position in the expression.	4/12/2013
146	Lesson 3, Question 1 Answer	Reads: Correct Answers as A, B, and C. Should read: Correct Answers as A, C, and D.	4/12/2013
154	Second code sample	Reads: SELECT S.shipperid, MAX(S.companyname) AS numorders, COUNT(*) AS shippedorders FROM Sales.Shippers AS S... Should read: SELECT S.shipperid, MAX(S.companyname) AS companyname, COUNT(*) AS numorders FROM Sales.Shippers AS S...	4/12/2013

Page	Location	Description	Date corrected																																																						
154	Bottom of page	<p>Add the following paragraph to bottom of page for more information about spatial aggregates:</p> <p>Note that with SQL Server 2012, you can create user defined aggregates (UDAs) by using .NET code based on the common language runtime (CLR). SQL Server 2012 provides some built-in CLR UDAs for the spatial data types GEOMETRY and GEOGRAPHY and also allows you to create new UDAs operating on spatial types as inputs. For more details please refer to books online.</p>	4/12/2013																																																						
155	The two example outputs	<p>Should read:</p> <pre>SELECT shipperid, YEAR(shippeddate) AS shipyear, COUNT(*) AS numorders FROM Sales.Orders WHERE shippeddate IS NOT NULL -- exclude unshipped orders GROUP BY GROUPING SETS ((shipperid, YEAR(shippeddate)), (shipperid), (YEAR(shippeddate)), ());</pre> <p>Should read:</p> <table> <thead> <tr> <th>shipperid</th> <th>shipyear</th> <th>numorders</th> </tr> <tr> <th>-----</th> <th>-----</th> <th>-----</th> </tr> </thead> <tbody> <tr><td>1</td><td>2006</td><td>36</td></tr> <tr><td>2</td><td>2006</td><td>56</td></tr> <tr><td>3</td><td>2006</td><td>51</td></tr> <tr><td>NULL</td><td>2006</td><td>143</td></tr> <tr><td>1</td><td>2007</td><td>130</td></tr> <tr><td>2</td><td>2007</td><td>143</td></tr> <tr><td>3</td><td>2007</td><td>125</td></tr> <tr><td>NULL</td><td>2007</td><td>398</td></tr> <tr><td>1</td><td>2008</td><td>79</td></tr> <tr><td>2</td><td>2008</td><td>116</td></tr> <tr><td>3</td><td>2008</td><td>73</td></tr> <tr><td>NULL</td><td>2008</td><td>268</td></tr> <tr><td>NULL</td><td>NULL</td><td>809</td></tr> <tr><td>3</td><td>NULL</td><td>249</td></tr> <tr><td>1</td><td>NULL</td><td>245</td></tr> <tr><td>2</td><td>NULL</td><td>315</td></tr> </tbody> </table>	shipperid	shipyear	numorders	-----	-----	-----	1	2006	36	2	2006	56	3	2006	51	NULL	2006	143	1	2007	130	2	2007	143	3	2007	125	NULL	2007	398	1	2008	79	2	2008	116	3	2008	73	NULL	2008	268	NULL	NULL	809	3	NULL	249	1	NULL	245	2	NULL	315	8/23/2013
shipperid	shipyear	numorders																																																							
-----	-----	-----																																																							
1	2006	36																																																							
2	2006	56																																																							
3	2006	51																																																							
NULL	2006	143																																																							
1	2007	130																																																							
2	2007	143																																																							
3	2007	125																																																							
NULL	2007	398																																																							
1	2008	79																																																							
2	2008	116																																																							
3	2008	73																																																							
NULL	2008	268																																																							
NULL	NULL	809																																																							
3	NULL	249																																																							
1	NULL	245																																																							
2	NULL	315																																																							

Page	Location	Description	Date corrected
157	Third paragraph	<p>Reads:</p> <pre>SELECT shipcountry, GROUPING(shipcountry) AS grpcountry, shipregion , GROUPING(shipregion) AS grpcountry, shipcity , GROUPING(shipcity) Agrpcountry, COUNT(*) AS numorders FROM Sales.Orders GROUP BY ROLLUP(shipcountry, shipregion, shipcity);</pre> <p>Should read:</p> <pre>SELECT shipcountry, GROUPING(shipcountry) AS grpcountry, shipregion , GROUPING(shipregion) AS grpregion, shipcity , GROUPING(shipcity) AS grpcity, COUNT(*) AS numorders FROM Sales.Orders GROUP BY ROLLUP(shipcountry, shipregion, shipcity);</pre>	4/12/2013
179	Last query	<p>Reads:</p> <p>This query generates the following output (shown here in abbreviated form).</p> <p>Should read:</p> <p>This query generates the following output (shown here in abbreviated form, with presentation ordering not guaranteed).</p>	11/12/2014
179	Second paragraph	<p>Reads:</p> <p>Because an explicit \ wasn't specified, both functions relied on the default offset of 1.</p> <p>Should read:</p> <p>Because an explicit offset wasn't specified, both functions relied on the default offset of 1.</p>	4/12/2013
193	Fifth paragraph	<p>Reads:</p> <p>You might need to restart SQL Server. After you restart it, check whether the filters were successfully installed by using the sys.sp_help_fulltext_components system procedure again</p> <p>Should read:</p> <p>You might need to restart SQL Server. After you restart it, check whether the filters were successfully installed by using the sys.sp_help_fulltext_system_components system procedure again</p>	8/23/2013
193	Code sample at foot of page	<p>Reads:</p> <pre>SELECT stoplist_id, names FROM sys.fulltext_stoplists</pre> <p>Should read:</p> <pre>SELECT stoplist_id, name FROM sys.fulltext_stoplists</pre>	4/12/2013

Page	Location	Description	Date corrected
198	First paragraph	<p>Reads: If the folder is C:\TK461, then you can use the following code directly; otherwise, change the folder in the OPENROWSET functions appropriately.</p> <p>Should read: If the folder is C:\TK70461, then you can use the following code directly; otherwise, change the folder in the OPENROWSET functions appropriately.</p>	4/12/2013
201	Top of the page	<p>Reads: You can create full-text catalogs and indexes by using SQL Server Full-Text Search and Sematinc Search.</p> <p>Should read: You can create full-text catalogs and indexes by using SQL Server Full-Text Search and Semantic Search.</p>	8/23/2013
201	Third bullet point in lesson summary	<p>Reads: You can use the sys.dm_fts_parser dynamic management view to check how Full-Text Search breaks your documents into words, creates inflectional forms of words, and more.</p> <p>Should read: You can use the sys.dm_fts_parser dynamic management object to check how Full-Text Search breaks your documents into words, creates inflectional forms of words, and more.</p>	8/23/2013
203	Paragraphs 3, 6, 7, 8	<p>The word "proximity" in the third paragraph, the word "generation" in the sixth paragraph, and the word "weighted" in the eighth paragraph should formatted as italic.</p> <p>The word "generation" in the seventh paragraph should not be italic.</p>	4/12/2013
225	Opening tag for the root element in the XML	<p>To go strictly with the rules and have an always correct XML, please replace the second and the third row in the XML document at the bottom half of the page 255.</p> <p>Reads: xmlns:xsd=http://www.w3.org/2001/XMLSchema</p> <p>xmlns:sqltypes=http://schemas.microsoft.com/sqlserver/2004/sqltypes</p> <p>Should read: xmlns:xsd="http://www.w3.org/2001/XMLSchema"</p> <p>xmlns:sqltypes="http://schemas.microsoft.com/sqlserver/2004/sqltypes"</p>	11/12/2014

Page	Location	Description	Date corrected
227	Paragraph in middle of page	<p>Reads: The WHERE clause of the query limits the output to two customers, with only every second order for each customer retrieved.</p> <p>Should read: The WHERE clause of the query limits the output to two customers, with only even-numbered orders for each customer retrieved.</p>	2/23/2015
250	Last paragraph	<p>Reads: The methods support querying (the query() method), retrieving atomic values (the value() method), checking existence (the exist() method), modifying sections within the XML data (the modify() method) as opposed to overriding the whole thing, and shredding XML data into multiple rows in a result set (the nodes() method).</p> <p>Should read: The methods support querying (the query() method), retrieving atomic values (the value() method), checking existence (the exist() method), modifying sections within the XML data (the modify() method) as opposed to overwriting the whole thing, and shredding XML data into multiple rows in a result set (the nodes() method).</p>	8/23/2013
258, 259	Exercises 2 and 3	<p>Pg 258, step 2 reads: Return all orders for the first customer as XML.</p> <p>Should read: Return all orders for the customer with @custid equal to 1 (the first customer in the XML document) as XML.</p> <p>Pg 259, step 3 reads: Shred all orders information for the first customer.</p> <p>Should read: Shred all orders information for the customer with @custid equal to 1.</p>	4/12/2013
271	Third paragraph	<p>Reads: Even though you can embed special characters such as @, #, and \$ in an identifier for a schema, table, or column name, that action makes the identifier delimited, no longer regular.</p> <p>Should read: If you embed special characters other than @, #, and \$ in an identifier for a schema, table, or column name, that action makes the identifier delimited, no longer regular.</p>	4/12/2013
271	Third paragraph	<p>Reads: or special characters other than those previously mentioned</p> <p>Should read: or non-alphanumeric characters other than @, \$, #, and _.</p>	8/23/2013

Page	Location	Description	Date corrected
274	Second paragraph, line eight	Reads: Many of the TSQL12 tabs Should read: Many of the TSQL2012 tables	4/12/2013
278	Exercise 2, step 2	Reads: ALTER COLUMN description NVARCHAR(500); Should read: ALTER COLUMN description NVARCHAR(500) NOT NULL;	4/12/2013
278	Exercise 2, step 1	Should read: CONSTRAINT PK_CategoriesTest PRIMARY KEY(categoryid) Note the additional word Test in the name.	4/12/2013
279	Exercise 2, step 4	Reads: ALTER COLUMN description NVARCHAR(500); Should read: ALTER COLUMN description NVARCHAR(500) NOT NULL;	4/12/2013
279	Problem 3	Reads: Test for the existence of anNULLs Should read: Inspect the table for the existence of any NULLs	8/23/2013
284	Quick Check Answer #2	Reads: withing Should read: within	5/21/2015
284	Quick Check, bottom of page	Question reads: 2. Can a primary key on one table have the same name as the primary key in another table in the same database? Should read: 2. Can a primary key on one table have the same name as the primary key in another table in the same database and in the same schema? Answer reads: 2. No, all table constraints must have unique names in a database. Should read: 2. No, all table constraints must have unique names withing the schema of a database.	4/12/2013
284	Second paragraph	Reads: By default, the index will be nonclustered. Should read: The unique index can be either clustered or nonclustered.	4/12/2013

Page	Location	Description	Date corrected
288	Fifth paragraph	<p>Reads: Remember that default constraints, like all other constraints, are database-wide objects. Their names must be unique across the entire database. No two tables can have default constraints named the same.</p> <p>Should read: Remember that default constraints, like all other constraints, are schema-wide objects. Their names must be unique across an entire schema. No two tables scoped to the same schema can have default constraints named the same.</p>	8/23/2013
295	Answers, Lesson 1, Question 3	<p>Reads: 3. Correct Answer: C A. Incorrect: Specifying NULL must come after the data type. B. Incorrect: PERMIT NULL is not a valid construct in the CREATE TABLE statement. C. Correct: You specify NULL right after the data type. D. Incorrect: ALLOW NULL is not a valid construct in the CREATE TABLE statement.</p> <p>Should read: 3. Correct Answer: D A. Incorrect: Specifying NULL must come after the data type. B. Incorrect: ALLOW NULL is not a valid construct in the CREATE TABLE statement. C. Incorrect: PERMIT NULL is not a valid construct in the CREATE TABLE statement. D. Correct: You specify NULL right after the data type.</p>	4/12/2013
295	Lesson 1, Answer 1.C.	<p>Reads: Correct Answers: A and D</p> <p>Should read: Correct Answers: A, C, and D</p> <p>Reads: C. Incorrect: A regular identifier cannot contain a dollar sign (\$).</p> <p>Should read: C. Correct: A regular identifier may contain a dollar sign (\$).</p>	8/23/2013
296	Last paragraph	<p>Reads: However, you can query sys.constraints</p> <p>Should read: However, you can query sys.key_constraints</p>	8/23/2013
300	Estimated lesson time: 20 minutes	<p>Reads: Estimated lesson time: 20 minutes</p> <p>Should read: Estimated lesson time: 45 minutes</p>	12/13/2013

Page	Location	Description	Date corrected
300	Lesson Summary	<p>Reads: After this lesson, you will be able to: Use the CREATE VIEW statement to create a table</p> <p>Should read: Use the CREATE VIEW statement to create a view</p>	4/12/2013
300	First paragraph	<p>Reads: In this lesson, you learn about these commands and their related options.</p> <p>Should read: In this lesson, you learn how to create and manage views and inline functions.</p>	4/12/2013
302	Ninth paragraph	<p>Reads: The view cannot have its schema definitions changed unless the view is dropped.</p> <p>Should read: The table cannot have its schema definitions changed unless the view is dropped.</p>	4/12/2013
307	Quick Check Answer paragraph. Answer 1.	<p>Reads: (using the UNION statemt)</p> <p>Should read: (using the UNION, UNION ALL, EXCEPT or INTERSECT statements)</p>	11/12/2014
307	First code sample	<p>Reads: SELECT SCHEMA_NAME,</p> <p>Should read: SELECT TABLE_SCHEMA,</p>	4/12/2013
309	Last paragraph	<p>Reads: The referenced objects cannot be altered</p> <p>Should read: The schemas of the referenced objects cannot be altered</p>	8/23/2013
309	Last sentence	<p>Reads: The referenced objects cannot be altered unless the view is dropped or the WITH SCHEMABINDING option is removed</p> <p>Should read: The referenced objects cannot be altered unless the function is dropped or the WITH SCHEMABINDING option is removed</p>	8/23/2013
309	Last sentence on bottom of page	<p>Reads: The referenced objects cannot be altered unless the view is dropped or the WITH SCHEMABINDING option is removed.</p> <p>Should read: The referenced objects cannot be altered unless the function is dropped or the WITH SCHEMABINDING option is removed.</p>	4/12/2013

Page	Location	Description	Date corrected
309	Fourth text (non-code) paragraph	<p>Reads: You can query the function but pass the year you want to see, as follows.</p> <p>Should read: You can query the function and pass the year you want to see, as follows.</p>	4/12/2013
312	Point 9	<p>Reads: and add a WHERE clause to filter the results</p> <p>Should read: and add a HAVING clause to filter the results</p>	4/12/2013
312	Exercise 1	<p>Reads: IF OBJECT_ID(' Sales.OrderTotalsByYearCustShip', N'V') IS NOT NULL</p> <p>Should read: IF OBJECT_ID(N'Sales.OrderTotalsByYearCustShip', N'V') IS NOT NULL</p>	4/12/2013
314	Lesson Review, 2.B	<p>Reads: B. The tables referred to in the view cannot be altered unless the view is first altered.</p> <p>Should read: B. The tables referred to in the view cannot be altered unless the view's SELECT statement is first altered.</p>	4/12/2013
317	Abstraction Layer - second paragraph	<p>Reads: For example, suppose the database DB1</p> <p>Should read: For example, suppose the database ReportsDB</p>	4/12/2013
320	Exercise 1, Item 6, Note reader aid	<p>Reads: Note that there is no unique key on the combination of columns in the GROUP BY of the Sales.OrderValues view.</p> <p>Should read: Note that there is no unique key on the combination of columns, minus the orderid column, in the GROUP BY of the Sales.OrderValues view.</p>	4/12/2013
324	Practice 2	<p>Reads: Products.Categories</p> <p>Should read: Production.Categories</p>	8/23/2013
331	First paragraph of "INSERT VALUES" section	<p>Reads: Sales.MyOrderValues</p> <p>Should read: Sales.MyOrders</p>	11/12/2014

Page	Location	Description	Date corrected
339	Exercise 1, Item 2	<p>Reads: SELECT ISNULL(CAST(custid AS INT), -1) AS custid, ...</p> <p>Should read: SELECT ISNULL(custid, -1) AS custid, ...</p>	4/12/2013
339	Step 5 in Exercise 1	<p>Reads: 5. After executing the previous INSERT statement, query the Sales.MyOrders table...</p> <p>Should read: 5. After executing the previous INSERT statement, query the Sales.MyCustomers table...</p>	4/12/2013
353	Update Statement in Step 3 of Exercise 1	<p>Reads: UPDATE TGT SET TGT.custid = SRC.custid , TGT.companyname = SRC.companyname , TGT.contactname = SRC.contactname , TGT.contacttitle = SRC.contacttitle, TGT.address = SRC.address , TGT.city = SRC.city , TGT.region = SRC.region , TGT.postalcode = SRC.postalcode , TGT.country = SRC.country , TGT.phone = SRC.phone , TGT.fax = SRC.fax FROM Sales.MyCustomers AS TGT INNER JOIN Sales.Customers AS SRC ON TGT.custid = SRC.custid;</p> <p>Should read: UPDATE TGT SET TGT.companyname = SRC.companyname , TGT.contactname = SRC.contactname , TGT.contacttitle = SRC.contacttitle, TGT.address = SRC.address , TGT.city = SRC.city , TGT.region = SRC.region , TGT.postalcode = SRC.postalcode , TGT.country = SRC.country , TGT.phone = SRC.phone , TGT.fax = SRC.fax FROM Sales.MyCustomers AS TGT INNER JOIN Sales.Customers AS SRC ON TGT.custid = SRC.custid;</p>	4/12/2013

Page	Location	Description	Date corrected
354	SQL statement	<p>Should read:</p> <pre> WITH C AS (SELECT TGT.companyname AS tgt_companyname , SRC.companyname AS src_companyname , TGT.contactname AS tgt_contactname , SRC.contactname AS src_contactname , TGT.contacttitle AS tgt_contacttitle, SRC.contacttitle AS src_contacttitle, TGT.address AS tgt_address , SRC.address AS src_address , TGT.city AS tgt_city , SRC.city AS src_city , TGT.region AS tgt_region , SRC.region AS src_region , TGT.postalcode AS tgt_postalcode , SRC.postalcode AS src_postalcode , TGT.country AS tgt_country , SRC.country AS src_country , TGT.phone AS tgt_phone , SRC.phone AS src_phone , TGT.fax AS tgt_fax , SRC.fax AS src_fax FROM Sales.MyCustomers AS TGT INNER JOIN Sales.Customers AS SRC ON TGT.custid = SRC.custid) UPDATE C SET tgt_companyname = src_companyname , tgt_contactname = src_contactname , tgt_contacttitle = src_contacttitle, tgt_address = src_address , tgt_city = src_city , tgt_region = src_region , tgt_postalcode = src_postalcode , tgt_country = src_country , tgt_phone = src_phone , tgt_fax = src_fax; </pre>	8/23/2013
358	First paragraph	<p>Reads:</p> <p>As you can see, the code uses an infinite loop (WHERE 1 = 1 is always true).</p> <p>Should read:</p> <p>As you can see, the code uses an infinite loop (WHILE 1 = 1 is always true).</p>	8/23/2013
361	Exercise 1, Item 3	<p>Reads:</p> <pre>LEFT OUTER JOIN Sales.Orders AS SRC</pre> <p>Should read:</p> <pre>LEFT OUTER JOIN Sales.MyOrders AS SRC</pre>	4/12/2013

Page	Location	Description	Date corrected
368	Top of page - 2.B. explanation	Reads: The DELETE statement does not reset an IDENTITY property. Should read: The DELETE statement does not reset an IDENTITY property.	8/23/2013
369	Exam objectives in this chapter - section	Reads: Modify data by using INSERT, UPDATE, and DELETE statements. Should read: * Modify Data * Modify data by using INSERT, UPDATE, and DELETE statements (use output clause). * Combining datasets	11/12/2014
380	Step 4 output	Should read: <pre> type start_value minimum_value ----- int 1 1 current_value increment is_cycling ----- 1 1 1 </pre>	4/12/2013
388	Second paragraph	Reads: TGT.custid = SRC.custid OR Should read: TGT.custid <> SRC.custid OR	4/12/2013
392	Query in step 3	Reads: WITH SRC AS MERGE INTO Sales.MyOrders AS TGT... Should read: MERGE INTO Sales.MyOrders AS TGT...	4/12/2013

Page	Location	Description	Date corrected
392	Query in step 3	<p>Reads:</p> <pre> WITH SRC AS (SELECT * FROM Sales.Orders WHERE shipcountry = N'Norway') MERGE INTO Sales.MyOrders AS TGT USING Sales.Orders AS SRC ON SRC.orderid = TGT.orderid WHEN MATCHED AND (TGT.custid <> SRC.custid OR TGT.empid <> SRC.empid OR TGT.orderdate <> SRC.orderdate) THEN UPDATE SET TGT.custid = SRC.custid, TGT.empid = SRC.empid, TGT.orderdate = SRC.orderdate WHEN NOT MATCHED THEN INSERT VALUES(SRC.orderid, SRC.custid, SRC.empid, SRC.orderdate); Should read: WITH SRC AS (SELECT * FROM Sales.Orders WHERE shipcountry = N'Norway') MERGE INTO Sales.MyOrders AS TGT USING SRC ON SRC.orderid = TGT.orderid WHEN MATCHED AND (TGT.custid <> SRC.custid OR TGT.empid <> SRC.empid OR TGT.orderdate <> SRC.orderdate) THEN UPDATE SET TGT.custid = SRC.custid, TGT.empid = SRC.empid, TGT.orderdate = SRC.orderdate WHEN NOT MATCHED THEN INSERT VALUES(SRC.orderid, SRC.custid, SRC.empid, SRC.orderdate); The change is in the USING clause: instead of "USING Sales.Orders AS SRC" it should be "USING SRC" </pre>	4/12/2013

403	Code block above "Lesson Summary"	<p>Reads:</p> <pre> IF OBJECT_ID('Sales.MyOrdersArchive') IS NOT NULL DROP TABLE Sales. MyOrdersArchive'; Should read: IF OBJECT_ID('Sales.MyOrdersArchive') IS NOT NULL DROP TABLE Sales.MyOrdersArchive; </pre>	4/12/2013
-----	-----------------------------------	--	-----------

Page	Location	Description	Date corrected
412	Last sentence before the last paragraph	<p>Reads: http://msdn.microsoft.com/en-us/library/ms186957(v=sql.90).aspx</p> <p>Should read: http://msdn.microsoft.com/en-us/library/ms186957(v=sql.110).aspx</p>	12/13/2013
415	First paragraph	<p>Line three reads: SQL Server maintains all its internal persistent system tables by using transactions that it classifies as system transactions. These transactions are not under user control.</p> <p>Should read: SQL Server uses <i><i>system transactions</i></i> to maintain all its internal persistent system tables; these types of transactions are not under user control.</p> <p>Line six reads: These are transactions created by users in the process of changing and even reading data, whether automatically, implicitly, or explicitly.</p> <p>Should read: Transactions created by users in the process of changing or reading data, whether automatically, implicitly, or explicitly, are called <i><i>user transactions</i></i>.</p>	4/12/2013
420	Fourth paragraph, Example	<p>Reads: BEGIN TRANSACTION Tran1;</p> <p>Should read: SELECT * FROM sys.dm_tran_active_transactions</p>	12/13/2013
421	Second paragraph, line 34	Italicize the bold title Distributed transactions	4/12/2013
421	Third line	<p>Reads: RESTORE DATABASE TSQ2012 FROM DISK = 'C:\SQLBackups\TSQL2012.bak'</p> <p>Should read: RESTORE DATABASE TSQ2012 FROM DISK = 'C:\SQLBackups\TSQL2012.bak'</p>	8/23/2013
423	Third paragraph	<p>Reads: If two sessions request an exclusive lock on the same resource, and one is granted the request, then the other session must wait until the first releases its exclusive lock</p> <p>Should read: A block occurs when one session has an exclusive lock on a resource, preventing another session from obtaining any kind of lock on the resource. The term 'block' is a keyword and should be italicized.</p>	4/12/2013

Page	Location	Description	Date corrected
424	First Quick Check box	<p>Reads:</p> <ol style="list-style-type: none"> 1. Can readers block readers? 2. Can readers block writers? <p>Should read:</p> <ol style="list-style-type: none"> 1. Can readers (shared locks) block readers? 2. Can readers block writers (exclusive locks)? 	4/12/2013
433	Table 12.8	<p>Reads:</p> <pre>SELECT lastname, firstname, region FROM HR.Employees;</pre> <p>Should read:</p> <pre>SELECT lastname, firstname, postalcode FROM HR.Employees</pre> <pre>WHERE empid = 1;</pre> <p>Reads:</p> <p><Results returned show region in original state for empid = 1></p> <p>Should read:</p> <p><Results returned show postalcode in original stat for empid = 1></p> <p>Reads:</p> <p><region for empid = 1 rolled back to original value></p> <p>Should read:</p> <p><postalcode for empid = 1 rolled back to original value></p>	4/12/2013
433	First paragraph	<p>Reads:</p> <p>Other than in this exercise, use the READ-COMMITTED table hint instead.</p> <p>Should read:</p> <p>Other than in this exercise, use the READUNCOMMITTED table hint instead.</p>	4/12/2013
437	Chapter 12, Lesson 2: Implementing Error Handling, p. 437	<p>Reads:</p> <p>'Error in % stored procedure</p> <p>Should read:</p> <p>'Error in %s stored procedure'</p>	4/12/2013
437	Penultimate paragraph, line 27	<p>Reads:</p> <p>You can issue RAISERROR with a severity level > 20</p> <p>Should read:</p> <p>You can issue RAISERROR with a severity level > 18</p>	4/12/2013

Page	Location	Description	Date corrected
440	Second sentence in last paragraph	<p>Reads: However, custom code for structured error handling can quickly become complex.</p> <p>Should read: However, custom code for unstructured error handling can quickly become complex.</p>	8/23/2013
443	Fourth paragraph	<p>Reads: To return the error number, you could add it to the @error_message string.</p> <p>Should read: To return the original error number, you could add it to the msg_str parameter of the RAISERROR statement.</p>	4/12/2013
445	Question 1 code-lines	<p>Reads: O</p> <p>Should read: GO IF @@TRANCOUNT <> 0 ROLLBACK TRANSACTION</p>	8/23/2013
448	Section 12-2 exercise 3.1, first paragraph	<p>Reads: by default will not roll back a transactin</p> <p>Should read: by default will not abort a transaction</p>	4/12/2013
453	Fourth paragraph	<p>Reads: the `ANSI standard and the SQL Server default</p> <p>Should read: the ANSI standard and the SQL Server default</p>	4/12/2013
453	Page 453, line 20 and line 33 (two places)	<p>Reads: addressFROM</p> <p>Should read: address FOM</p>	4/12/2013

Page	Location	Description	Date corrected
459 - 461	Step code blocks	<p>Pg 460, step 6 code, line 14: Reads: @statment Should read: @statement</p> <p>Pg 461, step 2 code, line 7: Reads: @statment Should read: @statement</p> <p>Pg 459, exercise 2, step 1 code, line 10/Pg 460, step 6/Pg 461, step 1: Reads: SET @SQLString = ' Should read: SET @SQLString = N'</p> <p>Pg 459, exercise 2, step 2/Pg 460 steps 3, 4 and 5/Pg 461 steps 7 and 8: Reads: @address = ' Should read: @address = N'</p>	4/12/2013
460	Practice exercise 2, step 5	<p>Reads: SELECT 1 -</p> <p>Should read: SELECT 1 --</p>	4/12/2013
465	Answer for question 2 in lesson 1	<p>Reads: Correct Answers: B and C</p> <p>Should read: Correct Answers:</p> <p>Reads: Correct: A single COMMIT commits only the innermost level of a nested transaction.</p> <p>Should read: C. Incorrect: A single COMMIT commits data only at the outermost level of a nested transaction.</p>	4/12/2013
466	Answer to Question 3	<p>Reads: A. If a T-SQL error with a severity level > 16 occurs, the transaction will be aborted.</p> <p>Should read: A. If a T-SQL error with a severity level > 16 occurs, the transaction will not be aborted.</p>	8/23/2013

Page	Location	Description	Date corrected
469	Chapters 13 and 14	Chapter 14 mistakenly preceded Chapter 13 in the first edition.	12/14/2012
470	Third bullet at top of page	Reads: ...which make calls to externally compiled data definition languages (DLLs) Should read: ...which make calls to externally compiled dynamic-link libraries (DLLs)	5/21/2015
471	Second paragraph	Reads: You cannot use CREATE AGGREGATE, RULE, DEFAULT, CREATE, FUNCTION, TRIGGER, PROCEDURE, or VIEW statements. Should read: You cannot use a CREATE statement with any of the following object types: AGGREGATE, RULE, DEFAULT, FUNCTION, TRIGGER, PROCEDURE, or VIEW.	4/12/2013
477	Code sample at bottom of page	Reads: When the IF and ELSE statements are used without BEGIN/END blocks Should read: When the IF or ELSE statements are used without BEGIN/END blocks	8/23/2013
479	Second Paragraph	Reads: The numbers 5 and 6 are skipped because of the CONTINUE statement; Should read: The numbers 5 and 6 are skipped because of the CONTINUE branch;	4/12/2013
485	Exercise 1, Item 7	Reads: THROW 50000, 'dbo.BackupDatabases: @databasename must be User or System', 0; Should read: THROW 50000, 'dbo.BackupDatabases: @databasetype must be User or System', 0;	4/12/2013
489	After step 8	After line 15, insert: 9. Drop the stored procedure.	8/23/2013
491	DML Triggers, second paragraph	Reads: A trigger executes only once for each DML statement, no matter how many rows Should read: A trigger executes only once for each DML statement, no matter how many rows	4/12/2013
492	Last paragraph	Reads: When an UPDATE or DELETE occurs and no rows are affected, Should read: When an INSERT, UPDATE or DELETE occurs and no rows are affected,	4/12/2013

Page	Location	Description	Date corrected
492	Fifth paragraph	<p>Replace the entire code snippet with:</p> <pre>CREATE TRIGGER TriggerName ON [dbo].[TableName] FOR DELETE, INSERT, UPDATE AS BEGIN SET NOCOUNT ON END</pre>	4/12/2013
492	Last paragraph	<p>Reads: In a trigger definition, the FOR can be replaced with either AFTER or INSTEAD OF to determine the type of trigger.</p> <p>Should read: In a trigger definition, AFTER is the default type of trigger when you specify FOR. But you can replace FOR with either AFTER or INSTEAD OF to determine the type of trigger.</p>	4/12/2013
493	Second paragraph	<p>Should read: Now add an existence test to the OBJECT_ID() function, using 'TR' as the object type. Define it on the TSQL2012 Sales.OrderDetails table, and call it Sales.tr_SalesOrderDetailsDML.</p> <pre>IF OBJECT_ID('Sales.tr_SalesOrderDetailsDML', 'TR') IS NOT NULL DROP TRIGGER Sales.tr_SalesOrderDetailsDML; GO CREATE TRIGGER Sales.tr_SalesOrderDetailsDML ON Sales.OrderDetails AFTER DELETE, INSERT, UPDATE AS BEGIN IF @@ROWCOUNT = 0 RETURN; -- Must be 1st statement SET NOCOUNT ON; END;</pre> <p>Now add SELECT statements on the inserted and deleted tables.</p>	4/12/2013
495	Code example for Instead Of dml trigger	<p>Reads: HAVING COUNT(*) > 1</p> <p>Should read: HAVING COUNT(*) > 0</p>	4/12/2013
495	Line 29	<p>Reads: JOIN Production.Categories A C</p> <p>Should read: LEFT JOIN Production.Categories AS C</p>	8/23/2013

Page	Location	Description	Date corrected
497	First paragraph	<p>Reads: Also, the feature can be disabled in sp_configure, so to do this exercise, make sure Disallow Results From Triggers is enabled.</p> <p>Should read: Also, the feature can be disabled in sp_configure, so to do this exercise, make sure Disallow Results From Triggers is disabled.</p>	4/12/2013
498	Exercise 2, step 2 - second sentence	<p>Reads: The trigger finds the violating row, which has a unitprice of 9.00 and a discount of 0.40.</p> <p>Should read: The trigger finds the violating row, which has a unitprice of 9.00 and a discount of 0.60.</p>	8/23/2013
498	Third paragraph	<p>Reads: in the TSQL12 database.</p> <p>Should read: in the TSQL2012 database.</p>	4/12/2013
507	Question 1, Part 1, second line	<p>Reads: Sales.SalesOrder</p> <p>Should read: Sales.OrderDetails</p>	11/12/2014
508	Fourth paragraph, Exercise 2 intro	<p>Reads: In this exercise, you write two table-valued UDFs: one as an inline function and the other as a multistatement function.</p> <p>Should read: In this exercise, you write an inline table-valued UDF.</p>	4/12/2013
509	Fourth paragraph	<p>Reads: Use the variable as the parameters</p> <p>Should read: Use the variables as parameters</p>	4/12/2013
513	WAITFOR Section	<p>Reads: For example, the following code waits until 11:45.</p> <p>Should read: For example, the following code waits until 11:46.</p>	4/12/2013
513	Question 2 Lesson 1	<p>Reads: C. Correct: Both passed parameters and temporary tables are visible to called stored procedures.</p> <p>Should read: C. Correct: The called procedure can see temporary tables and parameters passed to it from the calling procedure.</p>	8/23/2013

Page	Location	Description	Date corrected
518	Bottom of page	<p>Reads: WHERE A.col3 = constant 1 AND B.col4 = constant2</p> <p>Should read: WHERE A.col3 = constant1 AND B.col4 = constant2</p>	4/12/2013
520	Third paragraph, third sentence	<p>Reads: The actual execution is performed by the storage engine.</p> <p>Should read: The actual execution is performed by the execution engine of the relational engine and the storage engine together.</p>	2/23/2015
530	Bottom of page	<p>Reads:</p> <pre>SELECT C.custid, C.companyname, O.orderid, O.orderdate FROM Sales.Customers AS C INNER JOIN Sales.Orders AS O ON C.custid = O.custid SELECT C.custid, C.companyname, O.orderid, O.orderdate FROM Sales.Customers AS C INNER JOIN Sales.Orders AS O ON C.custid = O.custid WHERE O.custid < 5</pre> <p>Should read:</p> <pre>SELECT C.custid, C.companyname, O.orderid, O.orderdate FROM Sales.Customers AS C INNER JOIN Sales.Orders AS O ON C.custid = O.custid; SELECT C.custid, C.companyname, O.orderid, O.orderdate FROM Sales.Customers AS C INNER JOIN Sales.Orders AS O ON C.custid = O.custid WHERE O.custid < 5;</pre>	4/12/2013
540	Step 4 of exercise 1. Second sentence	<p>Reads: The function only needs two parameters: unitprice and qty.</p> <p>Should read: The function needs three parameters: @unitprice, @qty, and @discount.</p>	4/12/2013

Page	Location	Description	Date corrected
551	First paragraph	<p>Reads:</p> <p>Indexes are always organized as balanced trees. Other indexes, such as indexes that do not contain all of the data and serve as pointers to table rows for quick seeks, are called nonclustered indexes.</p> <p>Should read:</p> <p>Indexes are always organized as balanced trees. Indexes that do not contain all of the data like clustered indexes and serve as pointers to table rows for quick seeks are called nonclustered indexes.</p>	2/23/2015
552	Last code sample on	<p>Reads:</p> <pre>CREATE TABLE dbo.TestStructure (id INT NOT NULL, filler1 CHAR(36) NOT NULL, filler2 CHAR(216) NOT NULL);</pre> <p>Should read:</p> <pre>USE tempdb; GO CREATE TABLE dbo.TestStructure (id INT NOT NULL, filler1 CHAR(36) NOT NULL, filler2 CHAR(216) NOT NULL);</pre>	4/12/2013
552	Fifth paragraph	<p>Reads:</p> <p>You can find out how many pages are allocated for an object from the sys.dm_db_index_physical_stats dynamic management view or with the help of the dbo.sp_spaceused system procedure, as shown in the following code.</p> <p>Should read:</p> <p>You can find out how many pages are allocated for an object from the sys.dm_db_index_physical_stats dynamic management function or with the help of the dbo.sp_spaceused system procedure, as shown in the following code.</p>	8/23/2013
556	Second sentence in penultimate paragraph	<p>Reads:</p> <p>Note that the order data (the od column in the figure) is used for the clustering key.</p> <p>Should read:</p> <p>Note that the order date (the od column in the figure) is used for the clustering key.</p>	8/23/2013

Page	Location	Description	Date corrected
557	Figure 15-2, intermediate-level pages	<p>Page 1:47184 (first intermediate-level page) should be revised as follows: Third row reads: 20041210 1 1:47122</p> <p>Should read: 20041210 7 1:47122</p> <p>Page 1:47185 (second intermediate-level page) should be revised as follows: Last row reads: 20040206 170 1:47831</p> <p>Should read: 20050206 170 1:47831</p> <p>Page 1:36838 (last intermediate-level page) should be revised as follows: Last three rows read: 20071231 496 1:41477 20071231 536 1:41478 20071231 576 1:41479</p> <p>Should read: 20081231 496 1:41477 20081231 536 1:41478 20081231 576 1:41479</p>	8/23/2013
564	First paragraph	<p>Reads: orgnize</p> <p>Should read: organized</p>	4/12/2013
564	Last sentence on page	<p>Reads: If the clustering key is narrow—for example a 4-byte integer—than SQL Server can also accommodate more rows on a leaf-level page than when RID is used as the row locator.</p> <p>Should read: If the clustering key is narrow—for example a 4-byte integer—then SQL Server can also accommodate more rows on a leaf-level page than when RID is used as the row locator.</p>	8/23/2013

Page	Location	Description	Date corrected
566	Second sentence on page	<p>Reads: It is the same example of the customers' orders table; order data is used for the clustering key, and order ID is used for the key of the nonclustered inde</p> <p>Should read: It is the same example of th customers' orders table; order date is used for the clustering key, and order ID is used for the key of the nonclustered index</p>	8/23/2013
566	First paragraph	<p>Reads: It is the same example of the customers' orders table; order data is used for the clustering key, and order ID is used for the key of the nonclustered index.</p> <p>Should read: It is the same example of the customers' orders table; order date is used for the clustering key, and order ID is used for the key of the nonclustered index.</p>	4/12/2013
571	Exercise 1, Item 7	<p>Reads: DECLARE @i AS int = 0; WHILE @i < 24472 BEGIN SET @i = @i + 1; INSERT INTO dbo.TestStructure (id, filler1, filler2) VALUES (@i, FORMAT(@i,'0000'), 'b'); END;</p> <p>Should read: DECLARE @i AS int = 0; WHILE @i < 24472 BEGIN SET @i = @i + 1; INSERT INTO dbo.TestStructure (id, filler1, filler2) VALUES (@i, FORMAT(@i,'00000'), 'b'); END;</p>	4/12/2013

Page	Location	Description	Date corrected
572	Exercise 2, Item 3	<p>Reads:</p> <pre> DECLARE @i AS int = 0; WHILE @i < 28864 BEGIN SET @i = @i + 1; INSERT INTO dbo.TestStructure (id, filler1, filler2) VALUES (@i, FORMAT(@i,'0000'), 'b'); END; </pre> <p>Should read:</p> <pre> DECLARE @i AS int = 0; WHILE @i < 28864 BEGIN SET @i = @i + 1; INSERT INTO dbo.TestStructure (id, filler1, filler2) VALUES (@i, FORMAT(@i,'00000'), 'b'); END; </pre>	4/12/2013
573	Last paragraph	<p>Reads:</p> <p>You need to know which types of queries can benefit from indexes, and which types of queries does not use indexes, even if indexes exist.</p> <p>Should read:</p> <p>You need to know which types of queries can benefit from indexes, and which types of queries do not use indexes, even if indexes exist.</p>	4/12/2013
585	Last paragraph	<p>Add this sentence to the last bullet at the bottom of the page: "This option is turned off by default, and changing this option has no effect unless AUTO_UPDATE_STATISTICS is set to ON."</p>	4/12/2013
592	Practice Exercise 2, step 8	<p>Reads:</p> <pre> ALTER DATABASE TSQL2012 SET AUTO_CREATE_STATISTICS ON WITH NO_WAIT; EXEC sys.sp_updatestats; DROP STATISTICS Sales.Orders.st_shipcity; DROP INDEX idx_nc_custid_shipcity ON Sales.Orders; </pre> <p>Should read:</p> <pre> DROP STATISTICS Sales.Orders.st_shipcity; DROP INDEX idx_nc_custid_shipcity ON Sales.Orders; ALTER DATABASE TSQL2012 SET AUTO_CREATE_STATISTICS ON WITH NO_WAIT; EXEC sys.sp_updatestats; </pre>	4/12/2013
607	Last paragraph	<p>Reads:</p> <p>f2</p> <p>Should read:</p> <p>f squared</p>	4/12/2013

Page	Location	Description	Date corrected
632	Second paragraph	<p>Reads: If a table is organized as a heap, then the only access method available to SQL Server is a table scan</p> <p>Should read: If a table is organized as a heap and does not have any nonclustered indexes, then the only access method available to SQL Server is a table scan</p>	12/13/2013
633	Second paragraph	<p>Reads: SQL Server uses an allocation order scan for a clustered table if a query does not request any specific order, if the isolation level is Read Uncommitted, or if you are working in a readonly environment.</p> <p>Should read: SQL Server may use an allocation order scan for a clustered table if (the table has more than 64 pages and the query does not request any specific order) and (the isolation level is Read Uncommitted or you are working in a readonly environment).</p>	4/12/2013
639	Second paragraph	<p>Reads: Then the algorithm checks the next rows from the other side and adds them to the output until they match the predicate.</p> <p>Should read: Then the algorithm checks the next rows from the other side and adds them to the output as long as they match the predicate.</p>	4/12/2013
645	Exercise 1, step 6, second sentence	<p>Reads: As you probably expected, SQL Server scanned the clustered Sales.Customers table, then scanned the nonclustered covering index on the orderdate column of the Sales.Orders table, and then used the Merge Join iterator to join the data.</p> <p>Should read: As you probably expected, SQL Server scanned the clustered Sales.Customers table, then scanned the nonclustered index on the custid column of the Sales.Orders table, and then used the Merge Join iterator to join the data.</p>	12/13/2013
661	Last sentence on page	<p>Reads: You can specify query hints in an INSERT statement except when a SELECT clause is used inside the statement.</p> <p>Should read: You cannot specify query hints in an INSERT statement except when a SELECT clause is used inside the statement.</p>	8/23/2013