Microsoft® SQL Server® 2008 T-SQL Fundamentals

Itzik Ben-Gan (Solid Quality Mentors)

ISBN: 978-0-7356-2601-0

First printing: October, 2008

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
xi-xii	Table of Contents, page references of Chapter 10 through end	Beginning with the Chapter 10 page entries, all page entries are behind by 2. For example, 319 should be 321, 322 should be 324, and so on.	
15	Second bulleted item	Reads: Resource The Resource database was added as of SQL Server 2005 and it holds all system objects. When you query metadata information in a database, this information appears to be local to the database but in practice it resides in the Resource database. Should read: Resource The Resource database is a hidden, read-only database that was added as of SQL Server 2005 and it holds the definitions of all system objects. When you query system objects in a database, they appear to reside in the sys schema of the local database, but in practice their definitions reside in the Resource database.	
47	Third paragraph	The following paragraph is irrelevant and should be disregarded: Note that the reason that I specified the decimal value 100. (one hundred dot) in the expressions instead of the integer 100 is in order to cause implicit conversion of the integer values val and SUM(val) to decimal values. Otherwise, the division would have been an integer division and the fractional part would have been truncated.	
47	Last code block, fifth line	Reads: NTILE(100) OVER(ORDER BY val) AS ntile Should read: NTILE(10) OVER(ORDER BY val) AS ntile	
53	Step 3	Reads: 3. + (Positive), − (Negative), + (Add), (+ Concatenate), − (Subtract) Should read: 3. + (Positive), − (Negative), + (Add), + (Concatenate), − (Subtract)	

Page	Location	Description	Date corrected
57	Throughout both code blocks	All occurences that read: Less then 1000 Should read: Less than 1000	
62	Last paragraph, first sentence	Reads: You assume that SQL Server evaluates the expressions from left to right, and that if the expression col1 <> evaluates to FALSE, SQL Server will short-circuit; that is, it doesn't bother to evaluate the expression 10/col1 > 2 because at this point it is known that the whole expression is FALSE. Should read: You assume that SQL Server evaluates the expressions from left to right, and that if the expression col1 <> evaluates to FALSE, SQL Server will short-circuit; that is, it doesn't bother to evaluate the expression col2/col1 > 2 because at this point it is known that the whole expression is FALSE.	
63	First and third paragraphs	First paragraph, last sentence reads: You can see that if SQL Server decides to process the expression 10/col1 > 2 first, this query might fail because of a divide-by-zero error. Should read: You can see that if SQL Server decides to process the expression col2/col1 > 2 first, this query might fail because of a divide-by-zero error. Third paragraph, second sentence reads: Only if the first CASE expression does not evaluate to TRUE—meaning that col1 is not 0—does the second WHEN clause check whether the expression 10/col1 > 2 evaluates to TRUE. Should read: Only if the first CASE expression does not evaluate to TRUE—meaning that col1 is not 0—does the second WHEN clause check whether the expression col2/col1 > 2 evaluates to TRUE.	

Page	Location	Description	Date corrected
76	Paragraph after Table 2-1	Reads: The storage requirements for the last three data types in Table 2-1 (TIME, DATETIME2, and DATETIMEOFFSET) depend on the accuracy you choose. You specify the accuracy as an integer in the range 0 to 7 representing the fractional second precision. For example, TIME(0) means 0 fractional second precision—in other words, one-second accuracy. TIME(3) means one-millisecond accuracy, and TIME(7) means 100-nanosecond accuracy. If you don't specify a fractional second precision, SQL Server assumes 7 by default with all three aforementioned types. Should read: The storage requirements for the last three data types in Table 2-1 (TIME, DATETIME2, and DATETIMEOFFSET) depend on the precision you choose. You specify the precision as an integer in the range 0 to 7 representing the fractional second precision. For example, TIME(0) means 0 fractional second precision—in other words, one-second precision. TIME(3) means one-millisecond precision, and TIME(7) means 100-nanosecond precision. If you don't specify a fractional second precision, SQL Server assumes 7 by default with all three aforementioned types.	
77	First full paragraph, fourth sentence	Reads: SQL Server defi nes precedence among datatypes, and will usually implicitly covert… Should read: SQL Server defi nes precedence among datatypes, and will usually implicitly convert the…	
87	First "Syntax" line	Reads: DATEPART(dt_val, part) Should read: DATEPART(part, dt_val)	
88	First syntax entry	Reads: DATENAME(dt_val, part) Should read: DATENAME(part, dt_val)	
98	Solution 5, first sentence	Reads: Because the request involves activity in the year 2004, the query should have a WHERE clause with the appropriate date range filter (orderdate >= '20040101' AND orderdate < '20050101'). Should read: Because the request involves activity in the year 2007, the query should have a WHERE clause with the appropriate date range filter (orderdate >= '20070101' AND orderdate < '20080101').	

Page	Location	Description	Date corrected
117	Last paragraph, first sentence	Reads: The next step is to extend the previous query, adding a left outer join between Nums and the Orders tables. Should read: The next step is to extend the previous query, adding a left outer join between the Nums and Orders tables.	
141	Third paragraph, first sentence	Reads: This query returns the order ID 10274. The outer row's order ID—10248—is compared with the inner one—10274—and because there's no match in this case, the outer row is filtered out. Should read: This query returns the order ID 10739. The outer row's order ID—10248—is compared with the inner one—10739—and because there's no match in this case, the outer row is filtered out.	
157	Step 4, third sentence	Reads: ...because the same country can have more than one query. Should read: ...because the same country can have more than one customer.	
157	Solutions 3 and 4	Solution 3 reads: Write an outer query against the Employees table returning employees whose IDs appear in the set of employee IDs returned by the subquery. Should read: Write an outer query against the Employees table returning employees whose IDs do not appear in the set of employee IDs returned by the subquery. Solution 4 reads: Write an outer query against the Customers table that filters only customer rows where the country attribute appears in the set of countries returned by the subquery. Should read: Write an outer query against the Customers table that filters only customer rows where the country does not appear in the set of countries returned by the subquery.	
157	Third sentence of Item 4	Reads: ...because the same country can have more than one query. Should read: ...because the same country can have more than one customer.	

Page	Location	Description	Date corrected
174	Fourth line of last code sample	Reads: SELECT TOP (100) Should read: SELECT TOP (100) PERCENT	
188	Exercise 5-2	Reads: Using the CROSS APPLY operator and the function you created in Exercise 4-1, return, for each supplier, the two most expensive products. Should read: Using the CROSS APPLY operator and the function you created in Exercise 5-1, return, for each supplier, the two most expensive products.	
194	"The UNION Set Operation" paragraph, last sentence	Reads: The area marked with diagonal lines represents the result of the set operation. Should read: The area marked with a gray background represents the result of the set operation.	
209	Second line	Reads: ...before the rows from Customers,... Should read: ...before the rows from Suppliers,…	
209	Second line	Reads: ...before the rows from Customers,... Should read: ...before the rows from Shippers,…	

Page	Location	Description	Date corrected
215	Sample code output at bottom of page	Reads:	

```
orderid    orderdate                     empid   custid  qty
----------  ----------------------------  --------  -------  ----
10001      2007-12-24 00:00:00.000  2        A       12
10005      2007-12-24 00:00:00.000  1        B       20
10006      2008-01-18 00:00:00.000  1        C       14
20001      2008-02-12 00:00:00.000  2        B       12
20002      2009-02-16 00:00:00.000  1        C       20
30001      2007-08-02 00:00:00.000  3        A       10
30003      2009-04-18 00:00:00.000  2        B       15
30004      2007-04-18 00:00:00.000  3        C       22
30007      2009-09-07 00:00:00.000  3        D       30
40001      2008-01-09 00:00:00.000  2        A       40
40005      2009-02-12 00:00:00.000  3        A       10
```

Should read:

```
orderid    orderdate                     empid   custid  qty
----------  ----------------------------  --------  -------  ----
10001      2007-12-24                    2        A       12
10005      2007-12-24                    1        B       20
10006      2008-01-18                    1        C       14
20001      2008-02-12                    2        B       12
20002      2009-02-16                    1        C       20
30001      2007-08-02                    3        A       10
30003      2009-04-18                    2        B       15
30004      2007-04-18                    3        C       22
30007      2009-09-07                    3        D       30
40001      2008-01-09                    2        A       40
40005      2009-02-12                    3        A       10
```

Page	Location	Description	Date corrected
217	Last line on the page	Reads: You achieve this is by not applying the PIVOT operator to the original Should read: You achieve this by not applying the PIVOT operator to the original	
221	Second full sentence	Reads: You need apply a cross join between the EmpCustOrders table and a table that has a row for each customer. Should read: You need to apply a CROSS JOIN between the EmpCustOrders table and a table that has a row for each customer.	
230	Second paragraph, last sentence	Reads: The grouping set (a, c) is represented by the integer 10 ($1\times8 + 0\times4 + 1\times2 + 0\times1$), and so on. Should read: The grouping set (a, c) is represented by the integer 5 ($0\times8 + 1\times4 + 0\times2 + 1\times1$), and so on.	

Page	Location	Description	Date corrected
230	First paragraph, second sentence	Reads: For example, all rows where empid is 0 and custid is 0 are associated with the grouping set (empid, custid). Should read: For example, all rows where grpempis 0 and grpcustis 0 are associated with the grouping set (empid, custid).	
246	Third full paragraph	Reads: For example, the following code demonstrates how to insert a row to T1 with the explicit value 5 in keycol: Should read: For example, the following code demonstrates how to insert a row into T1 with the explicit value 5 in keycol:	
276	Second sentence in Solution 5	Reads: You can join Orders and Customers based on a match between the order's customer ID and the customers customer ID. Should read: You can join Orders and Customers based on a match between the order's customer ID and the customer's customer ID.	
288	Query output	Should be: spid most_recent_sql_handle ---- -- 52 0x01000800DE2DB71FB0936F0500000000000000000000000 53 0x0200000063FC7D052E09844778CDD615CFE7A2D1FB411802	
289	Second sentence, first paragraph	Reads: As for the blocker, in this example you can see the statement that caused the problem, but keep in mind that the blocker may continue work and that the last thing you see in the code isn't necessarily the statement that caused the trouble. Should read: As for the blocker, in this example you can see the statement that caused the problem, but keep in mind that the blocker may continue working and that the last thing you see in the code isn't necessarily the statement that caused the trouble.	
328	Last paragraph, last sentence (not including code at the bottom of the page)	Reads: For example, the following code runs a full backup of the sample database TSQLFundamentals2008 if today is the first day of the month, and a differential backup (changes since last full backup) if today is not the last day of the month. Should read: For example, the following code runs a full backup of the sample database TSQLFundamentals2008 if today is the first day of the month, and a differential backup (changes since last full backup) if today is not the first day of the month.	

Page	Location	Description	Date corrected
329	First code block, third line	Reads: PRINT 'Today is not the first day of the month.' Should read: PRINT 'Today is not the first day of the month.';	
329	Lines 5 and 6	Reads: BACKUP DATABASE TSQLFundamentals2008 TO DISK = 'C:\Temp\TSQLFundamentals2008_Diff.BAK' WITH INIT; Should read: BACKUP DATABASE TSQLFundamentals2008 TO DISK = 'C:\Temp\TSQLFundamentals2008_Diff.BAK' WITH INIT, DIFFERENTIAL;	
333	Code block, third-to-last line	Reads: @ordermonth DATETIME, Should read: @ordermonth AS DATETIME,	
347	First paragraph, first sentence	Reads: As of SQL Server 2005 you can chose whether to develop a routine with T-SQL or with .NET code based on the Common Language Runtime (CLR) integration in the product. Should read: As of SQL Server 2005, you can choose whether to develop a routine with T-SQL or with .NET code based on the Common Language Runtime (CLR) integration in the product.	
348	First paragraph, first sentence	Reads: The function calculates the age as the difference, in terms of years, between the birth year and the event year, minus 1 year in case within the year, the event month and day is smaller than the birth month and day. Should read: The function calculates the age as the difference, in terms of years, between the birth year and the event year, minus 1 year in case within the year, the event month and day is earlier than the birth month and day.	
371	Figure A-11	Middle table name reads: Sales.OrdersDetails Should read: Sales.OrderDetails Bottom-middle table name reads: Production, Products Should read: Production.Products	