Training Kit (Exam 70-463): Implementing a Data Warehouse with Microsoft SQL Server 2012

Dejan Sarka, Matija Lah, Grega Jerkic

ISBN: 978-0-7356-6609-2

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
	Throughout the book	The following term should be replaced throughout the book (occurs three times). Reads: SQL Server Development Tools Should read: SQL Server Data Tools Two other terms were used inconsistently throughout the book, Reads: Object Browser Action: Should be used only when referring to SQL Server Data Tools and the SQL Query Analyzer. Reads: Object Explorer Action: Should be used only when referring to SQL Server Management System.	8/16/2013
i-ii	Bottom of page i, top of page ii	Reads: Chapter 12 Lessons 1, 2, and 3 Should read: Chapter 12 Lessons 1 and 2	9/25/2014

Page	Location	Description	Date corrected
xxviii-xxix	From "Hardware and Operating System Requirements" section through "Using the Companion CD" section	Note About Companion Content & Sample Databases All of the companion content and sample databases are freely available here for download: http://go.microsoft.com/FWLink/?Linkid=260986 If you cannot access the files, please contact our Customer Service department at msinput@microsoft.com. An earlier version of the companion content webpage contained an incorrect URL, which is now fixed. You may need to refresh your browser's cache before the new page will load for you. We are sorry about any confusion caused by our error.	9/25/2014
xxviii-xxix	Second main bulleted paragraph	Reads: · Windows Software Development Kit (SDK) or Microsoft Visual Studio 2010 The Windows SDK provides tools, compilers, headers, libraries, code samples, and a new help system that you can use to create applications that run on Windows. You need the Windows SDK for Chapter 19, "Implementing Custom Code in SSIS Packages" only. If you already have Visual Studio 2010, you do not need the Windows SDK. If you need the Windows SDK, you need to download the appropriate version for your operating system. For Windows 7, Windows Server 2003 R2 Standard Edition (32-bit x86), Windows Server 2003 R2 Standard x64 Edition, Windows Server 2008, Windows Server 2008 R2, Windows Vista, or Windows XP Service Pack 3, use the Microsoft Windows SDK for Windows 7 and the Microsoft .NET Framework 4 from: http://www.microsoft.com/en-us/download/details.aspx?id=8279 Should read: · Microsoft Visual Studio 2010 or at least Microsoft Visual C# 2010 Express You will need Visual Studio 2010, or at least the Express Edition of Visual C# 2010, for Chapter 19, "Implementing Custom Code in SSIS Packages" only. If you do not have access to the complete version of Visual Studio 2010, you can download Visual C# 2010 Express from http://www.visualstudio.com/en-us/downloads/download-visual-studio-vs#DownloadFamilies_4.	9/25/2014
6	Third paragraph, last sentence	Reads: For example, you might have only the latest customer address, which might prevent you from calculating historical sales by country correctly. Should read: For example, you might have only the latest customer address (from which you extract the customer's current country), which might prevent you from calculating historical sales by country correctly.	8/16/2013

Page	Location	Description		Date corrected
9	"Snowflake Schema" section, second paragraph, last sentence	Reads: However, in the DimDate dimension, if you know the month, you obviously know the calendar quarter, and if you know the calendar quarter, you know the calendar year. Should read: However, in the DimDate dimension, if you know the month, you obviously know the calendar quarter, and if you know the calendar quarter, you know the calendar semester.		8/16/2013
16	Lesson review, question 2	Reads: 2. You are creating a quick POC project. Should read: 2. You are creating a quick POC DW project.		
18	Fifth paragraph, first sentence	Reads: …as you sow in the previous lesson. Should read: …as you saw in the previous lesson.		
20	First sentence after figure	Reads: Some attributes of the DimDate edimension include the following (not in the order shown in the figure Should read: Some attributes of the DimDate dimension include the following (not in the order shown in the figure):		8/16/2013
26	Exercise 2, fourth bulleted item	Reads: In the DimSalesReason dimension, it seems that there is a natural hierarchy: SalesReasonType --> SalesReasonName. Should read: In the DimSalesReason dimension, it seems that there is a natural hierarchy: SalesReasonReasonType --> SalesReasonName.		8/16/2013
28	Fifth paragraph, second sentence	Reads: In addition, the Source Order Details table has the ProductId foreign key column. The Quantity column is the measure. Format change: The word "Source" should not be capitalized and should not be in italics.		8/16/2013
50	Table, second and third rows, Column name and Data type columns	Reads: MaritalStatus NCHAR(1) Gender NCHAR(1) Should read: MaritalStatus NCHAR(5) Gender NCHAR(5)		8/16/2013

Page	Location	Description	Date corrected
50-51	Exercise 2, step 2 code block	Reads:	8/16/2013

Reads:

```
CREATE TABLE dbo.Customers
(
 CustomerDwKey INT NOT NULL,
 CustomerKey INT NOT NULL,
 FullName NVARCHAR(150) NULL,
 EmailAddress NVARCHAR(50) NULL,
 BirthDate DATE NULL,
 MaritalStatus NCHAR(1) NULL,
 Gender NCHAR(1) NULL,
 Education NVARCHAR(40) NULL,
 Occupation NVARCHAR(100) NULL,
 City NVARCHAR(30) NULL,
 StateProvince NVARCHAR(50) NULL,
 CountryRegion NVARCHAR(50) NULL,
 Age AS
 CASE
 WHEN DATEDIFF(yy, BirthDate, CURRENT_TIMESTAMP) <= 40
 THEN 'Younger'
 WHEN DATEDIFF(yy, BirthDate, CURRENT_TIMESTAMP) > 50
 THEN 'Older'
 ELSE 'Middle Age'
 END,
 CurrentFlag BIT NOT NULL DEFAULT 1,
 CONSTRAINT PK_Customers PRIMARY KEY (CustomerDwKey)
);
```

Should read:

```
CREATE TABLE dbo.Customers
(
 CustomerDwKey INT NOT NULL,
 CustomerKey INT NOT NULL,
 FullName NVARCHAR(150) NULL,
 EmailAddress NVARCHAR(50) NULL,
 BirthDate DATE NULL,
 MaritalStatus NCHAR(1) NULL,
 Gender NCHAR(1) NULL,
 Education NVARCHAR(40) NULL,
 Occupation NVARCHAR(100) NULL,
 City NVARCHAR(30) NULL,
 StateProvince NVARCHAR(50) NULL,
 CountryRegion NVARCHAR(50) NULL,
 Age AS
 CASE
 WHEN BirthDate IS NULL THEN NULL
 WHEN DATEDIFF(yy,BirthDate,CURRENT_TIMESTAMP) > 50
 THEN 'Older'
 WHEN DATEDIFF(yy,BirthDate,CURRENT_TIMESTAMP) > 40
 THEN 'Middle Age'
 ELSE 'Younger'
```

Page	Location	Description	Date corrected
		END CurrentFlag BIT NOT NULL DEFAULT 1, CONSTRAINT PK_Customers PRIMARY KEY (CustomerDwKey)); GO	
50-51	Step 2 code block	Should read: CREATE TABLE dbo.Customers (CustomerDwKey INT NOT NULL, CustomerKey INT NOT NULL, FullName NVARCHAR(150) NULL, EmailAddress NVARCHAR(50) NULL, BirthDate DATE NULL, MaritalStatus NCHAR(5) NULL, CREATE TABLE dbo.Products (ProductKey INT NOT NULL, ProductName NVARCHAR(50) NULL, Color NVARCHAR(15) NULL, Size NVARCHAR(50) NULL, SubcategoryName NVARCHAR(50) NULL, CategoryName NVARCHAR(50) NULL, CONSTRAINT PK_Products PRIMARY KEY (ProductKey)); GO	8/16/2013
51	Top section	The CASE statement's END clause is missing a trailing comma only in code files, not in the book itself. A corrected version of the code files is available from the book's webpage at http://go.microsoft.com/FWLink/?Linkid=260986	
60	Paragraph above Note reader aid	Reads: The query returns 6,343 rows and performs… Should read: The query returns 6,434 rows and performs…	8/16/2013
63	Second paragraph	Reads: There are three new catalog views you can use to gather information about columnstore indexes: * sys.column_store_index_stats * sys.column_store_segments * sys.column_store_dictionaries Should read: There are two new catalog views you can use to gather information about columnstore indexes: * sys.column_store_segments * sys.column_store_dictionaries	8/16/2013

Page	Location	Description		Date corrected
64	Quick Check reader aid	Reads: 2. No, you should use age compression only for data warehousing environments. Should read: 2. No, you should use page compression only for data warehousing environments.		8/16/2013
65	Table 2-5, second and third rows, Remarks column	The following text, which is currently in the CustomerKey row, should be moved to the FullName row: Concatenate FirstName and LastName from DimCustomer		9/25/2014
65	Table 2-5, sixth and seventh rows, Column name and Data type columns	Reads: MaritalStatus NCHAR(1) Gender NCHAR(1) Should read: MaritalStatus NCHAR(5) Gender NCHAR(5)		8/16/2013
67	Table 2-8, third and fourth rows, Remarks column	The following text, which is currently in the ProductKey row, should be moved to the DataKey row: OrderDateKey from FactInternetSales		9/25/2014
67	Exercise 1, step 4	Reads: SUBSTRING(CONVET(CHAR(8), FullDateAlternateKey,112), 5,2) Should read: FORMAT(MonthNumberYear,'00')		9/25/2014
74	Step 5	Reads: 5. Re-create the FactInternetSales table. Should read: 5. Re-create the InternetSales table.		8/16/2013
78	Lesson review, question 2, answer E	Reads: If you want to switch content from a nonpartitioned table to a partition of a partitioned table, what conditions must the nonpartitioned table meet? Should read: You have inserted data into an unpartitioned table and want to switch the content from this table ta a partition of a partitioned table. What conditions must the nonpartitioned table meet?		9/25/2014

Page	Location	Description	Date corrected
89	"Planning a Simple Data Movement" section, first sentence	Reads: To determine whether the and Export Wizard is the right tool for a particular data movement, ask yourself a few simple questions Should read: To determine whether the Import and Export Wizard is the right tool for a particular data movement, ask yourself a few simple questions	8/16/2013
99-100	Questions 1 and 2	The following sentence should be added at the end of questions 1 and 2: (Choose all that apply.)	8/16/2013
108	First question	Reads: SQL Server Development Tools Should read: SQL Server Data Tools	8/16/2013
124	Question 1, answer C	Reads: C. The Execute SQL Task Editor Should read: C. The Execute T-SQL Statement Task Editor	8/16/2013
128	Lesson 3, question 1, answer C	Reads: C. Incorrect: The Execute SQL Task Editor... Should read: C. Incorrect: The Execute T-SQL Statement Task Editor...	8/16/2013
133	First paragraph	The second occurrence of the phrase "mail servers" should be removed.	9/25/2014
150	Table 4-5, second row, Description column, first paragraph	The following sentence should be added to the end of the first paragraph: The SQLMOBILE provider is used for connections to SQL Server Compact Edition instances.	9/25/2014
150	Important reader aid	Should read: Important THE BULK INSERT TASK AND PERMISSIONS According to vendor documentation, the Bulk Insert task requires the user who is executing the SSIS package that contains this task to be a member of the sysadmin fixed server role. However, based on practical experience with SSIS 2012, it seems that the minimum set of permissions required to execute the task is far less than what is said in the documentation. The user executing an SSIS package, in which the Bulk Insert task is used, actually needs the following permissions: a) the user must be a member of the bulkadmin fixed server role; and b) must either have ALTER TABLE permissions on each destination table, or be a member of the db_ddladmin database role; and c) the user must have INSERT permissions on each destination table, or be a member of the db_datawriter database role.	2/25/2015

Page	Location	Description	Date corrected
180	Exam Tip reader aid	Reads: At run time, the data flow task builds an execution plan from the data flow, and the data flow engine executes the plan. Should read: At run time, the data flow task builds an execution plan from the data flow definition, and the data flow engine executes the plan.	8/16/2013
193	Step 16	Reads: 16. Execute the FillStageTables.dtsx package. Observe the execution to confirm successful completion of this exercise. Should read: 16. Execute the FillStageTables.dtsx package. Observe the execution to confirm successful completion of this exercise. If you get an error, please check if you have created a system DSN for the ODBC connection or try to set the project debugging property "Run64BitRuntime" to false.	9/25/2014
194	Step 1 code block	Should read: CREATE TABLE stg.CustomerInformation (PersonID INT NULL, EnglishEducation NVARCHAR(30) NULL, EnglishOccupation NVARCHAR(50) NULL, BirthDate DATE NULL, Gender NCHAR(5) NULL, MaritalStatus NCHAR(5) NULL, EmailAddress NVARCHAR(50) NULL);	8/16/2013
201	Table 5-4, sixth row, second column	Reads: Generates one or mode identical outputs... Should read: Generates one or more identical outputs…	8/16/2013

Page	Location	Description		Date corrected
209	Exercise 1, step 1	Should read: -- Customers dimension with a PK CREATE TABLE dbo.Customers (CustomerDwKey INT NOT NULL, CustomerKey INT NOT NULL, FullName NVARCHAR(150) NULL, EmailAddress NVARCHAR(50) NULL, BirthDate DATE NULL, MaritalStatus NCHAR(5) NULL, Gender NCHAR(5) NULL, Education NVARCHAR(40) NULL, Occupation NVARCHAR(100) NULL, City NVARCHAR(30) NULL, StateProvince NVARCHAR(50) NULL, CountryRegion NVARCHAR(50) NULL, Age AS CASE WHEN DATEDIFF(yy, BirthDate, CURRENT_TIMESTAMP) <= 40 THEN 'Younger' WHEN DATEDIFF(yy, BirthDate, CURRENT_TIMESTAMP) > 50 THEN 'Older' ELSE 'Middle Age' END, CurrentFlag BIT NOT NULL DEFAULT 1, CONSTRAINT PK_Customers PRIMARY KEY (CustomerDwKey)); GO		8/16/2013
210	Exercise 1, step 5	Reads: 5. Drag another OLE DB source adapter onto the workspace and rename it stgCustomer. In the OLE DB Source Editor, set the OLE DB connection manager to TK463 and select the stg.Customer table. Should read: 5. Drag another OLE DB source adapter onto the workspace and rename it stgCustomer. In the OLE DB Source Editor, set the OLE DB connection manager to TK463DW and select the stg.Customer table.		8/16/2013
212	Step 6	The following sentence should be added to the end of step 6: When prompted with the Input Output Selection dialog box, choose Lookup No Match Output from the Input drop-down list, and then click OK.		8/16/2013

Page	Location	Description	Date corrected
214	Step 14	Reads: Drag a Derived Column transformation from the SSIS Toolbox and connect the output arrow from the Check Customer Exists Lookup transformation to the new Derived Column transformation. Rename the Derived Column transformation Calc Columns. Should read: Drag a Derived Column transformation from the SSIS Toolbox and connect the output arrow from the Check Customer Exists Lookup transformation to the new Derived Column transformation. When prompted with the Input Output Selection dialog box, choose Lookup Match Output from the Input drop-down list, and then click OK. Rename the Derived Column transformation Calc Columns.	8/16/2013
226	Second paragraph, last sentence	The following sentence should be deleted: (If you are not familiar with the T-SQL MERGE statement, see Chapter 2.)	8/16/2013
227	Exercise 1, lesson 3, step 2	Reads: Click the Data Flow tab at the top of the SSIS Designer. Select the stgCustomer, stgPerson, Sort, Sort 1, Merge Join, and Sort 3 components, and press Delete on the keyboard. You will now replace the existing logic of sorting and merging data by doing this on the database layer. Should read: Click the Data Flow tab at the top of the SSIS Designer. Select the stgCustomer, stgPerson, Sort (sort step after stgPerson), Sort 1 (sort step after stgCustomer), Merge Join and Sort 3 (sort step after Merge Join) components, and press Delete on the keyboard. You will now replace the existing logic of sorting and merging data by doing this on the database layer.	8/16/2013
230	Table 5-10	Reads: EnglishEducation Should read: Education Reads: EnglishOccupation Should read: Occupation	8/16/2013
232	Question 2	Reads: How many Data Source adapters do you need? (Choose all that apply.) Should read: How many Data Source adapters could you use? (Choose all that apply.)	8/16/2013

Page	Location	Description		Date corrected
232	Lesson Review, question 1	Reads: Which data flow transformation would you use if had to combine data from two different database tables that exist on two different servers? Should read: Which data flow transformation would you use if you had to join data from two different database tables that exist on two different servers?		8/16/2013
235	Lesson 2, question 3	Reads: Correct Answers: B and D Should read: Correct Answers: A and D		8/16/2013
236	Last sentence	Reads: Then you would need to unpivot the columns for months to get a row for each month, and use the Lookup task to get appropriate surrogate keys from existing the dimension, and use the ODBC Destination adapter to write the data to the new fact table. Should read: Then you would need to unpivot the columns for months to get a row for each month, and use the Lookup task to get appropriate surrogate keys from the existing dimension, and use the ODBC Destination adapter to write the data to the new fact table.		8/16/2013
244	Table 6-1, fifth row, second column	The following sentence should be added after the existing text in the Description column: You can create additional namespaces for user-defined variables and change the name of the User namespace, but you cannot change the name of the System namespace, add variables to the System namespace, or assign system variables to a different namespace.		
262	Exercise 1, step 6	Reads: ...(between 1 for Monday and 7 for Sunday)... Should read: ...(between 1 for Monday and 7 for Sunday in most languages)...		9/25/2014
263	Step 3, Listing 6-2	Reads: DefaultBufferMaxRow Should read: DefaultBufferMaxRows		8/16/2013

Page	Location	Description	Date corrected
263	Exercise 2, after step 6	Add the following Note reader aid: Note LANGUAGE SETTINGS AND DATE/TIME FORMATS Depending on your environment, either Monday or Sunday is used as the first day of the week; by default, this is determined by the operating system regional settings. If Monday is used as the first day of the week (for instance, if the language used in your environment is British English) then Saturday and Sunday are represented by numbers 6 or 7, respectively. On the other hand, if Sunday is used as the first day of the week (for instance, if the language used in your environment is U.S. English) then Saturday and Sunday are represented by numbers 7 and 1, respectively. Please, make the appropriate adjustments to the expressions shown in Listings 6-2 and 6-3 to reflect your actual environment settings. You can find more information about SQL Server languages by inspecting the sys.syslanguages catalog view.	9/25/2014
269	Sixth paragraph, last sentence	Reads: The name of the parent package variable must match the name of the corresponding variable in the parent package. Should read: The name of the parent package variable must match the name of the corresponding variable in the child package.	8/16/2013
272	Exercise 2, step 5	Reads: ...create a new, project-scoped integer... Should read: ...create a new, package-scoped integer...	9/25/2014
274	Step 4	Should read: 4. Change the Evaluation operation constraint option from "Constraint" to "Expression and Constraint", leave "Success" as the evaluation value, then click the ellipsis (…) at the right of the Expression text box to open the Expression Builder dialog box, and build the expression shown in Listing 6-5.	8/16/2013
276	Question 1	The following sentence should be added to the end of question 1: (Choose all that apply.)	8/16/2013
295	Lesson 1, step 17	Reads: 17. Observe the created transformations and execute the package. Should read: 17. Observe the created transformations and execute the package. If you get an error, please change the data access mode on the destination adapter to "Table/View - fast load". There is an issue, because we are using a sequencer to populate the surrogate keys and the SCD wizard sets the destination by default to an access mode that does not allow a trigger with a sequencer on the destination table.	9/25/2014

Page	Location	Description	Date corrected
296	Step 21	Reads: NCHAR(5) Should read: NVARCHAR(5)	9/25/2014
306	Fourth sentence from bottom	Reads: If three UPDATE statements ware applied to the same row, you will get back three records. Should read: If three UPDATE statements were applied to the same row, you will get back three records.	8/16/2013
310	Steps 4 and 5	Steps 4 and 5 should be switched so that the instruction in step 5, "Ensure that the SQL Server Agent is running," comes before the instruction in step 4.	
310	Step 9	Reads: In the Configure ADO.NET Connection Manager dialog box, select New and create the connection to the TK463 database. Should read: In the Configure ADO.NET Connection Manager dialog box, select New and create the connection to the TK463DW database.	8/16/2013
312	Exercise 2, step 3	Reads: Add a new ADO.NET connection to the TK463 database. Should read: Add a new ADO.NET connection to the TK463DW database.	8/16/2013
314	Step 21	Reads: Inside the IncLoadStgSalesHeader package, add another OLE DB destination adapter. Connect the CDC Splitter transformation with the new OLE DB destination adapter and, in the Input Output Selection dialog box, select DeleteOutput in the Output drop-down list. Should read: Inside the IncLoadStgSalesHeader package, add another OLE DB destination adapter. Connect the CDC Splitter transformation with the new OLE DB destination adapter.	9/25/2014

Page	Location	Description	Date corrected
315	Step 33	The following new steps should be added after step 33 and before step 34 (so that step 34 becomes step 37): 34. Notice that there is a missing step to delete the data in stg.tmpDeleteSalesOrderHeader and stg.tmpUpdateSalesOrderHeader tables since each time you will run the package, additional data will be inserted in to these two tables. Drag an Execute SQL task from the SSIS Toolbox onto the control flow area. Rename the task "Truncate tables" and double-click the new task to open the editor. 35. Change the Connection property to use the TK463DW connection. In the SQL Statement property of the Execute SQL Task Editor dialog box, type the following code: TRUNCATE TABLE stg.tmpUpdateSalesOrderHeader; TRUNCATE TABLE stg.tmpDeleteSalesOrderHeader; 36. Set the Truncate tables Execute SQL task as the first task in the control flow by connecting it with the CDC Start CDC Control Task. 37. Write different INSERT, UPDATE, and DELETE statements against the stg.CDCSalesOrderHeader table and test your package.	9/25/2014
317	"Using Error Flows" section, first paragraph, third sentence	Reads: In addition to the data paths, however,, there are also error paths. Should read: In addition to the data paths, however, there are also error paths.	8/16/2013
320	Step 11	Reads: OLD DB Connection Manager Should read: OLE DB Connection Manager	8/16/2013
321	Step 18	Reads: In the Flat File Destination Editor, select the Errors For Flat File connection manager and click the Mappings tab to observe how columns will be mapped. Should read: In the Flat File Destination Editor, select the ErrorRows For Flat File connection manager and click the Mappings tab to observe how columns will be mapped.	8/16/2013
321	Step 19	Reads: Execute the package. Notice that five rows were transferred through the error path, and the package ran successfully. Should read: Execute the package. Notice that fifteen rows were transferred through the error path, and the package ran successfully.	8/16/2013
324	Lesson 3, question 1, answer A	Reads: A. Correct: All source adapters have an error flow. Should read: A. Correct: The OLE DB source adapter has an error flow.	8/16/2013

Page	Location	Description	Date corrected
331	Snapshot reader aid	Reads: Snapshot The data read within a transaction will not reflect changes made by other simultaneous transactions. The transaction uses the data row versions that exist when the transaction begins. No locks are placed on the data when it is read. Should read: Snapshot The data read within a transaction will not reflect changes made by other simultaneous transactions. The transaction uses the data row versions that exist when the transaction begins. No locks are placed on the data when it is read. The Snapshot value of the IsolationLevel property is incompatible with package transactions. Therefore, you cannot use the IsolationLevel property to set the isolation level of package transactions to Snapshot. Instead, use an SQL query to set package transactions to Snapshot.	8/16/2013
331	Third bulleted item	Reads: Chaos Behaves the same way as ReadUncommitted, but checks the isolation level of other pending transactions during a write operation so that transactions with more restrictive isolation levels are not overwritten. Should read: Chaos Behaves the same way as ReadUncommitted, but checks the isolation level of other pending transactions during a write operation so that transactions with more restrictive isolation levels are not overwritten. Furthermore, it is not supported on the SQL Server platform, is not an ANSI standard Isolation level, and cannot be rolled back.	8/16/2013
332	Note reader	Reads: Note Setting Isolation Levels of a Transaction in SSIS The IsolationLevel property in SSIS objects applies only when you explicitly open a transaction inside SSIS by setting the TransactionOption property to Required. Should read: Note Setting Isolation Levels of a Transaction in SSIS The IsolationLevel property in SSIS objects applies only when you explicitly open a transaction inside SSIS by setting the TransactionOption property to Required and the parent container has not already started a transaction.	8/16/2013
332	Quick Check Answer reader aid	Reads: You should remove the legacy task from the sequence container or set the TransactionOption property to Not Supported. Should read: You should remove the legacy task from the sequence container or set the TransactionOption property to NotSupported.	8/16/2013

Page	Location	Description		Date corrected
335	Step 21	Reads: Save the package by clicking the Save button on the toolbar. Should read: Save the package by clicking the Save button on the toolbar. Check if the MSDTS is running by looking at windows services if Distributed Transaction Coordinator service is running.		8/16/2013
340	Step 1	Reads: 1. If necessary, start SQL Server Data Tools, open the TK 463 Chapter 8 project, and open the package. Should read: 1. Start SQL Server Data Tools, open the TK 463 Chapter 8 project under C:\TK463\Chapter08\Lesson2\Starter\, and then open the TestCheckpoints.dtsx package.		9/25/2014
340	Step 9	Reads: In Windows Explorer, delete the checkpoint.xml file under C:\TK463\Chapter08 \Lesson2\Starter\. Should read: If you did not finish the previous exercise, please check if the checkpoint.xml file exists and then delete it (in Windows Explorer, delete the checkpoint.xml file under C:\TK463\Chapter08\Lesson2\Starter\ ).		9/25/2014
341	Step 13	Reads: 12. Click OK to close the SQL Task Editor, and rerun the package. Notice that only the last task was run, because you are using checkpoints. Should read: 12. Click OK to close the SQL Task Editor, and rerun the package. Notice that only the last task was run, because you are using checkpoints. Check if the checkpoint.xml file was automatically deleted and rerun the package.		9/25/2014
345	Step 12	Reads: 12. Drag a data flow task from the SSIS Toolbox onto the Sequence Container. Should read: 12. Drag a data flow task from the SSIS Toolbox onto the Event Handler design area.		8/16/2013

Page	Location	Description	Date corrected
345	Step 14	Reads: Set the OLE DB connection manager to TK463DW and, under Data Access Mode, select the stg.Customer table. Should read: Select TK463DW in the OLE DB Connection Manager list. From the Data Access Mode drop-down list, select Table Or View, and select the stg.Customer table in the Name Of The Table Or The View drop-down list.	8/16/2013
346	Step 17	The following step should be deleted: 17. On the Mapping tab of the OLE DB Destination Editor, check to make sure that all destination columns are mapped to input columns. Click OK.	9/25/2014
367	Question 3, answer D	Reads: The Lookup task cache mode propert Should read: The Lookup transformation SqlCommand property	8/16/2013
376	Step 1	Reads: If necessary, start SQL Server Data Tools, open the TK 463 Chapter 9 project, and open the FillStageTables_1.dtsx package. Should read: Start SQL Server Data Tools, open the TK 463 Chapter 9 project, and open the FillStageTables_1.dtsx package.	8/16/2013
446	Step 2	Reads: ... execute only the part for of the script for the Lesson 2 practice. Should read: ... execute only the part of the script for the Lesson 2 practice.	9/25/2014
449	Step 4	The following Note reader aid should be added after Step 4: Note SENSITIVE INFORMATION STORED IN SSIS PROJECTS When sensitive information is stored inside an SSIS project, it is encrypted by default using a user key based on the credentials of the user who created the project. As a result, other users might not be able to decrypt this information when viewing or editing the project or when they create a project deployment file; this way any sensitive data will remain protected. BIDS and other tools of the SSIS toolset might issue warnings about not being able to encrypt or decrypt sensitive data even if the project definitions do not actually contain any encrypted data, when used on projects created by another user. During the creation of the project deployment file, or when performing other exercises in this training kit, you might receive such warnings; however, because no encrypted data is actually present in the project definitions, you can safely ignore them.	9/25/2014

Page	Location	Description	Date corrected
475	Step 4	Reads: PowerShell_Example.ps2 Should read: PowerShell_Example.ps1	8/16/2013
481	"Principals" section, first paragraph	Reads: catalog.configure_server Should read: catalog.configure_catalog	8/16/2013
507	Quick Check Answer reader aid	Reads: On the SSISDB, project, and package levels Should read: On the SSISDB catalog, folder, project, and package levels	8/16/2013
508	Exercise 1, step 1	Reads: 1. If you are missing the database objects from Chapter 5, "Designing and Implementing Data Flow," execute the needed SQL code from that chapter to have all the stage and dimension tables available in the TK463DW database. Should read: 1. If you are missing the database objects from Chapter 5, "Designing and Implementing Data Flow," execute the needed SQL code from that chapter to have all the stage and dimension tables available in the TK463DW database. Start SSMS and connect to your SQL Server instance. Open a new query window by clicking the New Query button. Select the TK463DW database in the database drop-down list. Create the stg.SalesTerritory staging table by executing the following SQL statement. CREATE TABLE stg.SalesTerritory (TerritoryID INT NULL, Name NVARCHAR(50) NULL, CountryRegionCode NVARCHAR(10) NULL, [Group] NVARCHAR(50) NULL, ModifiedDate DATETIME NULL);	8/16/2013

Page	Location	Description	Date corrected
508	Exercise 1, step 3	Reads: If you are missing the database objects from Chapter 5, "Designing and Implementing Data Flow," execute the needed SQL code from that chapter to have all the stage and dimension tables available in the TK463DW database. Should read: If you are missing the database objects from Chapter 5, "Designing and Implementing Data Flow," execute the needed SQL code from that chapter to have all the stage and dimension tables available in the TK463DW database. Execute also the FillStageTables package to load the needed stage tables with sample data.	9/25/2014
515, 516, 526, 782	Multiple locations	Reads: DefaultMaxBufferRows Should read: DefaultBufferMaxRows	
538	Exercise 1, step 2, last bulleted item	The last bulleted item on the page, where DimDate appears for the second time, should be deleted.	8/16/2013
559	Question 1	The following should be deleted: (Choose all that apply)	9/25/2014
573	Exercise 1, step 3, last bulleted item	The last item in the bulleted list is duplicated and should be deleted: DimDate	9/25/2014
579	Exercise 1, first paragraph, last sentence	Reads: If you are using Windows Vista or Windows 7, use Control Panel | Programs And Features instead of Server Manager. Should read: If you are using Windows Vista or Windows 7, use Control Panel | Programs And Features and the Turn Windows Features on or off instead of Server Manager.	9/25/2014
594	Step 10	Reads: At the bottom of the page, click the Add Leaf Attribute button. Make the new attribute domain based and use the CountryRegion entity for the domain values. Do not change the display pixel width or enable change tracking. Save the attribute and then save the entity. Should read: At the bottom of the page, click the Add Leaf Attribute button. Name the attribute CountryRegion. Make the new attribute domain based and use the CountryRegion entity for the domain values. Do not change the display pixel width or enable change tracking. Save the attribute and then save the entity.	8/16/2013

Page	Location	Description		Date corrected
617	Lesson 2, paragraph above step 1	Reads: mdm.udpSecurityMemberProcessRebuildModel Should read: mdm.udpSecuritySetAdministrator		9/25/2014
628	Quick Check Answer reader aid	Reads: Quick Check Answer Use Excel 2012 with MDS Add-in for Excel to update batches of data in an MDS mod" Should read: Quick Check Answer Use Excel 2010 with MDS Add-in for Excel to update batches of data in an MDS model.		8/16/2013
644	Step 18	Reads: \p{L}+\d\d@ADVENTURE -WORKS \.COM Should read: \p{L}+\d@ADVENTURE -WORKS \.COM		8/16/2013
660	Question 2, answer D	Reads: Check the standard distribution of the values. Should read: Check the standard deviation of the values.		8/16/2013
674	Step 22	Reads: On the Create Test Set page, Should read: On the Create Testing Set page,		9/25/2014

Page	Location	Description	Date corrected
727	Exercise 3, step 1	Step 1 should be replaced with the following four steps. The current step 2 then become step 5. Reads: 1. Switch the solution configuration to Release, and build the TK463.CalculateCheckSum project again. Should read: 1. Before deploying the component, it needs to be released. If you were able to build the project successfully in the preceding exercise, you can now complete the building process by creating a release build. However, in order for the component to be used in production, it also needs to be digitally signed; for instance, using a strong name key. 2. In the Solution Explorer pane, right-click the TK463.CalculateCheckSum project, and in the context menu select Properties… 3. In the project properties editor, switch to the Signing tab at the bottom left. Check the Sign the assembly option; this will allow the component to be digitally signed. Expand the Choose a strong name key file list box, select Browse…, and in the Select file dialog navigate to the C:\TK463\Chapter19\Lesson3\Solution\TK463 Chapter 19\TK463.CalculateCheckSum folder. In the folder, locate the TK463.snk file, select it, and confirm the selection by clicking Open. 4. Save the solution, and close the project properties editor. Switch the solution configuration to Release, and build the TK463.CalculateCheckSum project again.	7/16/2015
728	Exercise 3, step 5	The following sentence should be added at the end of step 5: If in your particular environment elevated privileges are required to run the gacutil utility, right-click the file in Windows Explorer and then click Run as administrator to perform the operation.	9/25/2014
730	Questions 1, 2, and 3	The following text should be added to questions 1, 2 and 3: (Choose all that apply.)	9/25/2014
752	Exercise 2, step 4	Reads: ...CustomersClean table. Should read: ...CustomersDQSMatch view.	9/25/2014

Page	Location	Description	Date corrected
763	Exercise 2	The following paragraph should be added at the end of the intro text for the exercise 2, before the first step: If you did not enable MDS integration with DQS yet, the Data Quality group of the Master Data tab might not be activated within Excel. If this is the case, open the Master Data Services Configuration Manager console. You can find it in All Programs | Microsoft SQL Server 2012 | Master Data Services. Click "Web Configuration" in the left pane. In the Website drop-down in the right pane, select MDS. Lastly, click the "Enable integration with Data Quality Services" button in the bottom of the right pane. Click the "Exit" button. "	9/25/2014