



70-767^{Q&As}

Implementing a SQL Data Warehouse

Pass Microsoft 70-767 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4lead.com/70-767.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

You are designing the data warehouse to import data from three different environments. The sources for the data warehouse will be loaded every hour.

Scenario A includes tables in a Microsoft Azure SQL Database:

1.

Millions of updates and inserts occur per hour

2.

A periodic query of the current state of rows that have changed is needed.

3.

The change detection method needs to be able to ignore changes to some columns in a table.

4.

The source database is a member of an AlwaysOn Availability group.

Scenario B includes tables with status update changes:

1.

Tracking the duration between workflow statuses.

2.

All transactions must be captured, including before/after values for UPDATE statements.

3.

To minimize impact to performance, the change strategy adopted should be asynchronous.

Scenario C includes an external source database:

1.

Updates and inserts occur regularly.

2.

No changes to the database should require code changes to any reports or applications.

3.

Columns are added and dropped to tables in the database periodically. These schema changes should not require any interruption or reconfiguration of the change detection method chose.

4.

Data is frequently queried as the entire row appeared at a past point in time.



All tables have primary keys.

You need to load each data source. You must minimize complexity, disk storage, and disruption to the data sources and the existing data warehouse.

Which change detection method should you use for each scenario? To answer, drag the appropriate loading methods to the correct scenarios. Each source may be used once, more than once, or not at all. You may need to drag the split bar

between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Loading methods	Scenario	Loading method
Change Tracking	A	
Change Data Capture	B	
System-Versioned Temporal Table	C	

Correct Answer:

Answer Area

Loading methods	Scenario	Loading method
	A	System-Versioned Temporal Table
	B	Change Tracking
	C	Change Data Capture

Box A: System-Versioned Temporal Table

System-versioned temporal tables are designed to allow users to transparently keep the full history of changes for later analysis, separately from the current data, with the minimal impact on the main OLTP workload.

Box B: Change Tracking

Box C: Change Data Capture



Change data capture supports tracking of historical data, while that is not supported by change tracking.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/track-changes/track-data-changes-sql-server>

<https://docs.microsoft.com/en-us/sql/relational-databases/tables/temporal-table-usage-scenarios>

QUESTION 2

You have a Microsoft SQL Server data warehouse instance that supports several client applications.

The data warehouse includes the following tables: Dimension.SalesTerritory, Dimension.Customer, Dimension.Date, Fact.Ticket, and Fact.Order. The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. The Fact.Order table is optimized for weekly reporting, but the company wants to change it to daily. The Fact.Order table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows data loading.

All data in the data warehouse is stored on a shared SAN. All tables are in a database named DB1. You have a second database named DB2 that contains copies of production data for a development environment. The data warehouse has grown and the cost of storage has increased. Data older than one year is accessed infrequently and is considered historical.

You have the following requirements:

Implement table partitioning to improve the manageability of the data warehouse and to avoid the need to repopulate all transactional data each night. Use a partitioning strategy that is as granular as possible.

Partition the Fact.Order table and retain a total of seven years of data.

Partition the Fact.Ticket table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed.

Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.

Incrementally load all tables in the database and ensure that all incremental changes are processed.

Maximize the performance during the data loading process for the Fact.Order partition.

Ensure that historical data remains online and available for querying.

Reduce ongoing storage costs while maintaining query performance for current data.

You are not permitted to make changes to the client applications.

You need to configure data loading for the tables.

Which data loading technology should you use for each table? To answer, select the appropriate options in the answer area.

Hot Area:



Answer area

Table	Technology
Dimension.SalesTerritory	<ul style="list-style-type: none">Change Data Capture (CDC)Change TrackingTemporal tableMicrosoft SQL Server snapshot replication
Dimension.Customer	<ul style="list-style-type: none">Change Data Capture (CDC)Change TrackingTemporal tableMicrosoft SQL Server snapshot replication
Dimension.Date	<ul style="list-style-type: none">Change Data Capture (CDC)Change TrackingTemporal tableMicrosoft SQL Server snapshot replication

Correct Answer:

Answer area

Table	Technology
Dimension.SalesTerritory	<ul style="list-style-type: none">Change Data Capture (CDC)Change TrackingTemporal tableMicrosoft SQL Server snapshot replication
Dimension.Customer	<ul style="list-style-type: none">Change Data Capture (CDC)Change TrackingTemporal tableMicrosoft SQL Server snapshot replication
Dimension.Date	<ul style="list-style-type: none">Change Data Capture (CDC)Change TrackingTemporal tableMicrosoft SQL Server snapshot replication

Scenario: The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated



Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.

Box 1: Change TrackingBox 2: Change TrackingBox 3: Temporal Table

Temporal Tables are generally useful in scenarios that require tracking history of data changes.

We recommend you to consider Temporal Tables in the following use cases for major productivity benefits.

* Slowly-Changing Dimensions

Dimensions in data warehousing typically contain relatively static data about entities such as geographical locations, customers, or products.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/tables/temporal-table-usage-scenarios>

QUESTION 3

A SQL Server Integration Services (SSIS) project has been deployed to the SSIS catalog. The project includes a project Connection Manager to connect to the data warehouse. The SSIS catalog includes two Environments:

Development

QA

Each Environment defines a single Environment Variable named **ConnectionString** of type string. The value of each variable consists of the connection string to the development or QA data warehouses.

You need to be able to execute deployed packages by using either of the defined Environments.

Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Select and Place:

The question interface displays a list of six actions in yellow boxes on the left, and a vertical blue bar on the right with a right-pointing arrow, representing the answer area. A large green watermark 'www.Pass4Lead.com' is overlaid diagonally across the interface.

- In the deployed project, set the **ConnectionString** property of the Connection Manager to use the **ConnectionString** Environment Variable.
- In the deployed package, set the **ConnectionString** property of the Connection Manager to use the **ConnectionString** Environment Variable.
- Add the Environment reference to the deployed project.
- Deploy the project to the **Development** and **QA** Environments.
- Add permissions to the **Development** and **QA** Environments.
- Execute the deployed package by using the required Environment.



Correct Answer:

The screenshot shows a deployment wizard with the following steps:

- In the deployed package, set the **ConnectionString** property of the Connection Manager to use the **ConnectionString** Environment Variable.
- Deploy the project to the **Development** and **QA** Environments.
- Add permissions to the **Development** and **QA** Environments.
- Add the Environment reference to the deployed project.
- In the deployed project, set the **ConnectionString** property of the Connection Manager to use the **ConnectionString** Environment Variable.
- Execute the deployed package by using the required Environment.

QUESTION 4

You are developing a SQL Server Integration Services (SSIS) package to implement an incremental data load strategy. The package reads data from a source system that uses the SQL Server change data capture (CDC) feature.

You have added a CDC Source component to the data flow to read changed data from the source system.

You need to add a data flow transformation to redirect rows for separate processing of insert, update, and delete operations.

Which data flow transformation should you use?

- A. DQS Cleansing
- B. Merge Join
- C. Pivot
- D. Conditional Split

Correct Answer: D

QUESTION 5

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series.

Information and details provided in a question apply only to that question.



You are developing a Microsoft SQL Server Integration Services (SSIS) package.

You are importing data from databases at retail stores into a central data warehouse. All stores use the same database schema.

The query being executed against the retail stores is shown below:

```
SELECT *
FROM dbo.Sales
WHERE SalesDate >= CAST(date, GETDATE() -1)
ORDER BY ID
```

The data source property named IsSorted is set to True. The output of the transform must be sorted.

You need to add a component to the data flow.

Which SSIS Toolbox item should you use?

- A. CDC Control task
- B. CDC Splitter
- C. Union All
- D. XML task
- E. Fuzzy Grouping
- F. Merge
- G. Merge Join

Correct Answer: C

QUESTION 6

You are designing a complex SQL Server Integration Services (SSIS) project that uses the Project Deployment model.

The project will contain between 15 and 20 packages. All the packages must connect to the same data source and destination.

You need to define and reuse the connection managers in all the packages by using the least development effort.

What should you do?

- A. Copy and paste the connection manager details into each package.
- B. Implement project connection managers.
- C. Implement package connection managers.
- D. Implement parent package variables in all packages.

Correct Answer: B

**QUESTION 7**

You are developing a Microsoft SQL Server Integration Services (SSIS) package that loads a data warehouse. You need to inspect the data that is being processed by the package. What should you do first?

- A. Set a break point on the Control Flow path.
- B. Enable SQL Trace.
- C. Enable logging on the Data Flow path.
- D. Enable a data viewer on the Data Flow path.

Correct Answer: A

QUESTION 8

You are developing a SQL Server Integration Services (SSIS) package.

The package sources data from an HTML web page that lists product stock levels. You need to implement a data flow task that reads the product stock levels from the HTML web page.

Which data flow source should you use?

- A. Raw File source
- B. XML source
- C. Custom source component
- D. Flat File source

Correct Answer: C

QUESTION 9

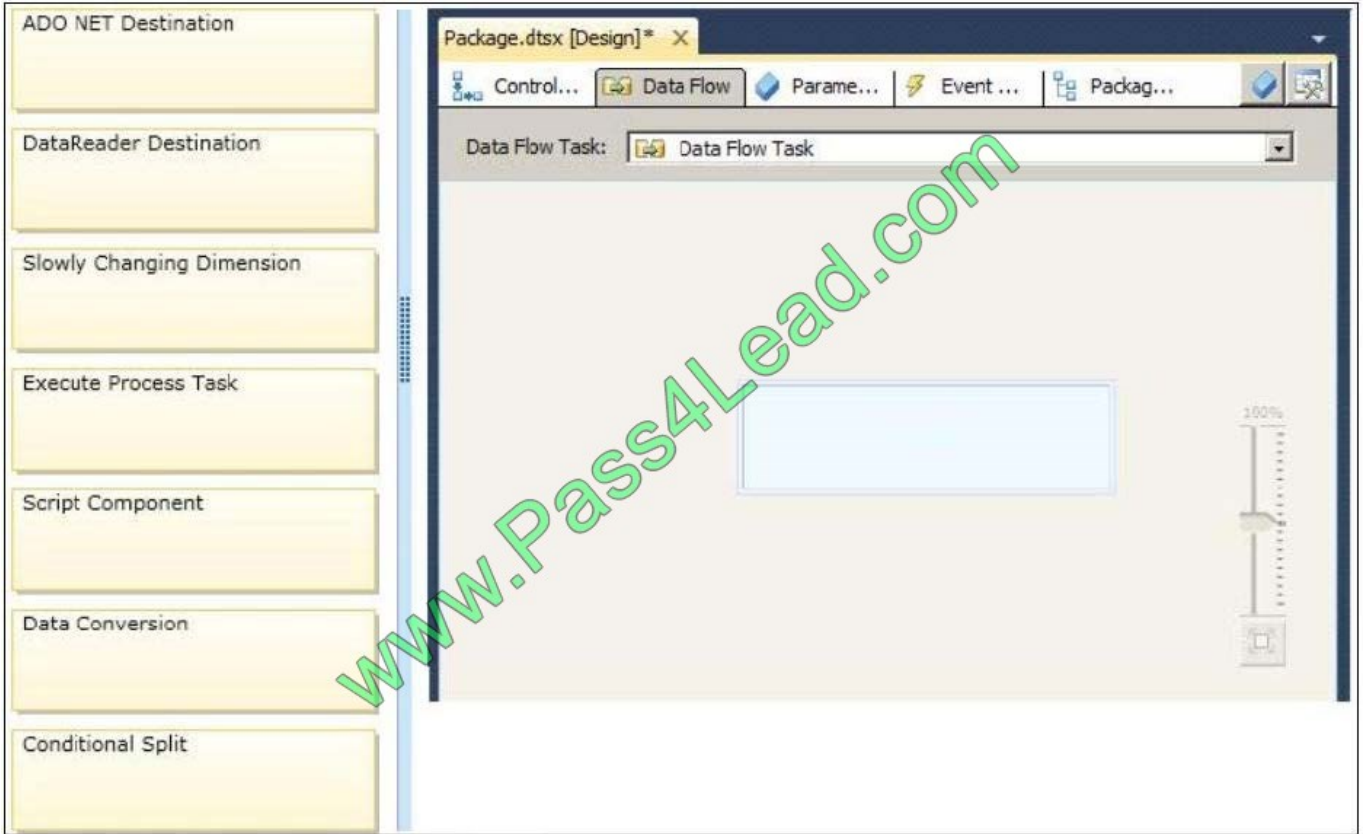
You are developing a SQL Server Integration Services (SSIS) package.

The package uses custom functionality that accesses a SQL Server database. The custom functionality must be implemented by using Language Integrated Query (LINQ).

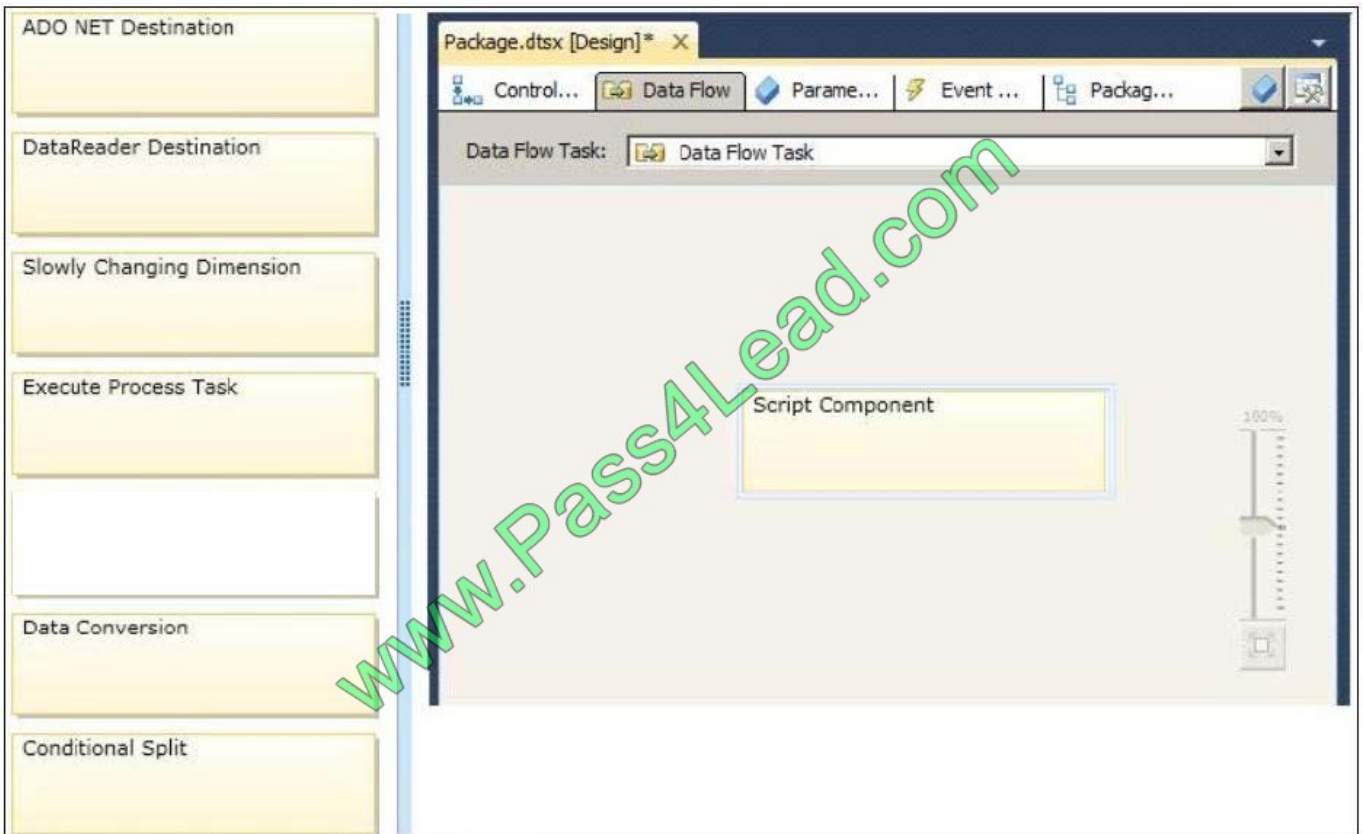
You need to ensure that the LINQ code can be debugged at design time.

What should you select from the SSIS Toolbox? (To answer, drag the appropriate item to the correct location in the answer area.)

Select and Place:



Correct Answer:



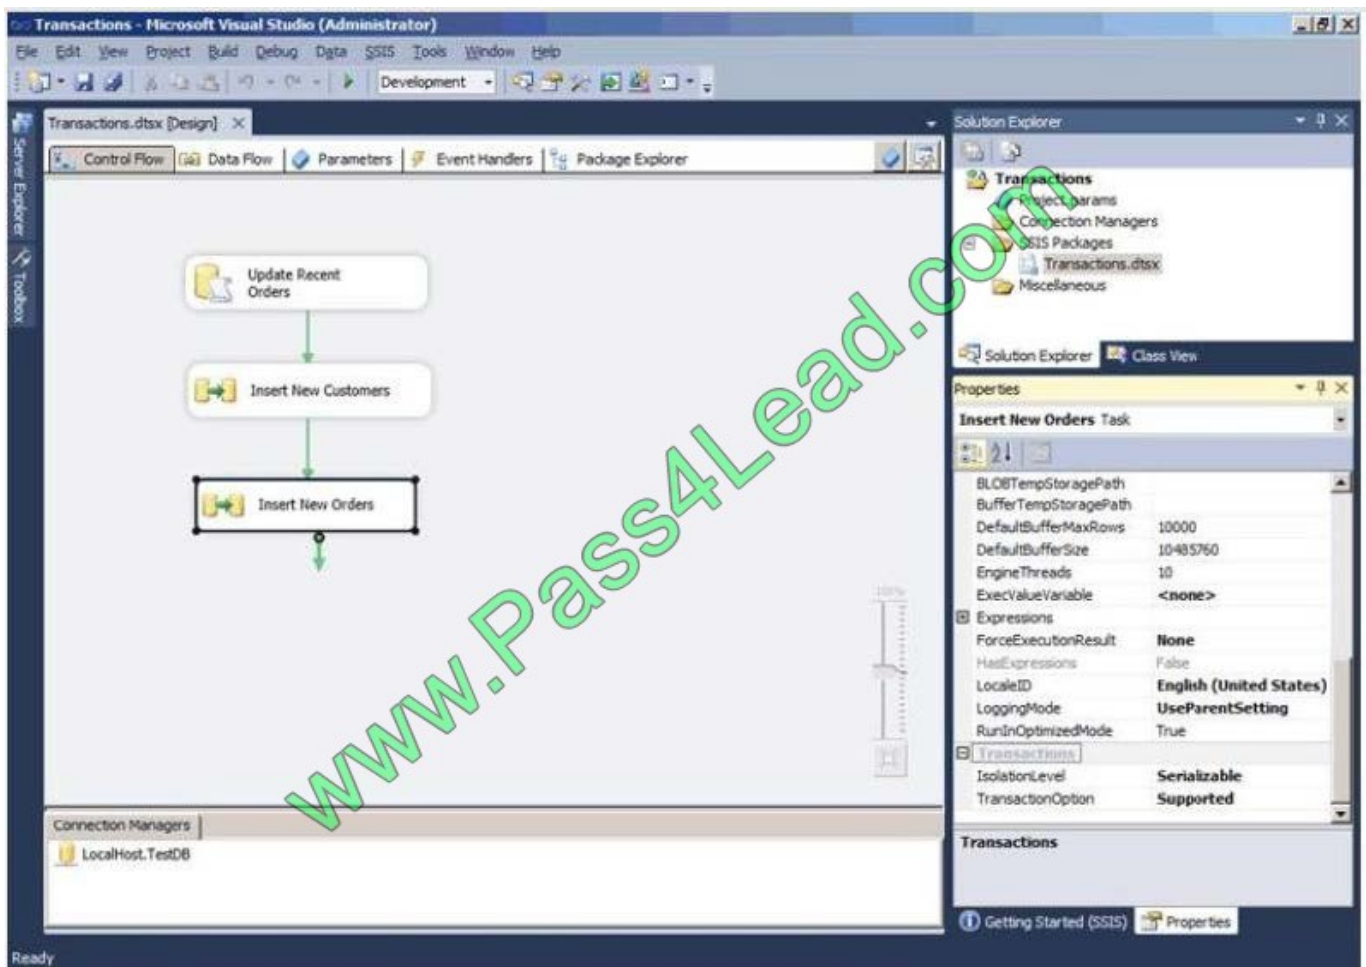


QUESTION 10

You are designing a package control flow. The package moves sales order data from a SQL Azure transactional database to an on-premise reporting database. The package will run several times a day, while new sales orders are being

added to the transactional database.

The current design of the package control flow is shown in the answer area. (Click the Exhibit button.)



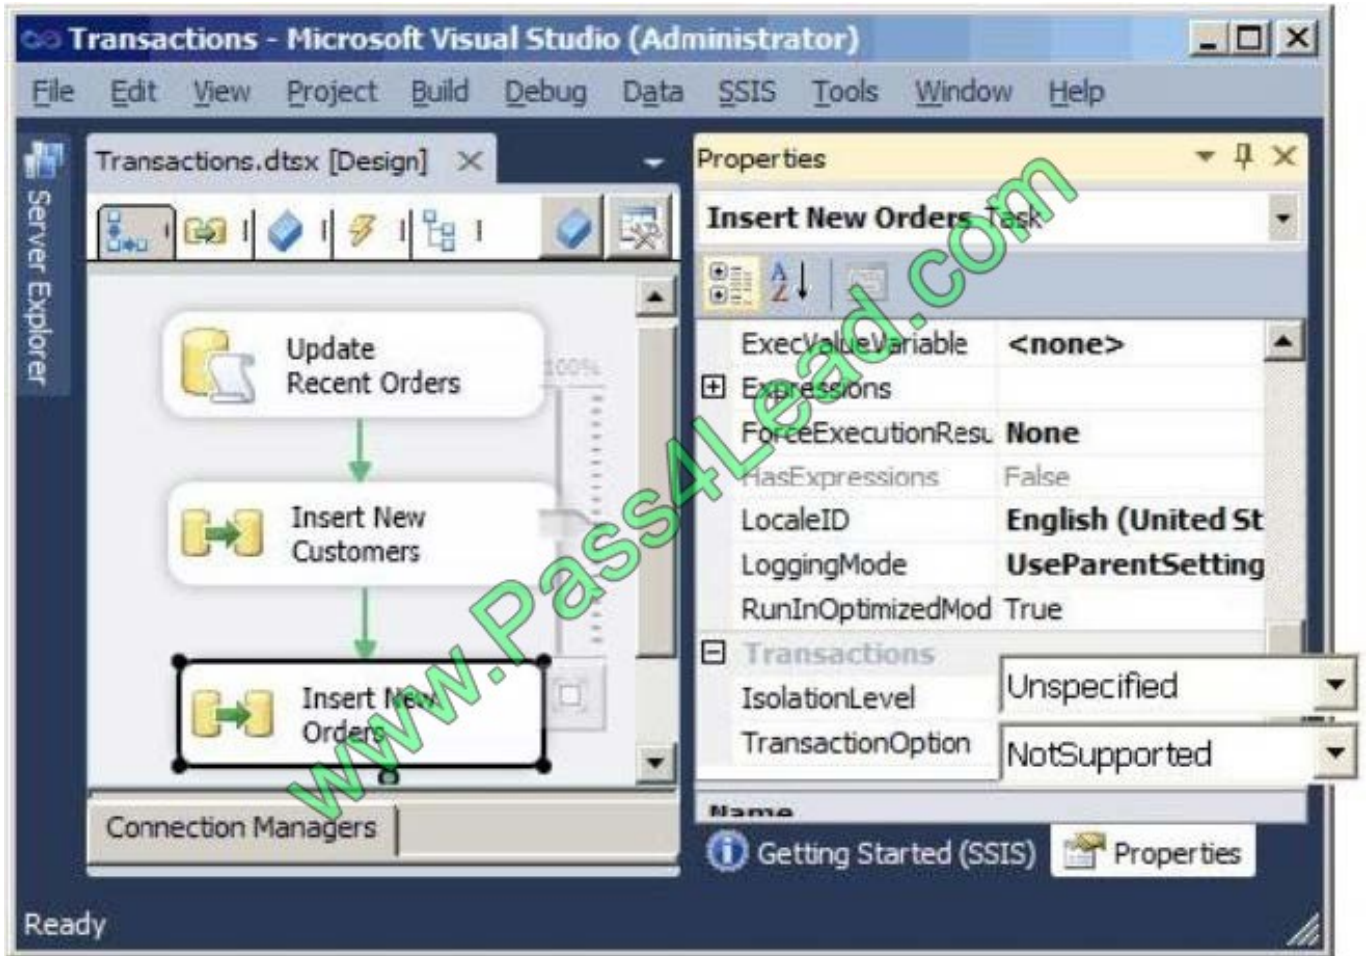
The Insert New Orders Data Flow task must meet the following requirements:

Usage of the tempdb database should not be impacted.

Concurrency should be maximized, while only reading committed transactions.

If the task fails, only that task needs to be rolled back.

You need to configure the Insert New Orders Data Flow task to meet the requirements.



How should you configure the transaction properties? To answer, select the appropriate setting or settings in the answer area.

Correct Answer:

IsolationLevel = ReadCommitted. TransactionOption = Required

QUESTION 11

Your company has several line-of-business applications. The applications use a server that has SQL Server installed and contains several databases.

You need that the business applications can access curated and validated data from the databases.

Which features should you deploy to the server?

- A. Data Quality Services (DQS)
- B. SQL Server Analysis Services (SSAS) dimension hierarchies.
- C. SQL Server Integration Services (SSIS)



D. Master Data Services (MDS) subscriptions.

Correct Answer: C

QUESTION 12

You are designing an enterprise star schema that will consolidate data from three independent data marts. One of the data marts is hosted on SQL Azure.

Most of the dimensions have the same structure and content. However, the geography dimension is slightly different in each data mart.

You need to design a consolidated dimensional structure that will be easy to maintain while ensuring that all dimensional data from the three original solutions is represented.

What should you do?

- A. Create a junk dimension for the geography dimension.
- B. Implement change data capture.
- C. Create a conformed dimension for the geography dimension.
- D. Create three geography dimensions.

Correct Answer: C

[Latest 70-767 Dumps](#)

[70-767 Practice Test](#)

[70-767 Exam Questions](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.pass4lead.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © pass4lead, All Rights Reserved.