



TestKingonline.com

No Pass No Pay!

MCSE, CCNA, CCNP, OCP, CIW, JAVA, Sun Solaris, Checkpoint

World No 1 Cert Exams

70-446

Congratulations!

You have purchased a TestKingonline. Study Guide.

This study guide is a complete collection of questions and answers that have been developed by our professional & certified team.

You must study the contents of this guide properly in order to prepare for the actual certification test. The average time that we would suggest you for studying this study guide is approximately 10 to 20 hours and you will surely pass your exam. We guarantee it!

GOOD LUCK!

DISCLAIMER

This study guide and/or material is not sponsored by, endorsed by or affiliated with Microsoft, Cisco, Oracle, Citrix, CIW, Checkpoint, Novell, Sun/Solaris, CWNA, LPI, ISC, etc. All trademarks are properties of their respective owners.

Guarantee

If you use this study guide correctly and still fail the exam, send a scanned copy of your official score notice at:

Sales@Testkingonline.com

We will gladly refund the cost of this study guide or give you an exchange of study guide of your choice of the Free.

This material is protected by copyright law and international treaties. Unauthorized reproduction or distribution of this material, or any portion thereof, may result in severe civil and criminal penalties, and will be prosecuted to the maximum extent possible under law.

Case Study # 1:

BACKGROUND

Fabrikam, Inc. is a manufacturing company that provides components for the engineering industry. Fabrikam has a sister company named Contoso, Ltd.

Work orders are processed by the following departments:

- The Production department handles request for components that are not in stock.

- The Warehouse department handles the storage and stock management of pre-manufactured components.


- The Packaging department handles the packaging of stock in the warehouse or the production department.

- The Distribution department handles the delivery of orders to clients.

The company uses a flexible resourcing structure where individual employees can be reassigned between departments. Productivity is higher when specific groups of staff members are assigned to a particular department.

EXISTING ENVIRONMENT

Fabrikam uses a Microsoft SQL Server 2005 database named FabrikamSQLDB to control and route client requests as work orders to the respective department. New Client and ProductOrder records can be added to the FabrikamSQLDB database throughout the day. These records must be reflected correctly in any reports with zero latency. The FabrikamSQLDB schema is shown in the following exhibit.



The performance of a department is measured against the following factors:

- Time to resolve a work order
- Number of work orders completed in a given period

Existing SSAS Solution

Fabrikam uses a live Microsoft SQL Server 2005 Analysis Services (SSAS) database named FabrikamASDB. At present, the company processes the Fabrikam SSAS solution every evening. This processing takes three hours. The Analysis Server Query Log is enabled.

The SSAS solution contains the following objects:

- A data source view. The data source view within the FabrikamASDB database sources data from the FabrikamSQLDB database.

- A cube named FabrikamCube. FabrikamCube exists within the FabrikamASDB database.

FabrikamCube contains the following measure groups:

- The ProductOrder measure group. This measure group is based on the ProductOrder table.
 - o The ProductOrder measure group has Regular relationships to the Client, Currency, and Product dimensions.
 - o The ProductOrder measure group uses real-time ROLAP that has no aggregations.
 - o Reports against the ProductOrder measure group currently tend to perform poorly when aggregates are produced across a large range of data.
 - o Users have drillthrough capability to individual rows within the ProductOrder measure group.
- The ProductOrderStatus measure group. This measure group is based on the ProductOrderStatus table.
 - o The ProductOrderStatus measure group has approximately 20 million rows.
 - o The ProductOrderStatus measure group is split into the following two partitions:
 - o The CurrentDay partition. This partition is designed to contain all order status data for the current working day and any future-dated orders.
 - o The Historical partition. This partition is designed to contain all historical data prior to the current working day.

- o Data within each partition is filtered by a query. Each query contains dates when the measure group was partitioned.
- o Data older than the current working day is used within highly aggregated reports.
- o Data for the current working day is used for individual order status reviews.

The EmployeeAssignment measure group. This measure group is based on the EmployeeAssignment table.

- o Records are added to the EmployeeAssignment data table continuously throughout the day.

FabrikamCube contains the following dimensions:

- A Client dimension. Clients are divided into three regions: North, South, and West. Each region is stored within the Region attribute of the Client dimension.

Case Study # 1 (Questions):

Question: 1

You need to ensure the Client dimension is able to support new orders and new client records. Which setting should you configure?

- A. ROLAP storage
- B. Lazy aggregations
- C. Write Enabled property
- D. Unknown Member functionality

Answer: A

Question: 2

You need to meet the requirements of the accounts team. What should you do?

- A. Set the Dimension type to Currency.
- B. Set the Currency member name as a unique attribute.
- C. Hide non-leaf data within the Currency dimension key attribute.
- D. Disable the All member of the Currency dimension key attribute.

Answer: D

Question: 3

You need to make structural changes to the cube. You also need to ensure that appropriate aggregations are created after the changes are made. What should you do?

- A. Record a weeks user activity by using SQL Profiler. Apply structural changes. Replay the activity and then run the Usage Based Optimization Wizard.
- B. Apply structural changes. Run the Database Engine Tuning Advisor.
- C. Apply structural changes. Run the Usage Based Optimization Wizard.
- D. Record a weeks user activity by using SQL Profiler. Apply structural changes. Run the Database Engine Tuning Advisor.

Answer: A

Question: 4

You need to select appropriate processing methods for the EmployeeAssignment measure group and the Employee dimension. Which method should you choose?

- A. Run a Process Update against the Employee dimension. Run a Process Full against the EmployeeAssignment measure group.
- B. Run a Process Full against the EmployeeAssignment measure group. Run a Process Update against the Employee dimension.
- C. Run a Process Full against the Employee dimension. Run a Process Incremental against the EmployeeAssignment measure group.

D. Run a Process Incremental against the EmployeeAssignment measure group. Run a Process Full against the Employee dimension.

Answer: A

Question: 5

You need to ensure that the new sales manager is able to view clients within only the North and South regions. What should you do?

- A. Add a Denied member set for the West clients to NorthRole.
- B. Add an Allowed member set for the South clients to NorthRole.
- C. Create a new role for the sales manager and add a Denied member set for the South clients.
- D. Create a new role for the sales manager and add an Allowed member set for the South clients. Remove the sales manager from SouthAndWestRole.

Answer: D

Question: 6

You need to meet the business requirements of the packaging department. What should you do?

- A. Add the ProductSize data table and the ProductSizeDescriptor column to the Product dimension. Set the order of the ProductSKU attribute to the ProductSizeDescriptor column.
- B. Add the ProductSize data table as a dimension and then add this dimension to the cube as a Referenced dimension. Set the order of the key attribute of the ProductSize dimension to the ProductSizeDescriptor column.
- C. Add the ProductSize data table and the ProductSizeDescriptor column to the Product dimension. Create a user-defined hierarchy that has ProductSKU as the top-level attribute and ProductSizeDescriptor as the bottom-level attribute.
- D. Add the ProductSize data table as a dimension and then add this dimension to the cube as a Referenced dimension. Create a user-defined hierarchy within the ProductSize dimension. Set the ProductSizeDescriptor attribute as the top-level attribute.

Answer: A

Question: 7

You need to ensure the Order Control team is able to track the number of orders that reach the production department. What should you do?

- A. Create a new Time dimension based on the Holiday table. Use derived columns for the month and the year.
- B. Create a Server Time dimension. Within this dimension, enable the time periods year, month, and date. Select the Additional Manufacturing calendar option.
- C. Create a new Time dimension based on the date values within the ProductOrderStatus table. Flag those days on which no orders were received as Holiday.
- D. Create a new Time dimension based on a data source view named query. Derive the data for this data source view from unique dates within the ProductOrderStatus table. Use a subquery to identify the holidays and derived columns for the month and the year.

Answer: D

Question: 8

You need to design a storage mode for the ProductOrderStatus measure group. What should you do?

- A. Use ROLAP for the Historical partition and MOLAP for the CurrentDay partition.
- B. Use MOLAP for the Historical partition and ROLAP for the CurrentDay partition.
- C. Use ROLAP for the Historical partition and Automatic MOLAP for the CurrentDay partition.
- D. Use HOLAP for the Historical partition and Low-latency MOLAP for the CurrentDay partition.

Answer: B

You need to maintain the partitions on the ProductOrderStatus measure group. What should you do?

- A. Create a Microsoft SQL Server 2005 Integration Services (SSIS) package that alters the partition queries and processes aggregated partitions every evening.
- B. Enable proactive caching on the partitions. Set the cache to update periodically with a latency of 24 hours.
- C. Set the rebuild interval of the partitions to 36 hours. Create a Microsoft SQL Server 2005 Integration Services (SSIS) package that processes aggregated partitions every evening.
- D. Enable proactive caching on the ProductOrderStatus measure group. Set scheduled polling that re-processes the measure group whenever the system dates change.

Answer: A

Question: 10

You need to ship a copy of the FabrikamASDB metadata to Contoso, Ltd. What should you do?

- A. Use the Synchronize Database Wizard.
- B. Describe the FabrikamASDB database in an XML for Analysis (XMLA) script. Ship the script to Contoso, Ltd.
- C. Copy the folder structure of the live SSAS solution to a compressed file. Ship the compressed file to Contoso, Ltd.
- D. Backup the FabrikamASDB database. Run a Process Full on the restored database. Ship the database to Contoso, Ltd.

Answer: B

Question: 11

You need to help the HR team improve the resource allocation process. What should you do?

- A. Perform data mining regression analysis to identify the correlation between productivity and specific employee groups.
- B. Set the GroupingBehaviour property within the Employee dimension.
- C. Partition the EmployeeAssignment measure group into groups by department.
- D. Use the Fuzzy Grouping task in Microsoft SQL Server 2005 Integration Services (SSIS).

Answer: A Case Study # 2:

BACKGROUND


Woodgrove Bank is a banking institution that offers deposit accounts and mortgage accounts. At present, the bank holds 1,000,000 accounts. On an average 50,000 new accounts are added every year. Ten percent of the accounts are mortgage accounts. The balances of mortgage accounts are recalculated every week. Some of the existing Microsoft SQL Server Reporting Services (SSRS) reports are double-counting transactions where there is more than one account holder for the same account.

EXISTING ENVIRONMENT

Woodgrove Bank has partially implemented a data warehouse named WoodgroveDW.

WoodgroveDW Data Store

The WoodgroveDW database schema is shown in the exhibit.



The WoodgroveDW database tables have the following functions:

Transaction: Fact table

Account: Dimension table

Person: Dimension table

AccountHolder: Resolver table between the Person table and the Account table

Mortgage: Dimension table

- ETLAudit: Log table for Microsoft SQL Server Intelligence Services (SSIS) package outcomes

These tables support the following business rules:

Each account may have multiple transactions.

Each account must have one or more account holders stored in the Person table.

Each person may have one or more accounts,

Transactions may not be associated with a specific person.

- All dimension tables must be Type 2 Slowly Changing Dimensions.

The Account Dimension table characteristics are as follows:

On average, the BandedAverageLast3MonthBalance data changes once per month.

On average, the address lines data changes once every two years.

On average, the AccountType data and the AccountStatus data change every five years.

- The AccountOpenedDate data does not change.

The Transaction Fact table characteristics are as follows:

Holds more than 80 million rows.

All foreign keys do not allow NULL values.

TransactionID is an Identity column. The Identity property is never overridden.

It is possible to raise a back-dated transaction up to two months after the required transaction date.

Ten percent of accounts are attached to mortgages.

Extract, Transform, and Load (ETL) Processes

A number of SSIS packages feed data from source systems to the WoodgroveDW database. An existing SSIS package populates the MortgageType table in the WoodgroveDW database. This SSIS package checks and populates the MortgageTypeCode column. This package uses a generic placeholder value for the MortgageTypeDescription field.

Woodgrove SQL Server Analysis Services (SSAS) Database

The SSAS solution has been partially implemented, and contains a database named WoodgroveSSAS. This database contains a data source view named WGDSV. WGDSV contains a named query that joins data from three tables in the bank's customer relationship management (CRM) database.

A measure group named Transaction has been created on the Transaction table. The following dimensions have been created:

Account

Person

Marketing Department

The marketing department is currently promoting three new products: WoodgroveLifeInsurance, WoodgroveHomeInsurance, and WoodgroveFraudProtection.

Case Study # 2 (Questions):

Question: 1

You need to reduce the duration of the daily full database backup process. What should you do?

- A. Perform transaction log backups every 30 minutes throughout the day.
- B. Partition the Transaction table by a range based on the TransactionID. Daily, back up only the file groups that have changed.
- C. Place the Transaction table in its own file group. Back up the Transaction table and the remaining file groups each evening.
- D. Partition the Transaction table by a range based on the MortgageID. Daily, back up only the file groups that have changed.

Answer: B

Question: 2

You need to identify the data mining algorithm that meets the requirements of the credit department. Which algorithm should you use?

- A. Microsoft Clustering
- B. Microsoft Decision Tree
- C. Microsoft Naive Bayes
- D. Microsoft Association

Answer: B

Question: 3

You need to deliver the three reports that detail the previous days transaction activity. What should you do?

- A. Create a new report that has three embedded subreports. Create one standard subscription.
- B. Create three new reports based on the same dataset. Create one data-driven subscription.
- C. Create three data-driven subscriptions on a single shared schedule. Ensure that the e-mail recipient list is identical across all three subscriptions.
- D. Create three standard subscriptions on a single shared schedule. Ensure that the mail recipient list and subject are identical across all three subscriptions.

Answer: A

Question: 4

You need to meet the fraud detection requirements. What should you do?

- A. Use Percentage Sampling in SSIS.
- B. Use a Neural Network Data mining algorithm.
- C. Create an SSAS Key Performance Indicator for each account.
- D. Monitor transactions by using SQL Server Profiler.

Answer: B

Question: 5

You need to record package and task outcome in the ETLAudit table. Which components of SSIS should you use?

- A. Custom event handlers
- B. Data flow error output
- C. SSIS log provider for SQL Server
- D. Execute SQL tasks with failure precedence constraints

Answer: A

Question: 6

You need to resolve the blocking issue that is affecting the performance of the CRM system. What should you do?

- A. Replace the named query with an SQL view.
- B. Set the transaction isolation level to serializable in the named query.
- C. Set the source database to use snapshot isolation.
- D. Set the transaction isolation level to repeatable read in the named query.

Answer: C

Question: 7

You need to resolve the double counting of transactions issue in the existing reports. What should you do?

- A. Apply a many-to-many relationship between the Transaction measure group and the Person dimension in the WoodgroveSSAS database.
- B. Apply a Conditional Split transformation in the Transaction data in the SSIS package.
- C. Delete the AccountHolder table in the WoodgroveDW database.
- D. Apply a foreign key relationship between the Person table and the Transaction table in the WoodgroveDW database.

Answer: A

Question: 8

You need to select the appropriate edition of SQL Server 2005 for the portable computers that are used in the Mortgage Data Cleansing project. Which edition should you use?

- A. Microsoft SQL Server 2005 Mobile Edition
- B. Microsoft SQL Server 2005 Express Edition
- C. Microsoft SQL Server 2005 Standard Edition
- D. Microsoft SQL Server 2005 Workgroup Edition


Answer: B Case Study # 3:

BACKGROUND

Fourth Coffee is a chain of coffee shops that operates in North America. Service personnel in the outlets receive commission for each item they sell. The outlets are grouped into geographical region.

EXISTING ENVIRONMENT**Fourth Coffee Database Schema**


The Fourth Coffee database schema is shown in the following exhibit.

**Sales Cube**

The SSAS solution database contains a sales cube with a single measure group based upon the Sale data table and the SaleItem data table. The Sales Cube measure group contains a quantity measure. The Region dimension is not included within the sales cube. The daily sales data for each store is loaded into the existing SSAS solution database everynight.

Product Dimension

The Product dimension supports the marketing department's Microsoft PivotTable. The existing Product dimension structure is shown in the following exhibit.



The marketing department uses the PivotTable that is connected directly to the existing SSAS solution to identify which products are selling more. This PivotTable combines the CategorySubCategory hierarchy from the Product dimension with Quantity measure. A blank row appears where there is no applicable product subcategory in the hierarchy.

Case Study # 3 (Questions):

Question: 1

You need to ensure that the TownStore hierarchical structure is displayed accurately in reports. How should you modify the Store dimension?

- A. Add Town as an attribute relationship to StoreID in the SSAS solution.
- B. Modify the StoreName attribute and set the MemberNameUnique property to true.
- C. Add the StoreID attribute to the bottom of the TownStore hierarchy in the SSAS solution.
- D. Use column binding to add Town to the Key column of the StoreName attribute in the SSAS solution.

Answer: D

Question: 2

You need to ensure the marketing department is able to implement cross-selling requirements. Which data mining model should you choose?

- A. Microsoft Time Series
- B. Microsoft Nave Bayes
- C. Microsoft Neural Networks
- D. Microsoft Association

Answer: D

Question: 3

You need to ensure the Point-Of-Sales system supports cross-selling functionality. What should you do?

- A. Use ADO.NET to execute a Transact-SQL statement to retrieve the data.
- B. Use ADOMD.NET to execute a Data Mining Extensions (DMX) query to retrieve the data.
- C. Use ADOMD.NET to execute a Multidimensional Expressions (MDX) query to retrieve the data.

D. Use Microsoft SQL Server 2005 Integration Services (SSIS) package that runs overnight to execute a data mining query task.

Answer: B

Question: 4

You need to calculate the average quantity of products sold by each store. What should you do?

- A. Create a calculated member within the Measures hierarchy. Divide the Quantity measure by a distinct count of store members.
- B. Create a calculated member within the StoreID hierarchy of the Store dimension. Divide the Quantity measure by a distinct count of store members.
- C. Create a calculated field within the SaleQuery table inside the data source view. Divide the QuantitySold column by a distinct count of the number of stores.
- D. Create a calculated field within the Store data table inside the data source view. Sum the QuantitySold column by that store and divide by the number of sales made.

Answer: A

Question: 5

You need to include the attributes of the Region dimension within the SSRS report. What should you do?

- A. Join the Region dimension to the Sales measure group by using a Fact relationship type.
- B. Join the Region dimension to the Sales measure group by using a Regular relationship type.
- C. Join the Region dimension to the Sales measure group by using a Referenced relationship type.
- D. Join the Region dimension to the Sales measure group by using a Many-to-Many relationship type.

Answer: C

Question: 6

You need to implement the security requirements of the SupplierData report. What should you do?

- A. Use the Internet Information Services (IIS) security to prohibit access to the new report.
- B. Within Report Manager, grant Browser rights and Report Builder rights on the new sales report to the users of the management group.
- C. Create a new user group and add all management users to this group. Within Report Manager, grant Browser rights on the report to the management group.
- D. Create a new user group and add all management users to this group. Use the file system security to grant access to the report for the management group.

Answer: C

Question: 7

You need to alter the StaffMember dimension and measure group relationship to resolve the processing errors. You also need to meet the requirements for handling NULL values in the StaffMember dimension. What should you do?

- A. Alter the relationship to map NULL to zero or blank. Add a default member to the StaffMember dimension.
- B. Alter the relationship to map NULL to Unknown. Design the Unknown Member functionality within the dimension.
- C. Alter the relationship to error upon NULL key values. Design the custom error configuration in the StaffMember dimension.
- D. Alter the relationship to preserve NULL values. Design the Missing Member Mode functionality of Multidimensional Expressions (MDX) in the dimension.

Answer: B

Question: 8

You need to design the Monthly Sales Items report for the regional managers. What should you do?

- A. Create a single report that uses the caching functionality of SSRS.
- B. Create a single report in SSRS by using a Matrix format that has the lower levels of grouping collapsed by default.
- C. Create a single parent report that displays the highest level of aggregate data. Embed subreports for each grouping of child data.
- D. Create a summary report. Add drillthrough functionality to the summary report to link to separate detailed reports for each grouping of data.

Answer: D

Question: 9

You need to ensure that all blank rows in the PivotTable reports are hidden. Which property of the Product dimension should you set?

- A. Set the AttributeHierarchyEnabled property to false.
- B. Set the HideMemberIf property to NoName.
- C. Set the AttributeHierarchyVisible property to true.
- D. Set the AttributeHierarchyVisible property to false.

Answer: B

Question: 10

You need to create a data mining model to support the SalesProjections report. How should you modify the SSAS solution?

- A. Create a mining model on the existing Time Of Day dimension by using the Microsoft Time Series algorithm.
- B. Create a Date dimension within the SSAS solution. Use the new Date dimension within a Microsoft Time Series data mining model.
- C. Create a Key Performance Indicator (KPI) that displays sales trends.
- D. Create a data mining model on the existing Store dimension by using the Microsoft Time Series algorithm.

Answer: B Case Study # 4:


BACKGROUND

Contoso, Ltd. is a vineyard and winery company. The company advertises its products through promotion campaigns. Customers might respond to more than one promotion. The company maintains a log of the catalogues that are sent in response to each promotion.

EXISTING ENVIRONMENT

ContosoDW

The company is designing a Microsoft SQL Server data warehouse named ContosoDW. The ContosoDW database is shown in the following exhibit.



LoadResellerSales. This package loads the ResellerSales file into the ContosoDW database.


ContosoETL. This Extract, Transform, and Load (ETL) package extracts data from the Contoso Sales database and loads it into the ContosoDW database.

ReloadWine. The Wine table in the ContosoDW database is updated periodically with wine descriptions. This updated data should be pushed into the ContosoSSAS solution immediately. This package is created by using default settings to address these update requirements. The package is shown in the following exhibit.

SalesExtract. This package extracts all sales data from the ContosoDW database. It is executed every night during a fixed time slot that cannot be altered. The package is used for projections and large aggregations. Performance of the data extract is more important than data accuracy.

ImportWineOrder. This package has three connection managers that point to three data sources: an OLE DB source, a flat file, and a Microsoft Office Excel source. Each data source contains the same columns, data types, and names.

SalesPersonBonus. This package handles the bonus approvals of sales personnel.



Case Study # 4 (Questions):

Question: 1

You need to modify the SalesExtract SSIS package to meet all business and technical requirements. Which transaction isolation level should you use?

- A. Serializable
- B. Read Committed
- C. Repeatable Read
- D. Read Uncommitted

Answer: D

Question: 2

You need to redesign the data flow for the ImportWineOrder SSIS package to meet the ETL requirements. Which transformation task should you use?

- A. Merge
- B. Multicast
- C. Union All
- D. Merge Join

Answer: C

Question: 3

You need to configure event handling within the TransferResellerSales SSIS package. What should you do?

- A. Add an OnError event handler at the package level to execute the CleanUp package. Add an OnError event handler to the FTPGetFile task to execute the ResetTransferFlags package.
- B. Add an OnError event handler for each task in the package to execute the CleanUp package. Add an OnError event handler to the FTPGetFile task to execute the ResetTransferFlags package.
- C. Add an OnError event handler at the package level to execute the CleanUp package. Add an OnError event handler to the FTPGetFile task to execute the ResetTransferFlags and CleanUp packages.
- D. Add an OnError event handler for each task in the package to execute the CleanUp package. Add an OnError event handler to the FTPGetFile task to execute the ResetTransferFlags and CleanUp packages.

Answer: A

Question: 4

You need to ensure that the ContosoSSAS measure groups remain available after the ReloadWine SSIS package is executed. What should you do?

- A. Set the Process Wine Dimension task to Process Full.
- B. Set the Process Wine Dimension task to Process Update.
- C. Add a Control Flow task to process the cube after the Process Wine Dimension task.
- D. Add a Control Flow task to process the cube before the Process Wine Dimension task.

Answer: B

Question: 5

You need to modify the CustomerBalance measure group to meet the reporting requirements. What should you do?

- A. Set the Date dimension to type Time. Set the AggregateFunction property of the CustomerBalance measure to LastChild.
- B. Set the Date dimension to type Time. Set the AggregateFunction property of the CustomerBalance measure to LastNonEmpty.
- C. Set the Customer dimension to type Customers. Set the AggregateFunction property of the CustomerBalance measure to LastChild.
- D. Set the MDXMissingMember property of the Date dimension to Ignore. Set the AggregationPrefix property of the CustomerBalance measure group to LatestBalance.

Answer: B

Question: 6

You need to redesign the flow sequence of the SalesPersonBonus SSIS package to ensure sales bonuses are automatically

approved if all requirements are met. Which two tasks should you perform? (Each correct answer presents part of the solution. Choose two.)

- A. Change the precedence constraint between the Days Absent task and the Bonus Approved task to OR.
- B. Change the precedence constraint between the Load Employee Sales Data task and the Process Sales Person Type A task to OR.
- C. Change the precedence constraint between the Process Sales Person Type A task and the Process Sales Employee task to OR.
- D. Change the precedence constraint between the Process Sales Person Type B task and the Process Sales Employee task to OR.
- E. Change the precedence constraint between the Days Absent task and the No Bonus task to OR.

Answer: C, D

Question: 7

You need to redesign the LoadResellerSales package to meet the ETL requirements. What should you do?

- A. Load the customer data and the wine order data simultaneously. Use a Lookup transformation to obtain the correct customer number.
- B. Disable the foreign key between the Customer table and the Wineorder table. Load data into both tables simultaneously. Restore the foreign key after the operation is complete.
- C. Load the customer data first and then write the wine order data to a Raw file. Load the wine order data from the Raw file and use a Lookup transformation to obtain the correct customer number.
- D. Load the wine order data first and then write the customer data to a Raw file. Load the customer data from the Raw file and use a Lookup transformation to obtain the correct customer primary key.

Answer: C

Question: 8

You need to create the PromotionResponse mining model for the promotions department. What should you do first?

- A. Add a foreign key to the Customer dimension and link it to the Promotion dimension.
- B. Create a CatalogueRequest fact table that links to the Promotion dimension and the Customer dimension.
- C. Add a column to the Customer table that indicates whether a customer has requested a catalogue before making a purchase.
- D. Add a column to the WineOrder table that indicates whether a customer has requested a catalogue before making a purchase.

Answer: B

Question: 9

You need to meet the processing requirements of the ContosoSSAS solution. What should you do?

- A. Use HOLAP storage.
- B. Enable lazy aggregations.
- C. Create usage-based aggregations.
- D. Use MOLAP with a 12-hour latency.

Answer: B Case Study # 5:

BACKGROUND

A. Datum Corporation is a software company that has manufactured a software application named ADatum Datacapture. This application is designed to capture real estate data.


EXISTING ENVIRONMENT

A. Datum Corporation has a live data mining structure named Bugs_DMStruct. This structure has several models attached to it.

CollectRemoteApplicationLogs SSIS Package

The company also has a Microsoft SQL Server 2005 data warehouse named ADatumSQLDB and a Microsoft SQL Server 2005 Integration Services (SSIS) package named CollectRemoteApplicationLogs.

The CollectRemoteApplicationLogs SSIS package is shown in the following exhibit.



The Select details for remote sites task derives a list of remote computers from the RemoteSitetable. This table has four columns: RemoteSiteID, RemoteSiteFTPAddress, RemoteSiteUsername, and RemoteSitePassword. The SSIS package uses this list to load application logs from remote computers into the RemoteApplicationLog table. If the CollectRemoteApplicationLogs SSIS package fails at any point, it must stop immediately so that an operator can check the package and restart it. The CollectRemoteApplicationLogs SSIS package must be able to restart from the remote target it failed against without re-collecting from the already successful remote targets.

SoftwareControl database

Versions of the company software are stored in a database named SoftwareControl. The SoftwareControl database contains the following tables:

The Software table. This table contains one row for each software product that the company manufactures. Rows in this table can be updated or appended.

The Version table. This table stores the latest version details of each software product. When a software product is upgraded, a row is removed from the Version table and a new row is added. Some software products have no associated version record.

SoftwareVersion table

The SoftwareVersion table is populated by the Software table and Version table. The SoftwareVersion table in the ADatumSQLDB data warehouse is shown in the following exhibit.

SoftwareVersion	
SoftwareVersionID	
SoftwareID	
SoftwareName	
VersionName	
VersionID	
NewestRecord	

Application Log

When the ADatum Datacapture application is installed at the client site, it produces a single logfile named Application.log. This log file contains user actions and errors from all client machines at the site. The Application.log file is shown in the following exhibit.

```

"SiteID","MachineName","DateTime","output"
"1","LP5100","2006-03-01 11:22:11","Record open"
"1","LP5100","2006-03-01 11:56:00","File Save"
"1","LP5100","2006-03-01 11:56:12","Update Field #DATEOFBIRTH"
"1","LP5100","2006-03-01 11:56:15","Error 100010 occurred!"
"1","PC1553","2006-03-01 10:55:03","File open"
"1","PC1553","2006-03-01 10:55:12","Update Field #GIVENNAME"
"1","PC1553","2006-03-01 10:55:14","Update Field #DOB"
"1","PC1553","2006-03-01 10:55:16","Update Field #DATEOFBIRTH"
"1","PC1553","2006-03-01 10:55:20","File Save"
"1","PC1553","2006-03-01 10:55:20","Close Application"
"1","XPX100","2006-03-01 12:05:15","File open"
"1","XPX100","2006-03-01 12:08:15","Browse"
"1","XPX100","2006-03-01 12:10:15","File Save"
"1","XPX100","2006-03-01 12:11:15","File Save"
"1","XPX100","2006-03-01 12:14:00","Backup"
"1","XPX100","2006-03-01 12:16:12","Browse"
"1","XPX100","2006-03-01 12:17:15","Error 100020 occurred!"
"1","YPZ423","2006-03-01 15:11:15","File open"
"1","YPZ423","2006-03-01 15:14:00","Backup"
"1","YPZ423","2006-03-01 15:16:12","Browse"
"1","YPZ423","2006-03-01 15:17:15","Error 100020 occurred!"

```

The CollectRemoteApplicationLogs SSIS package is used to transfer the Application.log files from all remote computers to the ADatumSQLDB data warehouse.

Case Study # 5 (Questions):

Question: 1

You need to control report rendering per the technical requirements. What should you do?



- A. Modify the RSReportServer.config file.
- B. Modify report settings in the Report Manager.
- C. Modify the Web.config file.
- D. Modify the report design in Business Intelligence Development Studio (BIDS).



E. Remove the PDF and TIF rendering extensions from the RSReportServer.config file.

Answer: A

Question: 2

You need to design the ETL process for the SoftwareVersion table. Which data flow diagram should you use?





Answer: A

Question: 3

You need to train the marketing teams clustering model. What should you do?

- A. Train the model on Set 1 data. Test on Set 3 data. Identify anomalous data by the cluster to which it belongs.
- B. Train the model on Set 2 data. Test on Set 1 data. Identify anomalous data by the cluster to which it belongs.
- C. Train the model on Set 3 data. Test on Set 2 data. Identify anomalous data by its low probability to fit into a cluster.
- D. Train the model on Set 2 data. Test on Set 1 data. Identify anomalous data by its low probability to fit into a cluster.

Answer: D

Question: 4

You are configuring the UploadSupportScreenShot package. You need to load the bitmap (BMP) files into a data flow. Which transformation task should you use?

- A. Lookup
- B. Copy column
- C. Import column
- D. Derived column

Answer: C

Question: 5

You need to recommend a strategy for maintaining the DatumReportSuite. What should you do?

- A. Use the Report Server Web service to develop the appropriate application.
- B. Use the Reporting Services Configuration tool to manage report deployment.
- C. Use the Report Definition Language (RDL) script to define the reports.
- D. Use Configuration Manager in Business Intelligence Development Studio (BIDS) to manage report deployment.

Answer: A

Question: 6

You need to train the Bugs_MSeq mining model that the application team uses to diagnose application errors. Which model should you choose?

A.

CASE TABLE	
CaseID	
1	
2	
3	

NESTED TABLE					
Siteid	MachineName	Datetime	Action	CaseID	NestedID
1	LP5100	01-Mar-06 11:22:11 AM	Record Open	1	1
1	LP5100	01-Mar-06 11:56:00 AM	File Save	1	2
1	LP5100	01-Mar-06 11:56:12 AM	Update Field #DATEOFBIRTH	1	3
1	LP5100	01-Mar-06 11:56:15 AM	Error 100010 Occured!	1	4
1	XPX100	01-Mar-06 12:11:15 PM	File Save	2	1
1	XPX100	01-Mar-06 12:14:00 PM	Backup	2	2
1	XPX100	01-Mar-06 12:16:12 PM	Browse	2	3
1	XPX100	01-Mar-06 12:17:15 PM	Error 100020 Occured!	2	4
1	YPZ423	01-Mar-06 03:11:15 PM	File Open	3	1
1	YPZ423	01-Mar-06 03:14:00 PM	Backup	3	2
1	YPZ423	01-Mar-06 03:16:12 PM	Browse	3	3
1	YPZ423	01-Mar-06 03:17:15 PM	Error 100020 Occured!	3	4

Notes:

Action column is set to predict

NestedID column is set to key sequence

CaseID column in the case table is set to key

B.

CASE TABLE				
Siteid	MachineName	Datetime	Action	SequenceID
1	LP5100	01-Mar-06 11:22:11 AM	Record Open	1
1	LP5100	01-Mar-06 11:56:00 AM	File Save	2
1	LP5100	01-Mar-06 11:56:12 AM	Update Field #DATEOFBIRTH	3
1	LP5100	01-Mar-06 11:56:15 AM	Error 100010 Occured!	4
1	XPX100	01-Mar-06 12:11:15 PM	File Save	5
1	XPX100	01-Mar-06 12:14:00 PM	Backup	6
1	XPX100	01-Mar-06 12:16:12 PM	Browse	7
1	XPX100	01-Mar-06 12:17:15 PM	Error 100020 Occured!	8
1	YPZ423	01-Mar-06 03:11:15 PM	File Open	9
1	YPZ423	01-Mar-06 03:14:00 PM	Backup	10
1	YPZ423	01-Mar-06 03:16:12 PM	Browse	11
1	YPZ423	01-Mar-06 03:17:15 PM	Error 100020 Occured!	12

Notes:

SequenceID column is set to key sequence

Action column is set to predict

C.

CASE TABLE			
Siteid	MachineName	Datetime	Action
1	LP5100	01-Mar-06 11:22:11 AM	Record Open
1	LP5100	01-Mar-06 11:56:00 AM	File Save
1	LP5100	01-Mar-06 11:56:12 AM	Update Field #DATEOFBIRTH
1	LP5100	01-Mar-06 11:56:15 AM	Error 100010 Occured!
1	XPX100	01-Mar-06 12:11:15 PM	File Save
1	XPX100	01-Mar-06 12:14:00 PM	Backup
1	XPX100	01-Mar-06 12:16:12 PM	Browse
1	XPX100	01-Mar-06 12:17:15 PM	Error 100020 Occured!
1	YPZ423	01-Mar-06 03:11:15 PM	File Open
1	YPZ423	01-Mar-06 03:14:00 PM	Backup
1	YPZ423	01-Mar-06 03:16:12 PM	Browse
1	YPZ423	01-Mar-06 03:17:15 PM	Error 100020 Occured!

Notes:

MachineName column is set to key

Action column is set to predict

Datetime column is set to time key

Answer: A

Question: 7

You need to create the TransferPurchaseOrders SSIS package per the ETL requirements. Which data flow destination configuration should you use?

A. DataReader destination

- B. SQL Server destination
- C. Flat File destination
- D. OLE DB destination with fast load

Answer: D

Question: 8

You need to create the TransferReceiptRecords package to meet the ETL requirements. What should you do?

- A. Use a single transaction.
- B. Use a single sequence container.
- C. Use the Repeatable Read setting.
- D. Use the RetainSameConnection property.

Answer: D

Question: 9

You need to ensure that the CollectRemoteApplicationLogs SSIS package does not reprocess remote sites. What should you do?

- A. Enable package Checkpointing. Ensure that no task within the package has dependencies upon OnPostExecute or OnError event handlers.
- B. Set the TransactionOption property of the ProcessRemoteSites container to Required.
- C. Add a column to the RemoteSite table to track the Processed/Not Processed status of each remote site. Set the Select Details From Remote Sites task to return only Not Processed sites. Add a task to the ProcessRemoteSites container to set status to Processed as each site is successfully processed. Add a task to execute after the ProcessRemoteSites container resets all status values to Not Processed.
- D. Add an OnPreExecute event handler to the FTP Get File From Remote Site task that skips the file when an FTP connection cannot be established. Add an OnError event handler to the ProcessRemoteSites container to restart the package when errors occur.

Answer: C

Question: 10

You need to apply business rules to the new valuation reports. Which method should you use?

- A. Create a custom assembly and use it in each new valuation report.
- B. Embed custom code into each new valuation report.
- C. Create and deploy a reporting model project.
- D. Modify the Report Definition Language (RDL) in each new valuation report.

Answer: A

Question: 11

The application team wants to test the mining model named Bugs_Timeseries. You need to recommend a processing method. What should you recommend?

- A. Run a Process Structure, and then run a Process Full on the new mining model.
- B. Run an Unprocess operation on the structure, and then run a Process Structure.
- C. Run a Process Full on the structure, and then run a Process Full on the new mining model.
- D. Run an Unprocess operation on the structure, and then run a Process Full on the structure.

Answer: A Case Study # 6:

BACKGROUND

City Power & Light is a company that supplies electricity to large corporate customers and individual households. Prices fluctuate and are set one day in advance for each hour in a day.

EXISTING ENVIRONMENT

Microsoft SQL Server Platform

A SQL Server 2005 data warehouse named CPLSQL contains the following fact tables:

ServiceCall. This holds around one million rows that track the status of service calls on equipment and infrastructure. Each record in the fact table relates to one service call.

- HouseholdBilling. This holds the billing information of each household, The fact table contains the following columns:
 - o HouseholdBillingRecordID (primary key)
 - o HouseholdID
 - o RegionID
 - o PeriodID
 - o BillingAmount

Microsoft SQL Sever Analysis Services (SSAS) Platform

An SSAS database named CPLSSAS is hosted on a server named CPLServer12.

The CPLSSAS database has the following dimensions:

Date. This holds a record for each day.

Time. This holds 24 members for each hour of the day.

Period. This holds the period to which a household bill relates.

PriceBand. This holds the corporate price of electricity for a given hour on a given date.

• CorporateCustomer. This holds information on the corporate customer. The CorporateCustomer dimension includes the following attributes:

- o CorpRefNumber
- o CompanyName
- o BusinessType
- o IndustryType
- o ConsumptionClassification
- o IsCurrentVersionFlag

The company uses CorpRefNumber to identify each corporate customer uniquely. The company tracks the electricity consumed by each corporate customer in kW per minute.

Each service call raised by a corporate customer is assigned a ServiceCallReference number. At present, the ServiceCallReference number is not stored in the CPLSSAS database.

- Household. This holds information about individual households and contains the following attributes:
 - o State
 - o County
 - o Town

County is unique to State. Town is unique to County and State. The planned reports sum data at various levels within this set of attributes.

Logical design has identified all dimensions as Type 2.


Microsoft SQL Sever Reporting Services (SSRS) Platform

A report named HouseholdRegion extracts data from the HouseholdBilling fact table. This report shows the total billing amount of a region for a specific period. The company has developed a Customer Account Reports suite. The reports in this suite are available on the Internet so that corporate clients can view the account details. The reports are sourced on the CPLSSAS database. Company

employees use Report Builder for ad hoc reporting. A report model named RMCPL is created. The report model provides data for employees of all departments. However, users report that this model contains too many unnecessary entities and attributes.

Microsoft SQL Server Integration Services (SSIS) Platform

An SSIS solution named CPLETL transforms data from the company's source systems to the CPLSQL database. The CPLETL solution extracts data based upon the business processing day, which is usually a day behind the current calendar day. This data value is stored as a variable in all packages of the CPLETL solution. A data flow named TransformPriceBand exists in the CPLETL SSIS solution. The TransformPriceBand data flow is shown in the following exhibit.



A perimeter server named CustomerServer was built and deployed outside a firewall. At present, the server has Microsoft Windows Server 2003 installed. The firewall permits only SSAS traffic between the CustomerServer server and the CPLServer12 server.

Case Study # 6 (Questions):

Question: 1

You need to source data for the Household Billing report. What should you do?

- A. Use a SQL query.
- B. Use a SQL Server view.
- C. Use a SQL Server trigger.
- D. Use a SQL Server stored procedure.

Answer: D

Question: 2

You need to redesign the RMCPL report model to meet the security requirements. What should you do?

- A. Configure two perspectives.
- B. Configure model item security.
- C. Configure two data source views.
- D. Configure entities and attributes in separate folders in the report model.

Answer: B

Question: 3

You need to modify the CorporateCustomer dimension table to support the logical design. What should you do?

- A. Add a Timestamp column.

- B. Add a Surrogate Key column.
- C. Add Version Start and End Date columns.
- D. Add a primary key constraint to the CorpRefNumber column.

Answer: B

Question: 4

You need to meet the engineering department requirements. What should you do?

- A. Add the ServiceCallReference number to the CorporateCustomer dimension.
- B. Add the ServiceCallReference number to the ServiceCall fact table. Use this fact table as a degenerate dimension.
- C. Create a new fact table to store the ServiceCallReference number. Reference this new fact table from the ServiceCall fact table.
- D. Create a new dimension table to store the ServiceCallReference number. Reference this new dimension table from the ServiceCall fact table.

Answer: B

Question: 5

You need to improve the performance of the HouseholdRegion report. Which method should you use?

- A. Create a view on the HouseholdBilling table that includes the PeriodID columns and the RegionID column.
- B. Create a view on the HouseholdBilling table that sums BillingAmount by PeriodID and RegionID.
- C. Partition the HouseholdBilling table into ranges based upon BillingAmount.
- D. Create a non-clustered index on the PeriodID column and the RegionID column of the HouseholdBilling table. Include the BillingAmount column at the leaf level of the index.

Answer: D

Question: 6

You need to redesign the TransformPriceBand data flow to record data lineage. What should you do?

- A. Add a Copy Column transformation.
- B. Add an Import Column transformation.
- C. Add an Audit transformation and a Derived Column transformation.
- D. Add an Audit transformation and an Import Column transformation.

Answer: C

Question: 7

You install the Reporting Services component on the CustomerServer server. You need to identify the additional SQL Server component that must be installed to ensure SSRS functions correctly. Which SQL Server component should you install?

- A. Analysis Services
- B. Integration Services
- C. SQL Server Database Services
- D. Notification Services

Answer: C

Question: 8

You need to design a fact table that records the amount of electricity consumed by corporatecustomers.Which fields should you include?

- A. Minute, Date, CorporateCustomer, and kWUsed

- B. Minute, Date, kWUsed, PriceBand, and TotalPrice
- C. Hour, Date, CorporateCustomer, kWUsed, and PriceBand
- D. Hour, Date, CorporateCustomer, kWUsed, and TotalPrice

Answer: C

Question: 9

You need to design the overnight processing of the CLPSSAS solution. What should you do?

- A. Use separate SSAS processing tasks within a single sequence container in SSIS.
- B. Use separate Analysis Services processing tasks within a single SSIS transaction.
- C. Use an XML for Analysis (XMLA) batch within an Analysis Services Execute DDL task in SSIS.
- D. Use separate SSAS processing tasks with corresponding unprocess tasks in OnError event handlers.

Answer: C

Question: 10

You need to design the Reporting Services infrastructure to support the security requirements of the CustomerAccount Reports suite. Which technology should you use?

- A. Reporting Services folder security
- B. Reporting Services role assignments
- C. Reporting Services security extensions
- D. Reporting Services security policy files


Answer: C

Case Study # 7: BACKGROUND Corporate Information

Wingtip Toys manufactures and supplies model aircraft to toy stores around the world.

EXISTING ENVIRONMENT Existing Data Warehouse Environment

The company uses a Microsoft SQL Server 2005 data warehouse that implements the following data model.



The ETL process updates the Calendar. IsInCurrentSalesperiod field during each data refresh. Each toy plane model is built from various components. The data that defines the construction of toy plane models is stored in XML documents as shown in the following sample.

```

<ToyPlane> <Name>Wi gtip2000</Name> <Component>
  <Name>3 volt lightweight single gear</Name> <Quantity> 1</Quantity>
</Component>

  <Component> <Name>Small Acrobatic Bodyframe</Name> <Quantity> 1 <Quantity>
    <Component>
      <Name>Small Standard Tail</Name> <Quantity> 1</Quantity> <Component>
        <Name>30" pushrod set</Name>
        <Quantity> 2</Quantity> </Component> <Component>
  
```

ToyPlaneModelID	UnitsSold 1WeekAgo	UnitsSold 2WeeksAgo	UnitsSold 3WeeksAgo	UnitsSold 4WeeksAgo	UnitsSold Last3Months
1	20	23	13	32	306
2	56	48	42	38	523
3	11	19	8	15	164
4	21	15	17	28	218

The data within SalesCube is refreshed overnight. This cube was initially deployed in a test environment where 26 queries were tested and logged against sample data. The cube was deployed to a production environment one day ago. Query logging is not currently enabled in the production environment.

Case Study # 7 (Questions):

Question: 1

You need to ensure that sales employees can drill through an Excel report to the CustomerDetail report. What should you do?

- A. Create a data mining structure based on the Customer dimension of the cube.
- B. Within the SalesCube cube, create a drillthrough action on all measure groups.
- C. Within the SalesCube cube, create a report action that has a target type of dimension members.
- D. Use SSRS navigation from within reports that include the Customer dimension to link to the Customer Detail report.

Answer: C

Question: 2

You need to source data for the Engineering departments forecast report. What should you do?

- A. Create a Data Mining Extensions (DMX) query by deriving source data from a SQL statement.
- B. Create a Multidimensional Expressions (MDX) query by deriving source data from a SQL statement.
- C. Create a Data Mining Extensions (DMX) query by deriving source data from a Multidimensional Expressions (MDX) statement.
- D. Create a Multidimensional Expressions (MDX) query by deriving source data from a Data Mining Extensions (DMX) statement.

Answer: A

Question: 3

You need to recommend an appropriate aggregation strategy for the production SSAS solution. What should you recommend?

- A. Use the Usage Based Optimization Wizard on the SalesCube cube.
- B. Use estimated counts within the dimensions and the Aggregation Design Wizard.
- C. Set each Dimension to use the LazyAggregations Processing mode. Reprocess the SalesCube cube.
- D. Set the AllMemberAggregationUsage property to Full for each dimension. Reprocess the SalesCube cube.

Answer: B

Question: 4

You need to ensure that the PhoneNumber attribute is available for the Customer Detail report. You also need to ensure that the PhoneNumber attribute is not available for filtering or aggregation. What should you do?

- A. Add the PhoneNumber attribute to the Customer dimension and set the AttributeHierarchyEnabled property to false.
- B. Add the PhoneNumber attribute to the Customer dimension and set the data type to PersonContact.
- C. Add the PhoneNumber attribute to the Customer dimension and set the MembersWithData property to NonLeafDataHidden.
- D. Add the PhoneNumber attribute to the Customer dimension and set the AttributeHierarchyDisplayFolder property to Member Property Only.

Answer: A

Question: 5

You need to generate a report that contains all the components that are required for a specific toy plane model. You design a Component dimension. What should you do next?

- A. Implement a user-defined hierarchy.
- B. Implement column binding.
- C. Implement a parent-child relationship.
- D. Implement a time dimension.

Answer: C

Question: 6

You need to generate the Sales versus Forecast report. What should you do?

- A. Create a new combined SalesAndSalesForecast measure group that joins Sales and SaleForecast at the ToyPlaneModel level.
- B. Create a new dimension named YearAndQuarter. Relate the SalesForecast measure group to this new dimension at the dimension key level.
- C. Redesign the data source view and the Calendar dimension to include a YearAndQuarter calculated field. Relate the SalesForecast measure group to the Calendar dimension at YearAndQuarter granularity.
- D. Create a new dimension on the SalesForecast data table exposing the YearAndQuarter attribute. Relate the SalesForecast dimension to the SalesForecast measure group by using a Fact relationship type.

Answer: C

Question: 7

You need to create a single measure group that generates aggregated sales volume for all reports that are required by senior management. What should you do?

- A. Create a data source view that contains a named query that joins the SaleLine table and the ToyPlaneModel table.
- B. Create a data source view that contains a named query that joins the SaleHeader table and the Customer table.
- C. Create a data source view that contains a named query that joins the SaleHeader table and the SaleLine table.
- D. Create a data source view that contains two named queries that expose the SaleLine table and the SaleHeader table separately.

Answer: C Case Study # 8:

BACKGROUND

Humongous Insurance is an insurance underwriter. Humongous Insurance provides motor vehicle insurance policies for individuals.

Humongous Insurance uses two sales channels:

- Direct sales: Employs sales personnel.
- Broker sales: Relies on third-party resellers.

EXISTING ENVIRONMENT

Existing Source Systems

Humongous Insurance has two online transaction processing (OLTP) source systems named Underwriting and Claims.

Underwriting system

Data for the Underwriting system is stored in a database named HIDB. The HIDB database contains a tblTXN table that holds records of policy transaction details.

Existing Reporting Platform

A Microsoft SQL Server 2005 Reporting Services (SSRS) platform has been deployed.

Policy Transaction Report

An existing Policy Transaction report sources data from the HIDW database. The Policy Transaction report contains the following columns:

Policy.PolicyNumber
Policy.PolicyTypeID
PolicyTransaction.TransactionDate
PolicyTransaction.TransactionType

The report accepts date range parameters that are applied to the TransactionDate column. The report limits the date range parameters such that they can never be more than three days apart.

Staging Platform

All data for the HIDW database is fed into a staging database on a separate server. This staging database is named HIStage. The StageClaims table in the HIStage database is shown in the following exhibit.



	ClaimID
	ClaimDate
	PolicyHolderID
	ClaimFault
	RowState

Existing Extract, Transform, and Load (ETL) Processes


An existing ETL process identifies new and updated claim rows in the Claims system. The ETL process inserts New or Updated rows into the StageClaims table and sets StageClaims.RowState to New or Updated.

A Microsoft SQL Server 2005 Information Services (SSIS) package named WebIntelligenceCollector processes a large amount of data from a third party.

Case Study # 8 (Questions):

Question: 1


You need to meet the incremental data feed requirements for the new Claims data. Which data flow design should you use?




Answer: B

Question: 2


You need to design the physical schema to support the logical business rules for policies, policyholders, and drivers. Which schema should you use?



B.



C.



Answer: A

Question: 3

You need to reduce the number of logical page reads that occur when the Policy Transactionreport is run. What should you do?

- A. Place the Policy table and the PolicyTransaction table in separate schemas.
- B. Place the Policy table and the PolicyTransaction table in separate file groups.
- C. Add indexes to the primary key columns of the Policy table and the PolicyTransaction table.
- D. Create a single indexed view of the Policy table and the PolicyTransaction table.

Answer: D

Question: 4

You need to add a checkpoint to the WebIntelligenceCollector SSIS package that allows the package to restart at the point of failure if it fails on the Web Service task. Which task should you use?

- A. Data Flow Task 1
- B. Execute SQL Task
- C. Web Service Task
- D. Foreach Loop Container

Answer: D

Question: 5
You need to meet the SSRS deployment requirements. What should you do?


- A. Use the RsConfig utility.



- B. Use the Reporting Services Configuration tool.
- C. Use the Multiple configuration deployment option in Business Intelligence Development Studio (BIDS).
- D. Design a deployment script and execute it by using the Rs utility.

Answer: D

Question: 6

You need to design the SSIS control flow items to meet the requirements of the ETL process for the ClaimDate and ClaimFault data. Which package should you use?





Answer: B

Question: 7

You need to design a subscription delivery plan that meets the requirements of the Sales Personnel report. What should you do?

- A. Create two data-driven subscriptions. Place e-mail recipient and report parameters in one table. Place filename and path in a second table.
- B. Create two standard subscriptions. Specify the delivery method and e-mail address list or file location appropriately in each subscription.
- C. Place delivery method, e-mail address or filename and path, and report parameters in a table for each user. Create one data-driven subscription based on the data in this table.
- D. Place delivery method, e-mail address, and report parameters in a table for each user. Place filename, path, and report parameters in a second table for each user. Create one data-driven subscription based on the data in the two tables.

Answer: A

Question: 8

You need to select the edition of SQL Server 2005 that will meet all technical and business requirements. Which edition should you use?

- A. Microsoft SQL Server 2005 Express Edition
- B. Microsoft SQL Server 2005 Standard Edition
- C. Microsoft SQL Server 2005 Enterprise Edition
- D. Microsoft SQL Server 2005 Workgroup Edition

Answer: C Case Study # 9:

BACKGROUND


Wide World Importers is an import company that has clients all over the world. The company stores its reporting data in a database named WWDataStore.

EXISTING ENVIRONMENT**WWDataStore**

The WWDataStore database is populated by an Extract, Transform, and Load (ETL) process. This ETL process starts every day at 14:00 GMT and completes by 05:00 GMT. Working hours at Wide World Importers are from 09:00 GMT to 17:00 GMT. As the ETL process overlaps working hours, the report outputs may change between one execution and the next.

WWDataStore schema

The WWDataStore data model is shown in the following exhibit.



The Customer table represents a Type 2 dimension and contains over a million records. The Shipment table contains over 250 million rows. The region descriptions are used to identify geographical locations. These region descriptions reside only within a Microsoft Access database. Business users currently use a region code from the Geography table of WWDDataStore to query region descriptions.

SSAS Solution

A Microsoft SQL Server 2005 Analysis Services (SSAS) solution contains dimensions for the SalesPerson, Date, Customer, Contact, and Geography tables. It also contains a cube that includes measure groups based on the Shipment table and the SalesTarget table. The Geography, Contact, and Date dimensions exist but are not yet referenced by the cube.

BUSINESS REQUIREMENTS

Business requirements are as follows:

- A new set of reports that will include aggregations of large volumes of data.
- Aggregated reports must not run directly against the underlying WWDDataStore.
- WWDDataStore and ETL processes cannot be changed.
- SSAS cube data must be filtered and aggregated by using the region descriptions.
- The SSAS cube must represent the state of the business as of the previous day.

Case Study # 9 (Questions):

Question: 1

You need to provide the finance department with a version of the cube that only exposes the data that the department requires. What should you do?

- A. Design a new perspective within your cube. Within the perspective, enable the measures and attributes that are used by the finance department.
- B. Design a new role within the SSAS solution and move the users from the finance department into this role. Within the role, enable the measures and attributes that are used by the finance department.
- C. Design new versions of every dimension within the SSAS solution and include only the attributes that the finance department uses. Make the new dimensions available to the users from the finance department only.
- D. Modify each dimension by setting the AttributeHierarchyVisible property to true for the attributes that are used by the finance department. Within the cube, set the DisplayFolder property to true for those measures that are used by the finance department.

Answer: A

Question: 2

You need to ensure that the sales department is able to compare sales against target sales. What should you do?

- A. Design a calculated measure that shows a current against historical comparison of sales against target sales.
- B. Design a calculated measure that shows sales minus target sales. Ensure that the Parent hierarchy is set to Measures.
- C. Design a Key Performance Indicator (KPI) that shows sales in the value expression, target sales in the goal expression, and sales against target sales in the status expression.
- D. Design a Key Performance Indicator (KPI) that shows sales in the status expression, sales against target sales in the goal expression, and target sales in the value expression.

Answer: C

Question: 3

You need to ensure that repeat invocations of the same report do not produce different results during working hours. What should you do?

- A. Use Automatic MOLAP storage and set the rebuild interval to 24 hours. Enable ROLAP aggregations.
- B. Use MOLAP storage and create a scheduled task to process the dimensions and cubes after the ETL processes are completed.
- C. Use RealTime ROLAP storage to update the cache every 24 hours. Use a tracking table to indicate the status of the ETL processes and set a notification on the tracking table.
- D. Use RealTime HOLAP storage to update the cache every 24 hours. Use a tracking table to indicate the status of the ETL processes and set a notification on the tracking table.

Answer: B

Question: 4

You add the SalesPerson dimension to the cube during development. You deploy and execute aProcess Full against the SSAS solution. You then run a Client Shipment report that summarizes shipments by date and customer. You notice that the Client Shipment report previously took four seconds to run, but now takes 40seconds. You need to improve the report response time of the Client Shipment report. Which strategy should you choose?

- A. Run a Process Index against the new dimension.
- B. Rerun the Aggregation Design Wizard to recreate aggregations.
- C. Recreate aggregations by using the Usage-based Optimization Wizard.
- D. Ensure that the new dimension has the AttributeHierarchyOptimizedState property set to FullyOptimized.

Answer: B

Question: 5

You need to ensure that all members of the sales teams can access the Sales data. What should you do?

- A. Use SSAS security roles for each Sales team by applying Denied Member Sets in Dimension Data Security.
- B. Use SSAS security roles for each sales team by applying Allowed Member Sets in Dimension Data Security.

- C. Use report security in Microsoft SQL Server 2005 Reporting Services (SSRS). Combine report security with custom code to restrict data access.
- D. Use role-based security in Microsoft SQL Server 2005 Reporting Services (SSRS) to restrict access to specific folders that contain the relevant report.

Answer: B

Question: 6

You need to make region descriptions available within the SSAS solution. You need to achieve this goal by using the least amount of administrative effort. Which three steps should you perform? (To answer, move the three correct steps from the list of procedures to the answer area and arrange them in the appropriate order.)

List of Procedures

Within SSAS, create a measure group based on the Region data table.

Within SSAS, create a Region dimension that is based on the Region data table.

Add the Region data table from the Access database as a dataset to all new reports.

Add the Region data table from the Access database to the data source view of the SSAS solution.

Use the Merge Data Flow transformation that has an Integration Services solution to add Region description to WWI Data Store.

Add the Region dimension to any applicable cubes within the SSAS solution. Create an appropriate relationship between measure groups and the Region dimension.

Answer Area

Answer:

List of Procedures

Within SSAS, create a measure group based on the Region data table.

Within SSAS, create a Region dimension that is based on the Region data table.

Add the Region data table from the Access database as a dataset to all new reports.

Add the Region data table from the Access database to the data source view of the SSAS solution.

Use the Merge Data Flow transformation that has an Integration Services solution to add Region description to WWI Data Store.

Add the Region dimension to any applicable cubes within the SSAS solution. Create an appropriate relationship between measure groups and the Region dimension.

Answer Area

Add the Region data table from the Access database to the data source view of the SSAS solution.

Within SSAS, create a Region dimension that is based on the Region data table.

Add the Region dimension to any applicable cubes within the SSAS solution. Create an appropriate relationship between measure groups and the Region dimension.

Question: 7

You need to process the Customer dimension by using the minimum amount of time possible. Which processing option should you use?

- A. Process Full
- B. Process Data
- C. Process Index
- D. Process Update

Answer: D Question: 8

You need to minimize the amount of space required for the SSAS cube. You must achieve this goal while maximizing the performance of business user queries. What should you do?

- A. Within the data source view, replace the Shipment data table with a named query that orders results by descending shipment date.
- B. Create a separate perspective within your SSAS solution. Ensure that the only selected measure group is the Shipment measure group. Ensure that only the Key attributes are included from each dimension.
- C. Use ROLAP storage on the SSAS cube.
- D. Create separate partitions on the Shipment measure group. Use a query in each partition to separate data on the basis of Shipment date. Ensure that the partition that contains the present calendar month is aggregated.

Answer: D**Question: 9**

You need to modify the cube to support the Shipping Work by Location report. What should you do?

- A. Add the Geography dimension twice to the cube. Create regular relationships between the two cube dimensions and the Shipment measure group.
- B. Add the Geography dimension once to the cube. Create a many-to-many Relationship type between the Geography dimension and the measure group.
- C. Add the Geography dimension once to the cube. Create a Referenced Relationship type between the Geography dimension and the measure group.
- D. Add the Geography dimension once to the cube. Set measure group bindings to the two Geography columns within the measure group tables.

Answer: A Case Study # 10:**BACKGROUND**

Blue Yonder Airlines is an international company.


The airline operates an air miles reward program. For every flight ticket purchased, the customer earns air miles. These air miles can then be redeemed for future ticket purchases. Blue Yonder has several partner airlines involved in the same air miles reward program. This means that air miles that are accrued from partner airlines are also redeemable against Blue Yonder ticket purchases.

EXISTING ENVIRONMENT

Separate SQL Servers are deployed to support development (DEV), functional testing (QA), and user acceptance testing (UAT). The airline has started a business intelligence (BI) project that contains the following:

Data Warehouse Data Store

A Microsoft SQL Server 2005 database named BlueDW is the data store. The BlueDW schema is shown in the exhibit.



All dimensions are Type 2. At present, accrued air miles are stored as discreet values against each customer record in the AirMilesTotal column in the Customer dimension table.

Ticket and Customer Data

Ticket data is extracted from a SQL Server 2005 database named BlueTicketing. Customer data is extracted from a Customer Relationship Management (CRM) database named BlueCRM. Data from the BlueCRM system is sent to an external security agency for processing. This process results in delays that prevent the customer data from being loaded at the same time as the ticket data. The processed data is returned to Blue Yonder and needs to be loaded into the BlueCRM system. A Microsoft SQL Server Integration Services (SSIS) package named SecurityAgencyLoad loads this data into the BlueCRM system. Accrued air miles are stored for each customer in the BlueCRM system. On average, a customer adds 500 air miles each time the customer books a flight with Blue Yonder or one of its partner airlines. Every month, partner airlines submit a file named PartnerMonthlyAirmiles.csv. This file details the partner airline miles for each customer. Blue Yonder adds this data to the BlueCRM system.

Extract, Transform, and Load (ETL) Processes

BlueDWLoad is an SSIS package that transfers data from the BlueTicketing database and the BlueCRM database to the BlueDW database. A separate SSIS package named CRMTicketCleansing cleanses the ticket and customer data. The PartnerAirmilesLoad SSIS package loads the PartnerMonthlyAirmiles.csv file into the BlueDW database.

Microsoft SQL Server Analysis Services (SSAS) Solution

The TicketSales measure group is refreshed by using Incremental Processing.

Microsoft SQL Server Reporting Services (SSRS) Environment

An AirportTickets report is deployed on the SSRS solution. This report is run each morning and displays data about ticket purchases for a specific airport. A filter is used at run time to ensure that only records for a specific airport are returned. The TicketSalesByFlight report currently shows summary records for a selected flight. Blue Yonder employees want to use the TicketSalesByFlight report to view detailed flight information by expanding the summary record of a given flight.

BUSINESS REQUIREMENTS

Reporting Requirements

The BlueTicketing system is capable of handling ticket cancellations. When a cancellation occurs, the Cancelled column in the Ticket table is set to true. The BI platform must accurately handle and reflect cancelled tickets.

Case Study # 10 (Questions):

Question: 1

You need to modify the BlueDWLoad SSIS package to ensure that TicketSale records can be loaded when the Customer data is not yet available. What should you do?

- A. Load an empty Customer dimension record and the TicketSale records. When the source data is available, update the Customer dimension record with the inferred member output of the Slowly Changing Dimension SSIS task.
- B. Load the TicketSales records without inserting an empty Customer dimension record. When the source data is available, insert a new Customer dimension record from the new member output of the Slowly Changing Dimension SSIS task.
- C. Load the TicketSales records and an empty Customer dimension record. When the source data is available, insert a Customer dimension record from the historical member output of the Slowly Changing Dimension SSIS task.
- D. Load the TicketSale records and an empty Customer dimension record. When the source data is available, delete the empty Customer dimension record and insert a row with the new member output of the Slowly Changing Dimension SSIS task.

Answer: A Question: 2

You need to reduce the number of inserts into the Customer dimension table in the BlueDW database. You also need to meet the reporting requirements. What should you do?

- A. Apply banding to the AirMiles attribute.
- B. Add the AirMiles attribute to the TicketSale fact table.
- C. Place the AirMiles attribute into a dedicated dimension table.
- D. Update the AirMiles attribute directly in the Customer dimension table.

Answer: A

Question: 3

You need to generate the Schedule.xls file from the BlueDW database. What should you do?

- A. Use Microsoft SQL Server Notification Services.
- B. Use Microsoft SQL Server Analysis Services.
- C. Use Microsoft SQL Server Management Studio.
- D. Use Microsoft SQL Server Reporting Services.

Answer: D

Question: 4

You need to modify the Customer dimension table to support the design and reporting constraints. What should you do?

- A. Add a version number to the Customer dimension table.
- B. Add a current version flag to the Customer dimension table.
- C. Add a tracking table to the BlueDW database to hold customer history.
- D. Add version start time and version end time to the Customer dimension table.

Answer: D

Question: 5

You need to modify the SecurityAgencyLoad SSIS package to meet the audit trail requirements. What should you do?

- A. Configure the Microsoft Windows Event Log logging provider. Use the OnPostExecute event of the task to capture the row count to log the Event Log.

- B. Use a row count transformation to store the total row count in a user variable. Configure a logging provider to use the variable to log the Microsoft Windows Event Log.
- C. Use an Execute SQL task to write a query to write the row count to a user variable. Configure a logging provider to use the variable to log the Microsoft Windows Event Log.
- D. Use a row count transformation to store the total row count in a user variable. Use a Script task to write the value of the variable to the DTS Log. Configure a logging provider to log the script task to the Microsoft Windows Event Log.

Answer: D

Question: 6

You need to modify the PartnerAirmilesLoad SSIS package to avoid duplicate customers in the BlueCRM system. Which SSIS transformation task should you use?


- A. Sort
- B. Pivot
- C. Fuzzy Grouping
- D. Conditional Split

Answer: C


Question: 7

You need to redesign the SecurityAgencyLoad SSIS package to meet the performance requirements. Which design should you use?


A.



B.



C.



Answer: A

Question: 8

You need to redesign the TicketSalesByFlight report. Which component should you use?

- A. Bookmarks
- B. Drilldown
- C. Drillthrough
- D. Document Map

Answer: B

Question: 9

You need to modify the Executive Dashboard suite reports to allow testing to occur per the requirements. What should you do?

- A. Create a single shared data source of type XML. Replace the connection string with an expression.
- B. Create a Report database and a ReportTempDB database on each SQL Server.
- C. Install reporting services on each of the SQL Server machines.
- D. Create an embedded report-specific data source of type SQL Server. Replace the connection string with an expression. Repeat for all reports.

Answer: D