



Exam : 070-442

**Title : PRO: Designing and Optimizing Data
Access by using Microsoft SQL
Server 2005**

Ver : 09.24.07

Topic 1, National Retailers, Scenario

COMPANY OVERVIEW

Background

National Retailers operates as a retailer that specializes in vintage dolls and handmade porcelain dolls as well as action figures. The National Retailers Company has four factories: one each in Chicago, Los Angeles, Miami, and Dallas. Chicago hosts the National Retailers head quarters. Apart from the four factories, National Retailers has sales representatives that works in 50 cities across the United States. These sales representatives telecommute and usually place orders from customers with the closest factory. Salespeople are paid on commission. Commission checks are generated monthly and include commission on all orders paid for in that month.

Products and Services on offer

National Retailers operates as a retailer that specializes in vintage dolls, handmade porcelain dolls, and action figures. National Retailers sells mostly to department stores, discount stores, and specialty stores. In future National Retailers hope to expand business and thus wants to start selling their exclusive doll line on their Web site.

Future Plans

Since its inception, National Retailers has grown rapidly from humble beginnings with a single factory in Chicago to a multi-location company that currently has four factories supplying the whole of the United States. The current database instruction is unable to handle the National Retailers requirements. As a result product shipments have been delayed due to stocking levels for popular products not meeting the demand. In the mean time the less popular doll ranges have been stored for too long and some has even been shipped to customers by mistake.

To this end National Retailers wants to upgrade its current system to Microsoft SQL Server 2005. They also want to centralize operations to accommodate more effective management of the factories as well as the order fulfillment process.

The National Retailers management wants to expand business by selling their exclusive doll line on the company Web site.

EXISTING ENVIRONMENT

Existing Application and Supporting Environment

The National Retailers network consists of two Active Directory forests, namely Corporate and Manufacturing respectively. Each of the factories represents its own domain. Currently all domain controllers are running Microsoft Windows 2000 Server. There are five servers that are used as database servers. All database servers are running Microsoft SQL Server 2000. NRCH-DB02 enjoys membership of the corporate domain. All other database servers belong to the Manufacturing domain. The following table illustrates the current Database Servers and their specifications:

Server	Operating System	Processor	Memory	Hard Disk
NRCH-DB01	Windows 2000 Server, latest service pack	Pentium IV, 1.5 GHz	512 MB	40 GB
NRLA-DB11	Windows NT 4.0, latest service pack	Pentium III, 600 MHz	512 MB	30 GB
NRMI-DB12	Windows 2000 Server, latest service pack	Pentium IV, 850 MHz	256 MB	80 GB
NRDA-DB13	Windows XP Server, latest service pack	Pentium IV, 1.2 GHz	1 GB	
NRCH-DB02	Windows 2000 Server, latest service pack	Pentium III, 800 MHz	256 MB	

The Chicago head quarters have two database administrators and three network administrators. These employees vacillate between the head quarters and the factories when required.

At present the inventory application allows users to enter criteria. It also builds Transact-SQL queries dynamically by concatenating the data entered by users.

At present the order management application calls stored procedures that accept a Transact-SQL string as the only parameter.

Connectivity between office and factories

Only the Chicago office is connected to the Internet. The National Retailers factories are not currently connected to the Internet. A server named ISA-Chicago, in the Chicago office is configured with Internet Security and Acceleration Service (ISA). This ISA server separates the internal network from the Internet.

The factories are making use of dedicated links to connect to the Chicago office. Each factory is responsible for the management of its own product tracking. They all make use of the same custom application for product tracking. The following table illustrates the current database storage requirements:

Server	Storage Capacity	Data Quantity
NRCH-DB01	60 GB	4 years
NRLA-DB11	15 GB	1 year
NRMI-DB12	20 GB	2 years
NRDA-DB13	30 GB	1 year

Regularly at the end of each month, reports are generated manually and e-mailed to the Chicago office. Monthly sales and invoice aging information is included in these reports. The Chicago office handles debt collection. All customers with an outstanding balance over 60 days are also sent to the Chicago office. The accounting department makes use of a collections application that imports XML data.

The ProductLine and ProductLineDetails tables contain information about each product and the parts and components it contains. The ProductLineDetails table contains a foreign key constraint that references the Parts and Components table. The Parts and Components table contains a foreign key constraint that references the Suppliers table. The Orders and OrderDetails tables contain information about each order. The Orders table has a foreign key constraint that references the Customers table. The OrderDetails table has a foreign key constraint that references the Products table. The Products table is referenced by the ProductLine table's foreign key constraint.

Envisaged Environment

The National Retailers management wants all database servers to be upgraded to Microsoft SQL Server 2005.

The future online Sales application will be outsourced. NRCH-DB02 will host a database to support the e-commerce application. The e-commerce application will make use of Simple Object Access Protocol (SOAP) to retrieve product information and submit orders.

Consolidation of data and revision of the existing applications will be done when required.

The factories will not be provided with Internet access.

BUSINESS REQUIREMENTS

An annual growth rate of 10% per annum is envisaged for the next five years.

Accessibility of data:

Accounting data

All accounts data, i.e. accounts receivable and payable, must remain immediately accessible for a period of six months after the end of the fiscal year in which the invoice was paid. After this period, the same data must be stored for another seven year period. Usually this data will then hardly be accessed if ever.

Information about customer payments and invoices must be available in the corporate office within 24 hours of the order shipping.

Manufacturing data

Manufacturing data and product quality data must be accessible for a five year period. In reality however, data older than 6 months is rarely accessed.

Sales data

National Retailers Sales Representatives must continue to be allowed to place orders via the nearest factory. In the event of a factory unable to fulfill the order, product inventory at other factories should be checked and an e-mail order sent to the manager of the factory that has the product in stock. The manager will then verify inventory and submit the order locally.

Sales Representatives should also be allowed to check the current commissions due to them.

National Retailers wants its sales representatives to be able to place orders on a Web site. In the event of a salesperson attempts to place an order for a customer who has a balance over 60 days past due, an e-mail should be sent to the salesperson and the customer, informing them about order rejection.

Product data

Industry regulations require an audit trail for all components and parts that go into a doll or action figure. The audit trail must be able to use the UPC number on the package to trace all the components and parts, who supplied the components and parts, and what date the components and parts were shipped. Components and parts information must be entered at the time the product line is created and cannot be altered after the quality inspection. Attempts to modify components and parts information after the inspection time must be logged.

Only product lines with an intact audit trail can be shipped. If an audit trail is not available, the product line must be destroyed. The product line application includes a verification routine that checks a specific product line's audit trail. The application makes use of structured error handling to react to a product line that must be destroyed.

A product line takes two days to complete. All components and parts are collected during the first half a day. A custom application automatically logs the components and parts information to the database when components and parts are collected. The Chicago factory can produce two lines at once. They are synchronized so that the second line is always started one day after the first line.

TECHNICAL REQUIREMENTS

Online customers have to be catered for. Thus credit card numbers must be encrypted in the database and when sent across the Internet.

Availability:

The tables related to tracking a product line must be available as quickly as possible if the system fails during the factory's hours of operation. The accounting and sales tables do not require quick restoration.

The tables related to tracking a product line must support point-in-time recovery.

In the event of a factory's database server down during business hours, the inventory and order management application must be able to use one of the other factories' database servers until the local server is recovered.

Topic 1, National Retailers (11 Questions)

QUESTION 1

You work as a contract developer for National Retailers. You are currently working on the online order application the National Retailers database developer has informed you that the product parts and components will be stored in a xml data type column. National Retailers employees are allowed to request product information regarding a doll or action figure line by clicking a button on the Web form. The data will only be retrieved if the employee requests details.

You need to design the appropriate data access technologies and must thus choose the object that you need to use to store the query results for product parts and component information.

What should you do?

- A. Make use of a DataSet object.
- B. Make use of an XmlElement object.
- C. Make use of an XmlDocument object.
- D. Make use of an XmlDocumentFragment object.

Answer: D

Explanation: An XmlDocumentFragment object can store either a well-formed XML document or a fragment of a document. This would be consistent with the xml data type which has the ability to store well-formed documents as well as fragments.

1. The National Retailers management wants to expand business by selling their exclusive doll line on the company Web site

Incorrect Answers:

A: A DataSet object is used when the method returned a relational data set, akin to the method that is generated from a SELECT statement.

B: An XmlElement object is used to store the results of a SELECT ... FOR XML query, not the contents of a single value.

C: An XmlDocument object must store well-formed XML Documents. It is thus not compatible with the xml data type.

QUESTION 2

You work as the database developer for National Retailers. You need to optimize the indexing strategies and are thus designing the indexes for the Sales.Orders table. The query in the following exhibit is frequently executed, though it is not the most commonly executed query.

```
SELECT Salesrepresentative, SUM(Commission)
FROM Sales.Orders
WHERE Date BETWEEN @ startDate AND @ endDate
GROUP BY Salesrepresentative
ORDER BY Salesrepresentative
```

You need to use the appropriate statement to create the best index to accommodate this query.

What should you do?

A. Use the following statement:

```
CREATE INDEX ix_Commission
ON Sales.Orders(Salesrepresentative, Date, Commission)
```

B. Use the following statement:

```
CREATE CLUSTERED INDEX ix_Commission
ON Sales.Orders(Salesrepresentative, Date)
```

C. Use the following statement:

```
CREATE INDEX ix_Commission
ON Sales.Orders(Date)
INCLUDE (Salesrepresentative);
```

D. Use the following statement:

```
CREATE INDEX ix_Commission
ON Sales.Orders(Date, Salesrepresentative)
INCLUDE (Commission);
```

Answer: D

Explanation: The Date column is used to select the records and the Salesrepresentative column is used to group and order the records. This means that both these columns have to be key columns. The Date column is used in a BETWEEN comparison, thus it should also be the first in the query. Furthermore, it has higher selectivity than the Salesrepresentative column.

1. Sales Representatives should also be allowed to check the current commissions due to them.

Incorrect Answers:

A: The column used for equality or BETWEEN comparisons should be listed first. This must then be followed by the most selective column, then the rest of the predicate

columns in order of decreasing selectivity. And, although it is possible that you can create an index by using a computed column as the key column, it is recommended that key columns be kept as narrow as possible. This means that making use of an included column for Commission is a better option.

B: A Clustered index should have high selectivity. The Salesrepresentative column does not have high selectivity. This means that this clustered index will not be appropriate for any other queries done against the table like retrieving order information for instance. A table can only have one unique index.

C: The Salesrepresentative column is used in the GROUP BY clause. Thus it would be better suited as the key column rather than an included column. Furthermore, this is not a covering index as it does not include Commission as an included column.

QUESTION 3

You work as the database developer for National Retailers. While busy designing queries for the retrieval of data from XML sources, you are writing a script that will generate an XML file to be imported into the collections application. This application assumed XML data includes both elements and attributes. You need to make a choice of the most appropriate type of Transact-SQL statement to use to retrieve the data from the database. What should you do?

- A. Make use of the OPENXML Transact-SQL statement.
- B. Make use of the sp_xml_preparedocument stored procedure.
- C. Make use of the SELECT ... FOR XML AUTO Transact-SQL statement.
- D. Make use of the SELECT ... FOR XML PATH Transact-SQL statement.

Answer: D

Explanation: The FOR XML PATH clause of the SELECT statement will allow you to use XPath to define the structure for the XML data that is returned. This structure can include both elements and attributes.

1. The Chicago office handles debt collection. All customers with an outstanding balance over 60 days are also sent to the Chicago office. The accounting department makes use of a collections application that imports XML data.

Incorrect Answers:

A: The OPENXML statement is used to insert XML data into relational tables. Not to retrieve a resultant set formatted as XML from relational tables.

B: The sp_xml_preparedocument system stored procedure is used prior to calling OPENXML.

C: The FOR XML AUTO clause of the SELECT statement will generate a result set of nested elements and as such do not make allowance for a mix of elements and attributes.

QUESTION 4

You are busy designing the database query strategy for National Retailers that will retrieve the result set used to report on monthly sales trends. To this end you are designing the stored procedure. This stored procedure will be used by an

application that makes use of Microsoft Visual C# .NET.

Of the requirements that you need to keep in mind is that the analysts need the ability to retrieve a result set with approximately 100 records and scroll through them to gather information. These analysts will need the ability to locate records based on either factory or product ID. They do NOT update any data. You thus need to design the most appropriate cursor strategy for National Retailers.

What should you do?

- A. Create a server-side static cursor.
- B. Create a server-side dynamic cursor.
- C. Create a client-side forward-only cursor.
- D. Use a default result set and do not create a cursor.

Answer: D

Explanation: The Default Result Set caches all records in the result set to the client. This will thus reduce round trips across the network to one and prevent data from being stored in the tempdb.

1. Regularly at the end of each month, reports are generated manually and e-mailed to the Chicago office. Monthly sales and invoice aging information is included in these reports.

Incorrect Answers:

A: If you create a Server-side static cursor it will consume server resources and require a round-trip across the network each time the client fetches data. Furthermore the cursor will then be stored in tempdb.

B: A Server-side dynamic cursor does require a round-trip across the network every time a client fetches data.

C: Creating a Client-side forward-only cursor will not allow analysts to scroll through the data the way that they require to. These cursors only support the ability to fetch the next record and not to move to a specific record.

QUESTION 5

You work as the database developer for National Retailers. You are currently designing the code that will retrieve the product information for the online order application. To this end you need to determine the appropriate type of object that you should create to return to the Web application. The Web application will require that information such as product name, product description, and product price information be returned in response to a search request.

What should you do?

- A. Use a Transact-SQL stored procedure.
- B. Use an extended stored procedure.
- C. Use a table-valued user-defined function.
- D. Use a Common Language Runtime (CLR) user-defined function.

Answer: A

Explanation: The online application accesses data via a Native XML Web service. This means that you are limited to using an object that can be exposed as a Web method. A Transact-SQL stored procedure can return the required information and can be exposed as a Web method.

1. The National Retailers management wants to expand business by selling their exclusive doll line on the company Web site.
2. The future online Sales application will be outsourced. NRCH-DB02 will host a database to support the e-commerce application. The e-commerce application will make use of Simple Object Access Protocol (SOAP) to retrieve product information and submit orders.

Incorrect Answers:

B: An Extended stored procedure cannot be exposed as a Web method.

C: A Table-valued user-defined function cannot be exposed as a Web method.

D: A CLR user-defined function is not required because the data can be retrieved using Transact-SQL.

QUESTION 6

You are designing the error-handling routines for National Retailers. You are creating a function that will be used to check the audit trail for the products. This function must raise a custom error in the event of the audit trail not being verifiable. Do bear in mind that the code to check the audit trail requires a query and a cursor.

What should you do?

A. Use the following error-handling method:

```
-- Perform query and open the cursor
```

```
-- Run code that iterates through the cursor and checks the audit trail
```

```
IF @@ERROR
```

```
RAISERROR @@ERROR
```

B. Use the following error-handling method:

```
BEGIN TRY
```

```
-- Perform query and open the cursor
```

```
-- Run code that iterates through the cursor and checks the audit trail
```

```
END TRY
```

```
BEGIN CATCH
```

```
RAISERROR (@myCustomError, 15)
```

```
END CATCH
```

C. Use the following error-handling method:

```
--Perform query and open the cursor
```

```
--Fetch the first row
```

```
WHILE /*not the last row*/
```

```
IF /*Code that checks the audit trail*/
```

```
--Fetch the next row
```

```
ELSE
```

```
SET @@ERROR = @myCustomError
```

```
END WHILE
```

D. Use the following error-handling method:

```
--Perform query and open the cursor
--Fetch the first row
WHILE /*not the last row*/
IF /*Code that checks the audit trail*/
-- Fetch the next row
ELSE
RAISERROR (@myCustomError, 15)
END WHILE
```

Answer: D

Explanation: You need to pass the error message back to the application via an IF ELSE statement to check the values in a row, then calling RAISERROR in the event of the values being invalid. The RAISERROR function will allow you to send an error message up the calling chain.

1. Industry regulations require an audit trail for all components and parts that go into a doll or action figure. The audit trail must be able to use the UPC number on the package to trace all the components and parts, who supplied the components and parts, and what date the components and parts were shipped. Components and parts information must be entered at the time the product line is created and cannot be altered after the quality inspection. Attempts to modify components and parts information after the inspection time must be logged.

2. Only product lines with an intact audit trail can be shipped. If an audit trail is not available, the product line must be destroyed. The product line application includes a verification routine that checks a specific product line's audit trail. The application makes use of structured error handling to react to a product line that must be destroyed.

Incorrect Answers:

A: The requirements will not be met if you use @@ERROR after iterating through the cursor. The @@ERROR global variable contains the error number generated by the last statement. Thus it will only include an error number if the last statement to process the cursor is responsible for the error.

B: Using a TRY ... CATCH block within a function is not possible.

C: @@ERROR is a global variable that is automatically set by SQL Server every time a statement executes. Thus you cannot set @@ERROR.

QUESTION 7

You are designing the error-handling routines for the National Retailers Web application. In the event of a customer failing to enter valid payment information, you need to send a custom error to the calling application. Furthermore, this error message has to be localized to the user's language. The Web application will pass the language when it opens a connection to the database server. You thus need to decide which would be the most appropriate way to support error messages in multiple languages.

What should you do?

- A. A CASE statement should be added to the CATCH block.
Then call RAISERROR with a different error message for each language.
- B. A different error number should be used for each language version.
Then a CASE statement should be added to the trigger and call RAISERROR with the language-specific error.
- C. Add a message for each language using sp_addmessage.
A different error number should be used for each language version.
Raise the error in the @@ERROR global variable.
- D. Add a message for each language using sp_addmessage.
The same error number should be used for each language version.
Call RAISERROR from the trigger.

Answer: D

Explanation: The same error numbers should be used when adding an error message for each language. This will allow SQL Server to match the session language to the error message language when you call RAISERROR. The sp_addmessage stored procedure is used to add an error message.

1. The National Retailers management wants to expand business by selling their exclusive doll line on the company Web site.
2. The future online Sales application will be outsourced. NRCH-DB02 will host a database to support the e-commerce application. The e-commerce application will make use of Simple Object Access Protocol (SOAP) to retrieve product information and submit orders.

Incorrect Answers:

- A: A CASE statement should not be added to the CATCH block and then call RAISERROR with a different error message for each language. This will make maintenance quite difficult since you will then be required to modify the code that calls RAISERROR if a message changes or you need to add a language.
- B: There is no need to use a different error number for each different language version and then add a CASE statement to the trigger. This option will work, though it will be difficult to maintain since the trigger will then have to be modified each time a language is added. It will also require more programming effort instead of allowing SQL Server to choose the right language for the error message.
- C: If you raise the error in @@ERROR, then the error associated with the last statement will be raised and not the custom error message. Also, there is no need to define a different error number for each language version.

QUESTION 8

You are designing the error-handling routines for National Retailers. You thus need to design the code that will validate the audit trail for a product batch.

Users that do not enjoy membership of the sysadmin role will run the fn_ValidateBatch function. You need to use the appropriate statements to ensure that in the event of an invalid audit trail, an error is raised to the calling application and an event is written to the application log.

What should you do?

A. Use the following statements:

```
sp_addmessage (@msgnum=50888, @severity=15, @msgtext="Batch" + @b "cannot be  
validated.")
```

and

```
RAISERROR(50888, @batchnum, WITH LOG)
```

B. Use the following statements:

```
sp_addmessage (@msgnum=50888, @severity=16, @msgtext="Batch %b cannot be  
validated.")
```

and

```
RAISERROR (50888, WITH SETERR, @batchnum)
```

C. Use the following statements:

```
sp_addmessage (@msgnum=50888, @severity=16, @msgtext="Batch %b cannot be  
validated.")
```

and

```
RAISERROR (50888, %batchnum)
```

D. Use the following statements:

```
sp_addmessage (@msgnum=50888, @severity=16, @msgtext="Batch %b cannot be  
validated.")
```

and

```
RAISERROR(50888, @batchnum, WITH LOG)
```

E. Use the following statements:

```
sp_addmessage (@msgnum=50888, @severity=21, @msgtext="Batch %b cannot be  
validated.")
```

and

```
RAISERROR(50888, @batchnum, WITH LOG)
```

Answer: D

Explanation: By specifying the 16 severity level, you ensure that the users who do not belong to the sysadmin role have the ability to raise the error. Also this option displays the correct syntax for using arguments in a message text string. By default, only messages with a severity level of over 19 are logged to the application log. This can be overridden by specifying WITH LOG when calling RAISERROR.

1. Industry regulations require an audit trail for all components and parts that go into a doll or action figure. The audit trail must be able to use the UPC number on the package to trace all the components and parts, who supplied the components and parts, and what date the components and parts were shipped. Components and parts information must be entered at the time the product line is created and cannot be altered after the quality inspection. Attempts to modify components and parts information after the inspection time must be logged.

Incorrect Answers:

A: This option would be correct except for the severity level of 15. By default, only messages with a severity level of over 19 are logged to the application log. This option does not offer any overriding probabilities to have the error logged to the application log.

B: This option is only partly correct. But you should not raise the error by specifying

SETERR. This will result in @ERROR and ERROR_NUMBER to be set to 50000. You do not make use of a concatenated string to use an argument in the message text.

C: The RAISERROR (50888, \$batchnum) is the incorrect syntax. A local variable starts with an @ and not a %. Also this option will not cause the application to be logged to the application log.

E: You should not specify a severity of 21. Messages with this severity level can only be raised by members of the sysadminrole.

QUESTION 9

You work as the database developer for National Retailers. You are currently designing the indexes that will be used for the OnlineSales.Customers table. The following exhibit illustrates the most common query that will be used.

```
SELECT FirstName, LastName, Address, City, State, Zip
FROM OnlineSales.Customers
WHERE CustomerID = @emailAddress
```

You thus need to select the best statement that should be used to create this index that will accommodate this query.

What should you do?

A. Use the following statement:

```
CREATE INDEX ix_custInfo
ON OnlineSales.Customer(CustomerID)
INCLUDE (FirstName, LastName, Address, City, State, Zip);
```

B. Use the following statement:

```
CREATE INDEX ix_custInfo
ON OnlineSales.Customer(LastName, FirstName)
INCLUDE Address, City, State, Zip;
```

C. Use the following statement:

```
CREATE CLUSTERED INDEX ix_custInfo
ON OnlineSales.Customer(CustomerID, FirstName, LastName, Address, City, State,
Zip);
```

D. Use the following statement:

```
CREATE CLUSTERED INDEX ix_custInfo
ON OnlineSales.Customer(CustomerID)
INCLUDE (LastName, Address, FirstName, City, Zip, State);
```

Answer: A

Explanation: Making use of this statement will create an index with a small unique key that covers the entire query. For optimal performance a small key size that includes the columns used in the

WHERE clause and other predicates or in joins are recommended. In this scenario the only required key column would be the CustomerID column. Then you can include the rest of the columns as non-key columns. This means that you need to create the index as a non-clustered index because non-key included columns are not supported for clustered indexes.

1. The National Retailers management wants to expand business by selling their exclusive doll line on the company Web site.

Incorrect answers:

B: This index would not be used for the query because the CustomerID column is not listed as a key column.

C: This statement, if created will result in an index with 344 bytes in the key column. Though it is possible, it is not optimal since it is better to create an index with a smaller key column and included columns.

D: This is incorrect because you cannot make use of included columns in a clustered index.

QUESTION 10

National Retailers require a database that will support the e-commerce application. You need to find a way to provide the e-commerce application with information about the existing inventory. The database must determine the nearest factory that has the product ordered in stock and send the order to that specific factory to have it fulfilled.

You thus need to determine how the inventory information should be made available to the e-commerce application while providing the best possible performance when placing the order.

What should you do?

A. Query the database server at each factory using OPENROWSET from a stored procedure run on NRCH-DB02.

B. At each factory implement an HTTP endpoint on the database server and expose a Web method that checks inventory.

C. Replicate the data from the database servers at the different factories to the NRCH-DB02 server using merge replication.

D. Use a linked server for the database server at each factory at NRCH-DB02.

Answer: C

Explanation: Using merge replication to replicate all data from the factories to NRCH-DB02 will make the data available locally and improve query performance when an order is placed. This is the best type of replication to use in this case because there could be conflicts by orders placed over the Web and order places by the Sales representatives at each factory. Also in a case where the different sites/factories are not well-connected making use of merge replication makes sense. And in this case the factories are connected to the Chicago office by demand-dial links.

1. The National Retailers management wants to expand business by selling their exclusive doll line on the company Web site.

2. The future online Sales application will be outsourced. NRCH-DB02 will host a database to support the e-commerce application. The e-commerce application will make use of Simple Object Access Protocol (SOAP) to retrieve product information and submit orders.

Incorrect answers:

A: Making use of OPENROWSET to query each database server at each factory from a stored procedure on NRCH-DB02 will result in one or more remote queries to be sent across the network for each product ordered.

B: Making use of an HTTP endpoint will result in more overhead than is required. There is already a demand-dial connection between the Chicago office and the different factories. Thus there is no firewall that restricts incoming protocol to only HTTP. Making use of a SOAP request will result in even more overhead than a remote query.

D: Making use of a linked server for the database server at each factory and performing a remote query is not the best solution because each product that is ordered will cause at least one, and possible more, remote queries to be sent across the network.

QUESTION 11

National Retailers requires an order fulfillment application. This application needs to perform several tasks. These tasks include:

1. Displaying pending orders
2. Allowing order items to be marked as shipped, back-ordered, of out-of-stock
3. Displaying and allowing for updates to inventory levels

You make use of Microsoft Visual Basic .NET to implement the application. Now you need to select the most optimal way to access the data.

What should you do?

A. Make use of Microsoft SQL Native Client Library and create two separate connections.

B. Make use of Microsoft SQL Native Client Library and create three separate connections.

C. Make use of Microsoft Data Access Components (MDAC) and enable Multiple Active Result Sets (MARS) on a single connection.

D. Make use of Microsoft SQL Native Client Library and enable Multiple Active Result Sets (MARS) on a single connection.

Answer: D

Explanation: MARS allows one to establish multiple open queries on a single connection. This is better than establishing separate connections because a connection inevitably results in the usage of resources on the server. However, you will require using SQL Native Client Library to be able to make use of MARS.

1. In the event of a factory unable to fulfill the order, product inventory at other factories should be checked and an e-mail order sent to the manager of the factory that has the product in stock. The manager will then verify inventory and submit the order locally.

Incorrect Answers:

A: If you do not make use of MARS, you will need three connections. If you do make use of MARS, you will only use a single connection.

B: When using earlier versions of SQL Server, you would have to have three separate connections. However, since Microsoft SQL Server 2005 supports MARS, you are allowed to issue multiple queries on a single connection. Thus three separate connections

are not required.

C: MDAC does not support MARS queries and will thus not work in this case.

Topic 2, Stanford Finance, Scenario

BACKGROUND:

Company overview

Stanford Finance is a private company that specializes in the provision of investment and asset management services for its clients. Stanford Finance has ten offices located in New York, Boston, Chicago, Atlanta, London, Manchester, Madrid, and Lisbon. Each office is currently responsible for maintaining its own customers. Stanford Finance offers its customers a number of investment services, including investment in real estate, equities, and private share listing.

Planned Changes

The company wants to create an online investment tracking system that will centralize investment portfolios at the eight offices. The company also wants to implement automated investment through its public Web site.

EXISTING ENVIRONMENT:

The Stanford Finance network consists of a single Windows Server 2003 Active Directory domain that is protected by a firewall. The domain controller for the network is named SF-DC01 and is located behind the firewall in New York.

The network also contains a SQL Server 2005 database server named SF-DB04 that is protected by the firewall. SF-DB04 hosts two databases named SF_Accounts and SF_Personnel. Both Microsoft Windows and Microsoft SQL Server authentication is configured on SF-DB04. Each employee who needs access to SF-DB04 has a login. Guest access to the online investment system is allowed through an application role. Three database roles named Managers, Consultants, and Posters have also been created.

BUSINESS REQUIREMENTS:

General Requirements

Each office offers four different kinds of investments: real estate, equities, private share listing, and derivatives. Each type of investment can be either fixed deposit or monthly deposit fixed term. Prospective investors must be able to search for investment opportunities by city based on type of investment, minimum required deposit, and other features. Prospective investors must be able to enter a word or phrase and SQL Server should return a match of any tense of that word.

Investors must provide a credit card, address, e-mail, and phone number when applying online. This information must be kept confidential at all times. Only employees in the Stanford Finance Accounting department should have access to the investors' credit card information. Also, when a potential investor makes an online investment, the investment opportunity must be reserved immediately so that it is not available to other potential investors. If the investor provides his or her credit card number, the credit card must be verified using a third-party Component Object Model (COM) component. If the investor's credit card is declined, the investor should be sent an e-mail and the investment must be rolled back.

Investment consultants at an office must also be able to access the investment database for that office from a Web application. They must have up-to-date information about

investment availability whenever they query for availability.

Commission on investments differs from one city to the other and the company often offers special rates at specific offices. When an investor cashes out his or her investment, the commission must be based on the initial commission charged when the investor signed up. Only the office manager can change commission rates, but only for the office they manage.

When an investor cashes out, they can pay commission by cash, check, or credit card.

Commission should not be deducted from the investment pay out. If they cash out online, they must pay by credit card. The online system should provide five questions about investor satisfaction but investors can choose not to respond to these questions. The investor's response to the questions must be saved to the database.

TECHNICAL REQUIREMENTS:

Application Development

The Stanford Finance wants to use Visual Studio 2005 and ASP.NET 2.0 to develop Web applications for the company.

Performance

Investment consultants must be able to determine the availability of an investment opportunity using as few round trips to the server as possible.

When an investor provides his or her credit card number, the credit card must be verified and authorized using a third-party Component Object Model (COM) component. When the credit card authorization system is busy or the lines are down, the investor should not be required to wait while their credit card is authorized. They should be sent an e-mail when the credit card is verified or declined.

Security

SF-DB04 must be protected against SQL Injection attacks.

Topic 2, Stanford Finance (14 Questions)

QUESTION 1

You need to design the procedure that will reserve investments. The procedure must be used by both Web pages that allow potential investors to invest online and Web pages that allow investment consultants to secure investments for their customers.

You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What should you do?

A. Start a transaction.

Decrement the investment opportunity availability.

Add the investment reservation.

Call the credit card authorization component.

Commit or rollback the transaction.

B. Start a transaction.

Add the investment reservation.

Decrement the investment opportunity availability.

Send a credit card authorization message to a Service Broker queue.

Commit the transaction.

Write a queue reader that reverses the changes if credit card authorization fails.

C. Decrement the investment opportunity availability.

Start a transaction.

Add the investment reservation.

Call the credit card authorization component.

Commit or rollback the transaction.

D. Send a credit card authorization message to a Service Broker queue.

Write a queue reader that starts a transaction.

Decrement the investment opportunity availability.

Add the investment reservation.

Commit or rollback the transaction.

Answer: B

Explanation: You need to reserve the investment without waiting for credit card authorization. Therefore, you should use a Service Broker queue to allow the transaction to be committed immediately while a separate asynchronous process verifies the credit card data.

Incorrect Answers:

A: You need to commit the transaction without waiting for the credit card authorization component to complete.

C: Performing changes outside of the transaction will prevent you from being able to roll the changes back should the credit card authorization fail.

D: You should use the Service Broker queue to reverse the transaction after it has been committed, rather than use to Service Broker queue to commit the transaction and then reverse it later.

QUESTION 2

You need to design an indexing strategy for the online investment system that will provide the best performance. The initial indexes have already been created. You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What could you do?

A. Run sys.dm_db_missing_index_group_stats, or use SQL Server Profiler with the Tuning template to capture a workload and use the Database Engine Tuning Advisor to analyze the workload.

B. Run sys.dm_db_index_usage_stats, or use SQL Server Profiler with the Tuning template to capture a workload and use the Database Engine Tuning Advisor to analyze the workload.

C. Run sys.dm_db_index_operational_stats, or use SQL Server Profiler with the Replay template to capture a workload and use the Database Engine Tuning Advisor to analyze the workload.

D. Run sys.dm_db_index_usage_stats, or use SQL Server Profiler with the Replay template to capture a workload and use the Database Engine Tuning Advisor to analyze the workload.

Answer: A

Explanation: To get a list of recommended indexes, you can either use the sys.dm_db_missing_index_group_stats dynamic management view or use SQL Server Profiler with the Tuning template to capture a workload and use the Database Engine Tuning Advisor to analyze the workload. The sys.dm_db_missing_index_group_stats dynamic management view provides less detail than the SQL Server Profiler.

Incorrect Answers:

B: The sys.dm_db_index_operational_stats dynamic management view returns information about indexes that are currently in use. It does not provide recommendations for additional indexes.

C: The sys.dm_db_index_operational_stats dynamic management view returns information about index locks and latches. It does not provide recommendations for additional indexes. You also cannot use Database Engine Tuning Advisor to analyze a workload based on the Replay template.

D: The sys.dm_db_index_operational_stats dynamic management view returns information about indexes that are currently in use. It does not provide recommendations for additional indexes. You also cannot use Database Engine Tuning Advisor to analyze a workload based on the Replay template.

QUESTION 3

You need to design the index strategy for the InvestmentTypes table. You need to ensure that your solution provides the best possible performance when potential investors search for investment opportunities based on features. You need to ensure that your solution meets Stanford Finance's technical and business requirements. What should you do?

- A. Create a full-text index on the InvestmentTypeID and Description columns.
- B. Create a non-clustered index on the InvestmentTypeID and Description columns.
- C. Create a unique clustered index on the InvestmentTypeID column and a full-text index on the Description column.
- D. Create a unique clustered index on the Description column and a non-clustered index on the InvestmentTypeID column.
- E. Create a unique clustered index on the InvestmentTypeID and Description columns.

Answer: C

Explanation: You need to implement full-text indexing on the table; but, before you can use a full-text index, you must first implement a unique index. For best performance, you should use the smallest column for the unique index. You can then use the unique index as the key index for the full-text index.

Incorrect Answers:

A: Before you can use a full-text index, you must first implement a unique index. The unique index is used as the index key for the full-text index.

B: A non-clustered index will not allow users to search by features. You should implement a full-text index to enable users to search by features.

D: For best performance, you should use the smallest column for the unique index. You should also implement a full-text index to enable users to search by features.

E: A unique clustered index can only be created on one column. You should also implement a full-text index to enable users to search by features.

QUESTION 4

You need to design a stored procedure that returns the top investment consultant for the month based on the highest investment amount for a single investment and the top investment consultant for the month based on total invested amount for investments secured by the consultant in that month. You must ensure that the stored procedure is able to return the data even if the Investments table is locked by another session. Your solution must allow for maximum concurrency.

What statement should you do?

- A. Set the transaction isolation level to **SERIALIZABLE**.
- B. Set the transaction isolation level to **SNAPSHOT**.
- C. Set the transaction isolation level to **READ UNCOMMITTED**.
- D. Set the transaction isolation level to **REPEATABLE READ**.

Answer: B

Explanation: You should use the snapshot isolation level. Snapshot isolation prevents both uncommitted dependencies (dirty reads) and nonrepeatable reads. It uses row versioning by reading data into tempdb and accessing data from there. It is also a lower isolation level than Serializable and therefore less likely to result in blocking locks and deadlock conditions.

Incorrect Answers:

A: Serializable isolation prevents both uncommitted dependencies (dirty reads) and nonrepeatable reads but it serializable is also the highest isolation level and is the most likely to result in deadlocks.

C: Read uncommitted isolation is the least likely to cause deadlock conditions, but prevents neither uncommitted dependencies (dirty reads) nor nonrepeatable reads. This may result in data inconsistencies.

D: Repeatable read isolation will take shared locks; therefore the transaction will be blocked if the Investments table is locked by another session.

QUESTION 5

You are designing the online investment system. During testing you discover that the system sometimes times out when online selection changes are made. You check the Windows performance log on SF-DB04 and discover a long processor queue during heavy use. You need to determine which queries are being executed.

What should you do?

- A. Create a workload that includes the Investments.LoggedTime column and analyze the

workload using SQL Server Profiler.

B. View the SQL Server log and the performance log using Event Viewer.

C. Create a performance log that includes SQL Server:Plan Cache Object counter for the SQL Plans instance.

D. Use SQL Server Profiler with the Replay template to capture a workload and use the Database Engine Tuning Advisor to analyze the workload.

Answer: A

Explanation: You should create a workload that includes a column that records the time and a performance log. You can then use SQL Server Profiler to correlate the workload with a performance log taken at the same time. This will indicate which queries were executing when a resource usage spike occurred.

Incorrect Answers:

B: Event Viewer cannot be used to identify which queries are executed and when they are executed.

C: The SQL Server:Plan Cache Object counter records the frequency with which the cache is accessed. It does not indicate which queries were executing.

D: You also cannot use Database Engine Tuning Advisor to analyze a workload based on the Replay template.

QUESTION 6

You need to design a Web application that will be used by investment consultants. Your Web application will cache data regarding investment availability. You want to use the best method to cache the data. You also need to ensure that your solution meets Stanford Finance's technical and business requirements. What should you do?

A. Set the sqldependency attribute for the Web page to AvailableInvestments.

B. Add a button to the Web page and have the button call the Response.RemoveOutputCacheItem function.

C. Call the Response.RemoveOutputCacheItem function in the OnPageLoad subroutine of the Web page.

D. Set the outputcache duration attribute for the Web page to 1.

Answer: A

Explanation: By setting the sqldependency attribute to AvailableInvestments, will inform the page each time the data in the Available column of the Investments table changes. This will render the cached data as invalid and will cause the data to be refreshed the next time the data is displayed on the Web page.

Incorrect Answers:

B: Using a button to refresh the data will require that the investment consultant manually refresh the data. This could cause unnecessary round trips to the database server and will have a negative impact on performance.

C: Refreshing the data each time the page loads will not ensure that the data changes that

occur after the page has loaded is reflected in the cache.

D: Setting the outputcache duration attribute to 1 will cause the cache to expire every second. This will cause unnecessary round trips to the database every second.

QUESTION 7

You need to design a Web page that allows office managers to add and manage logins and database users. You need to select the appropriate object model. You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What should you do?

- A. User the Microsoft Data Access Components (MDAC).
- B. User the SQL Distributed Management Objects (SQL-DMO).
- C. User the SQL Native Client (SQLNCLI).
- D. User the SQL Management Objects (SMO).

Answer: D

Explanation:

The Stanford Finance wants to use Visual Studio 2005 and ASP.NET 2.0 to develop Web applications for the company. The SQL Management Objects (SMO) can be used to perform server administration for a .NET programming language.

Incorrect Answers:

A: Microsoft Data Access Components (MDAC) is a legacy client library that is used for data access. It is not used to perform server administration for a .NET programming language.

B: SQL Distributed Management Objects (SQL-DMO) is a COM library that can be used when programming in legacy languages, such as Visual Basic 6.0 or C++.

C: SQL Native Client (SQLNCLI) is a client library that is used for data access. It is not used to perform server administration for a .NET programming language.

QUESTION 8

You need to design stored procedure that will return the average length of an investment at each office. You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What Transact SQL query should you use?

- A.

```
SELECT o.City, AverageInvestmentPeriod =  
AVG(DateDiff(day,i.StartDate,i.EndDate))  
FROM Investments AS i  
INNER JOIN Offices AS o  
ON o.OfficeID = i.OfficeID  
GROUP BY o.OfficeID
```
- B.

```
SELECT o.City, AverageInvestmentPeriod =  
AVG(DateDiff(day,i.StartDate,i.EndDate))  
FROM Investments AS i
```

```
INNER JOIN Offices AS o
ON o.OfficeID = i.OfficeID
GROUP BY o.City
C. SELECT o.City, AverageInvestmentPeriod = AVG(i.StartDate - i.EndDate)
FROM Investments AS i
INNER JOIN Offices AS o
ON o.OfficeID = i.OfficeID
GROUP BY o.OfficeID
D. SELECT o.City, AverageInvestmentPeriod = AVG(i.StartDate - i.EndDate)
FROM Investments AS i
INNER JOIN Offices AS o
ON o.OfficeID = i.OfficeID
GROUP BY o.City
```

Answer: B

Explanation: The DateDiff function can be used to return the length of an investment period and the AVG aggregate function can be used to average the length of an investment period returned by the DateDiff function. When you use the AVG aggregate function, you can only include the aggregate and columns in the GROUP BY clause; therefore you need to group by city name.

Incorrect Answers:

A: When you use the AVG aggregate function, you can only include the aggregate and columns in the GROUP BY clause; therefore you cannot group by o.CityID as the o.CityID column is not retrieved by the SELECT statement.

C: You cannot use the minus operator on a datetime or smalldatetime column.

D: You cannot use the minus operator on a datetime or smalldatetime column. Also, when you use the AVG aggregate function, you can only include the aggregate and columns in the GROUP BY clause; therefore you cannot group by o.CityID as the o.CityID column is not retrieved by the SELECT statement.

QUESTION 9

You need to design a security feature that will ensure that the database user has the permission to change the commission rate. You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What code should you use?

A. IF IS_MEMBER('Consultants')

<allow the changes>

ELSE

<do not allow the changes>

B. IF IS_MEMBER('Posters')

< do not allow the changes>

ELSE

< allow the changes>

C. IF IS_MEMBER('Posters')

```
<allow the changes>
ELSE
<do not allow the changes>
D. IF IS_MEMBER('Managers')
<allow the changes>
ELSE
<do not allow the changes>
```

Answer: D

Explanation: Only office managers should be able to change commission rates; therefore your code should check if the user is a member of the 'Managers' database role and allow the changes if the user is a member of the 'Managers' database role. The IS_MEMBER function can be used to check membership to a database role.

Incorrect Answers:

A, C: Only office managers should be able to change commission rates; therefore your code should check if the user is a member of the 'Managers' database role rather than the 'Consultants' or 'Posters' database roles.

B: This code check if the user is a member of the 'Posters' database group and denies them the ability to make the changes. If the y arte not members of the 'Posters' database role, the changes are allowed. However, only office managers should be able to change commission rates.

QUESTION 10

You need to design a parameterized stored procedure that will allow users to search for investment opportunities based on keywords. The stored procedure will accept Keyword and OfficeID as input parameters. You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What stored procedure should you use?

A. SELECT t.Description, t.CommissionRate
FROM InvestmentType AS t
JOIN AvailableInvestments AS a
ON a.OfficeID = @officeid
WHERE CONTAINS (t.Description, 'FORMSOF(INFLECTIONAL, @keyword)')
AND a.NumRealEste > 0 OR a.NumStock > 0

B. SELECT t.Description, t.CommissionRate
FROM InvestmentType AS t
INNER JOIN AvailableInvestments AS a
ON a.OfficeID = @officeid
WHERE FREETEXT (t.Description, @keyword)
AND a.NumRealEste > 0 OR a.NumStock > 0

C. SELECT t.Description, t.CommissionRate
FROM InvestmentType AS t
INNER JOIN AvailableInvestments AS a
ON a.OfficeID = @officeid

```
WHERE t.Description, LIKE ('FORMSOF(INFLECTIONAL, @keyword)')
AND a.NumRealEste > 0 OR a.NumStock > 0
D. SELECT t.Description, t.CommissionRate
FROM InvestmentType AS t
JOIN AvailableInvestments AS a
ON a.OfficeID = @officeid
WHERE CONTAINS (t.Description, LIKE (@keyword) )
AND a.NumRealEste > 0 OR a.NumStock > 0
```

Answer: A

Explanation: To use full text searches to retrieve results for any tense of a word, you should use the CONTAINS clause with the FORMSOF function.

Incorrect Answers:

B: FREETEXT is used to return synonyms of a word, not different tenses.

C: The LIKE clause does not accept a FORMSOF function.

D: You cannot use LINK in a CONTAINS clause.

QUESTION 11

You need to collect performance statistics for the online investment system. You want to collect baseline performance data before and after tuning the database to compare query execution statistics. You are particularly interested in identifying the time it takes for queries to execute before and after tuning, and the worst performing queries before and after tuning. What should you do?

- A. Run the Windows Performance applet and collect a counter log abased on SQL Server performance objects.
- B. Run the SQL Server Profiler with the Replay template and collect the statistics as a trace table.
- C. Run the SQL Server Profiler with the Duration template and collect the statistics as a trace table.
- D. Run the SQL Server Profiler with the Tuning template and collect the statistics as a trace table.

Answer: C

Explanation: The SQL Server Profiler with the Duration template allows you to identify query duration and the worst performing query. You must collect trace tables for before and after tuning so that you can compare the two.

Incorrect Answers:

A: Windows Performance counters do not collect statistics on the execution of individual Transact-SQL queries.

B: The SQL Server Profiler with the Replay template collects data that can be used for replay. It does not provide statistics on the execution of queries.

D: The SQL Server Profiler with the Tuning template collects information as a workload

file that can be used by the Database Engine Tuning Advisor to tune the database. It does not provide statistics on the execution of queries.

QUESTION 12

You need to design the functionality for the online investment system. You need to implement a data access method that you can use to submit queries to SQL Server 2005. You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What should you do?

- A. Create a Common Language Runtime (CLR) function.
- B. Create a Transact-SQL user-defined function.
- C. Create ad hoc queries by using the DataSet and DataSource objects.
- D. Create ad hoc queries by using the SqlConnection and SqlCommand objects.

Answer: B

Explanation: You need to protect the database server against SQL injection attacks; therefore you should create Transact-SQL user-defined functions.

Incorrect Answers:

A: You should only use Common Language Runtime (CLR) functions when a Transact-SQL user-defined function cannot perform the required tasks. In this scenario, a Transact-SQL user-defined function can and should be used.

C, D: Ad hoc queries are vulnerable to SQL Injection attacks and should not be used.

QUESTION 13

You need to design a parameterized stored procedure that will decrement the AvailableInvestments table when a real estate investment is secured. The stored procedure must accept the OfficeID, InvestmentTypeID, StartDate and EndDate as input parameters. You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What stored procedure should you use?

- A. UPDATE AvailableInvestments
SET NumRealEstate = NumRealEstate - 1
WHERE OfficeID = @officeid
AND InvestmentTypeID = @investmenttypeid
AND date BETWEEN @startdate and @enddate
- B. UPDATE AvailableInvestments
SET NumRealEstate = NumRealEstate - 1
WHERE OfficeID = officeid
AND InvestmentTypeID = investmenttypeid
AND date BETWEEN startdate and enddate
- C. UPDATE AvailableInvestments
SET NumRealEstate = NumRealEstate - 1
WHERE OfficeID = @officeid


```
AND InvestmentTypeID = @investmenttypeid
AND date >= @startdate
AND date <= @enddate
D. DECLARE @date
WHILE @date < EndDate
UPDATE AvailableInvestments
SET NumRealEstate = NumRealEstate - 1
WHERE OfficeID = @officeid
AND InvestmentTypeID = @investmenttypeid
AND date = @date
BREAK
```

Answer: A

Explanation: You need to protect the database server against SQL injection attacks; therefore you should create Transact-SQL user-defined functions.

Incorrect Answers:

B: Input parameters require the 'at' symbol (@) as the first character.

C: You cannot use the compare operators (> < =) with a datetime or smalldatetime column.

D: Executing the change in a loop is not as efficient as executing that changes in a single statement.

QUESTION 14

You need to design the functionality for the online investment system. You need to use ADO.NET objects to display the investment opportunities that match the user's search queries. You want to use data bound controls. You need to ensure that your solution meets Stanford Finance's technical and business requirements.

What should you do?

A. Use SqlConnection and DataSet objects.

B. Use SQLDataSource objects.

C. Use DataReader objects.

D. Use SqlCommand, DataSet and DataAdapter objects.

Answer: B

Explanation: The SQLDataSource objects allow you to simplify data binding in a Web application.

Incorrect Answers:

A, D: SqlConnection, SqlCommand, DataSet and DataAdapter objects were required in ASP.NET 1.0 and ASP.NET 1.1. SQLDataSource simplify these objects in ASP.NET 2.0.

C: DataReader objects retrieve data from the database each time a record is fetched and is only suitable for well-connected applications. In this scenario, offices are connected to the database via Internet connections that may go down.

Topic 3, Bilco Engineering, Scenario

BACKGROUND:


Company overview

Bilco Engineering is a government owned west wing military project which is responsible for creating, reverse engineering and modifying advanced weapon systems and is the primary source for weapons systems in locations throughout North America. Bilco Engineering has recently started contracting for weapons with three ammunition weapon stores named National Weapon Trade, Real-Weapon Technology and Weapon Skills Learning who each have their own administrative and instructional staffs and manages its own systems.

Bilco Engineering knows that the all the stores offer traditional weapons training and Real-Weapon Technology with Weapon Skills Learning additionally offer online training and short-duration satellite demonstration classes in non-traditional locations. The employees at the three stores all come from different backgrounds, but the majority of the employees are adults improving job skills or looking to make career changes.

TECHNICAL ENVIRONMENT:

Bilco Engineering is currently consisting of only 20 distribution centers with each location having its own internal network that includes a database server running SQL Server 2005. Bilco Engineering has the internal network isolated from the Internet by a perimeter network whilst the Web servers are hosted in the main office in Atlanta. The Bilco Engineering Inc main network infrastructure is shown below:



Bilco Engineering has additionally ensured that the Web servers are the only servers located in a private perimeter network whilst each location maintains its own inventory, with the database servers configured as federated database servers. Bilco Engineering has recently ensured that the Web servers host local copies of all offices' inventories in a custom designed database in a local SQL Server instance and civilian registration information which is also written to the local database servers.

Bilco Engineering has configured the Atlanta office to support the company's operations servers and two distribution warehouse servers. Bilco Engineering makes use of the main office for maintaining information about all weapons and defense systems information and a master registered civilian list. Bilco Engineering has identified the registered civilians as either visitors or testers-only. The Bilco Engineering military left wing project currently has one database server used as a weapons archival records server at the main office. The SQL Server used as the weapons archival records server is also used as a data source for data transform and load to a business intelligence server.

Bilco Engineering Inc has made sure that all the internal network database servers are deployed as Active Directory member servers and that the servers in the Atlanta main

office are configured for Windows authentication only. Bilco Engineering has decided to have their Web servers configured as stand-alone servers. Bilco Engineering has additionally configured the SQL server on the Web servers are configured for mixed mode authentication.

TECHNICAL REQUIREMENTS:

Bilco Engineering wants to have their left wing government military project company to be identifying the requirements for local office Web servers which should take place after local deployment to ensure each Web server maintains a current copy of the company's inventory without cost information. Bilco Engineering has also decided to configure using the simple recovery model is enough protection for the Web server database server instances.

Bilco Engineering wants to ensure that all the deployed servers to any local network must make use of Active Directory user accounts in order to access to the server. Bilco Engineering also wants to have the user credentials to be protected against unauthorized disclosure.

All of the left wing government military project company Bilco Engineering database servers are configured with the hard disk configuration shown below:

1. 80 GB IDE hard disk as drive C:
2. 80 GB IDE hard disk as drive D:
3. 270 GB SCSI RAID 5 storage unit as drive E:

Bilco Engineering has recently announced that they require having all the company local database servers to automatically report serious errors to the Atlanta main office and immediately paging all IT staff members is imperative if and when errors occur. Bilco Engineering has also expressed concern that the IT staff requires periodic information about selected resource use.

Bilco Engineering has decided that they will make use of replication on the local database servers for provide inventory updates to the Web servers whilst ensuring that any orders generated on the Web servers are submitted to the local database servers using the locally deployed Web services configured at each local network.

Bilco Engineering has planned to have all the Web servers moved to the local networks as soon as possible. Bilco Engineering will have the Web servers then moved to the local offices after which the local offices will use the same network configuration currently used in Atlanta.

Bilco Engineering has the database servers shown below at the Atlanta main office:

1. BILCO-SR01
2. BILCO-SR02
3. BILCO-SR03
4. BILCO-SR04

Bilco Engineering makes use of BILCO-SR02 to contain the information about the weapons systems records, secure employee records, civilian credit information, and other sensitive data. Bilco Engineering additionally makes use of BILCO-SR01 which maintains the summary civilian records, purchase order records, consolidated civilian orders, and other information needed for daily operations. Bilco Engineering makes use of BILCO-SR03 and BILCO-SR04 as distribution centre servers which are configured like the servers for the other local office database servers.

Bilco Engineering has recently announced that the server BILCO-SR01 will have one

database named WeaponTech. BILCO-SR01 is configured with each filegroup containing a single read-write file.

Bilco Engineering has decided that the WeaponTech database shall be configured for the full recovery model and backups are run as shown below:

1. Bilco Engineering runs a full database backup is run on the first Saturday of each month.
2. Bilco Engineering runs a full backup of the Primary filegroup is run on Sunday night.
3. Bilco Engineering runs a full backup of the Secondary filegroup is run on Wednesday night.
4. Bilco Engineering runs a differential backup of each filegroup is run before start of business each morning.

Bilco Engineering has recently revealed that the Atlanta main office database servers are required to be available at all times whilst potential single points of failure should be minimized.

BUSINESS REQUIREMENTS:

1. Bilco Engineering has planned to have the detailed civilian information is maintained locally by each office whilst the summary civilian information and sensitive information such as credit card information is maintained at the Atlanta main office. Bilco Engineering has recently stated that each civilian should be registered with and appears in only one local database. Bilco Engineering requires having the users be able to access combined detailed civilian information through a distributed view named Sales.v_All_Civil.

1. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order. Bilco Engineering has recently expressed concern about orders for the current quarter being kept separate from orders for previous quarters.

1. Bilco Engineering has also recently stated that their distribution center inventory is maintained through just-in-time ordering with most inventories coming from local vendors. Bilco Engineering has insured that their purchase order stocking points are determined automatically, and vendor purchase orders are placed from the Atlanta main office. The local branch offices are authorized by Bilco Engineering to order to fill spot shortages and for specialty items that require pick-up from the vendor.

1. Bilco Engineering has also recently stated that their company policy specifies that purchase orders can be kept locally for up to 10 days, but must be kept until the order is either received (and given a status of Closed) or canceled. Bilco Engineering wants to ensure that any order not received within 10 days should be canceled whilst the archival purchase order records are used to update BILCO-SR02 SQL server.

1. Bilco Engineering wants to have the replication administered locally from the archival server whilst requiring the local offices to have control over specific aging requirements for order archives. The Bilco Engineering local offices sometimes require changing the length of time that they hold purchase orders in cases when there is a dispute with a vendor over charges.

1. Bilco Engineering has also recently announced that all periodic reporting requirements

rely on data maintained in the Atlanta office.

1. Bilco Engineering has recently decided to deploy local Web servers at each location and is considered a priority project. Bilco Engineering wants to ensure when this happens that each office will be responsible for posting its own orders. Bilco Engineering additionally wants to have the orders delivered to Atlanta nightly when this happens.

Topic 3, Bilco Engineering (13 Questions)

QUESTION 1

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing performance tuning a database and database application. You will be responsible optimizing the indexing strategies for Bilco Engineering. You are currently in the process of optimizing application performance for applications that reference the Weapons database. You have decided to restart the stop and restart the SQL Server service after hours.

Bilco Engineering has asked you to collect index statistics from selected dynamic management counters after the normal day's operations. You decided to make use of the information to create new nonclustered indexes for the tables accessed by applications. You are required to collect information about missing equality columns and missing include columns whilst selecting which dynamic management views provide the information required.

What should you do? (Choose TWO.)

- A. You should make use of the sys.dm_db_missing_index_details.
- B. You should make use of the sys.dm_db_missing_index_columns.
- C. You should make use of the sys.dm_db_missing_index_group_stats.
- D. You should make use of the sys.dm_db_index_usage_stats.
- E. You should make use of the sys.dm_db_index_operational_stats.
- F. You should make use of the sys.dm_db_missing_index_groups.

Answer: A, B

Explanation: You should remember in the scenario that when we create a new index, the CREATE INDEX statement list equality columns first then inequality columns and that any columns that are listed as included columns should be added as included columns.

1. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers.

Incorrect Answers:

C, F: You should not consider making use of these options in the scenario because

neither of the options used will provide detailed columns information as required for design specific indexes.

D, E: You should not consider making use of these options in the scenario because these options are used to provide information about existing indexes and their current use.

QUESTION 2

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing performance tuning a database and database application.

You will be responsible optimizing the indexing strategies for Bilco Engineering.

Bilco Engineering knows that the Weapon records management application accesses several tables. You have recently discovered that the performance has degraded noticeably when running reports that return information sorted by date.

Each of the tables which contain weapons information has a clustered index based on the table's primary key. Bilco Engineering makes use of the primary key in most of the cases as an identity value generated as rows are added. Bilco Engineering also knows that most of the tables also contain nonclustered indexes based on columns that appear in multiple categories. You are required to determine if the tables are optimized specifically for reporting purposes which are the most resource-intensive operation.

What should you do?

- A. You should make use of SQL Profiler to collect trace information while running reports on a development server containing a copy of the database and make use of the Database Engine Tuning Wizard to recommend the best indexes to use.
- B. You should make use of SQL Profiler to and a Windows counter log to collect a day's activity and use a correlated view of the data collected to determine queries needing the most resources.
- C. You should make use of SQL Profiler to collect a day's worth trace data without reports running and make use of the Database Engine Tuning Wizard to recommend the best indexes to use.
- D. You should make use of SQL Profiler to collect trace information while running reports in addition to normal system activity and make use of the Database Engine Tuning Wizard to recommend the best indexes to use.

Answer: D

Explanation: You should remember in the scenario that taking the action in the answer will ensure that the collected data is respectively as possible when the reports are running by making use of the data collected as the workload source for the Database Engine Tuning Wizard to identify the appropriate index changes to tune for the queries.

1. Bilco Engineering always makes use of the Web Server which generates orders

making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers.

Incorrect Answers:

A: You should not consider making use of this answer in the scenario because the only thing running would be the reports and the servers use different hardware platforms making the representation of the environment inaccurate.

B: You should not consider taking this action in the scenario because even though the answer provides a way to determine the queries requiring the most resources there is no actions taken to what should be done for improving the performance.

C: You should not consider taking the action described in this answer in the scenario because the load is put on resources by the reports and this should be run then the trace is collected.

QUESTION 3

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing performance tuning a database and database application.

You will be responsible optimizing the indexing strategies for Bilco Engineering.

You are in the process of designing the tables for storing weapons information on WeaponTech. Bilco Engineering wants to have the table design to optimize access to the weapons information. The weapons content is stored in the content column.

Bilco Engineering wants to ensure that each content table has a primary key with the VersionNumber as the key column. Bilco Engineering wants the index used for the primary key constraint created as a nonclustered index. You are required to design an indexing strategy for optimizing access to the weapons content whilst ensuring your solution minimizes disk space requirements.

What should you do?

- A. You should add the Content column to the index created for the primary key as an included column.
- B. You should create a PATH secondary XML index on the Content column.
- C. You should create a new clustered index with the Content column as the index key column.
- D. You should create a primary XML index on the Content column.

Answer: A

Explanation: You should remember in the scenario that we should add the Content column to the index created as an included column which enables large object types which includes XML columns and non-key columns in nonclustered indexes.

1. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the

orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order.

Incorrect Answers:

B: You should not consider taking this action in the scenario because indexes are used to help queries that collect information from XML columns and this answer is not use full when retrieving the full XML document stored in the column.

C: You should not consider taking this action in the scenario because you are not able to specify an XML column as the key column for a clustered index.

D: You should not consider taking this action in the scenario because a primary XML index must first be created in order to create a secondary index. XML indexes are used to help queries that collect information from XML columns.

QUESTION 4

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing performance tuning a database and database application.

You will be responsible optimizing and tuning queries for performance for Bilco Engineering. You are currently in the process of looking for ways to improve performance when querying weapons information in the consolidated database.

You plan to replace the ad hoc queries currently used with parameterized queries; the ad hoc query you are modifying is shown below:

```
SELECT c.WepaonID, WepaonName, WeaponType, WeaponStatuys, g.gpa  
FROM Combine.Wepaon c, Combine.Grade g  
WHERE WeaponID = 'A4124' AND WeaponStatus = 'C'
```

You are required to identify situations where parameters can be used in queries. What should you do?

- A. You should replace the literal WeaponID value with an unnamed parameter.
- B. You should replace the literal WeaponID value with unnamed parameters.
- C. You should add an additional column to the SELECT clause as an unnamed parameter.
- D. You should replace the existing column in the SELECT clause with an unnamed parameter.

Answer: A

Explanation: You should remember in the scenario that a parameter query may be used when you require running a query several times with different values in the search argument and there for this is the proper choice in the scenario.

1. Bilco Engineering has configured the Atlanta office to supports the company's operations servers and two distribution warehouse servers. Bilco Engineering makes use of the main office for maintaining information about all weapons and defense systems

information and a master registered civilian list. Bilco Engineering has identified the registered civilians as either visitors or testers-only. The Bilco Engineering military left wing project currently has one database server used as a weapons archival records server at the main office. The SQL Server used as the weapons archival records server is also used as a data source for data transform and load to a business intelligence server.

2. Bilco Engineering has also recently stated that their distribution center inventory is maintained through just-in-time ordering with most inventories coming from local vendors. Bilco Engineering has insured that their purchase order stocking points are determined automatically, and vendor purchase orders are placed from the Atlanta main office.

Incorrect Answers:

B: You should not consider taking this action in the scenario because you should remember that unnamed parameter queries only support one parameter per query.

C, D: You should not consider taking the actions in the options here in the scenario because the parameter queries support the use of parameters in search conditions (WHERE clause) only.

QUESTION 5

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing error-handling routines for Bilco Engineering. You will be responsible for designing code that validates input data and permissions. You are in the process of designing the online registration for online civilians. The online weapons training is currently limited to United States residents only. The civilians are required to provide the information below during the initial registration:

1. Full name - to be entered to the FullName column.
2. E-mail address - to be entered to the Email column.
3. Screen name - to be entered to the UserName column.
4. ZIP code - to be entered to the ZIPCode column.

Bilco Engineering requires each of the civilians to provide a unique e-mail address as the company policy does not allow online civilians to share an e-mail address.

Bilco Engineering makes use of a stored procedure to write the required information to the WebCivilian.NewRegistration table. The additional required information is requested if the civilian chooses to register for any online weapons training.

You have recently received additional instruction that the civilians should be able to change their screen names at any time. You are required to ensure each registration is unique and the screen name should be unique whilst ensuring your solution requires minimal database server and Web server resources.

What should you do?

A. You should set the WebCivilian.NewRegistration table Email column as the table primary key and place a unique constraint on the UserName column.

- B. You should have the Web application query the WebCivilian.NewRegistration table to determine if the E-mail or screen name values already exist.
- C. You should cache the WebCivilian.NewRegistration table locally at the Web server and query the local copy to determine if the E-mail or screen name values already exist.
- D. You should create a primary key on the WebCivilian.NewRegistration table based on Email and UserName as key columns.

Answer: A

Explanation: You should consider taking the actions in this answer in the scenario because taking the actions ensures that if SQL Server will report an error when the stored procedure attempts to enter a duplicate value in either column which can be captured and handled appropriately by the stored procedure.

1. Bilco Engineering has decided that they will make use of replication on the local database servers for provide inventory updates to the Web servers whilst ensuring that any orders generated on the Web servers are submitted to the local database servers using the locally deployed Web services configured at each local network.
2. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order.

Incorrect Answers:

B: You should not consider taking this action in the scenario because you are required to keep resource usage to a minimum and this option will cause the resource usage in the scenario to increase.

C: You should not consider taking the action of this option in the scenario because this action puts more resource usage on the Web server and we are required to keep resource requirements to a minimum.

D: You should not consider taking the action of this option in the scenario because this would only validate the combination of the two values is unique and does not validate each value individually.

QUESTION 6

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing efficient access to a SQL Server service. You will be responsible for designing caching strategies for Bilco Engineering. You are in the process of designing the Web application that online civilians will use to access weapons materials. Bilco Engineering wants to have the application retrieve the weapons from the table in which it is stored and service all further access requirements from the cached data.

Bilco Engineering has the weapons content for individual weapons training courses

stored as an XML document. The solution you design should minimize interaction between the Web server and database server. You are required to identify the best way to meet data access requirements whilst ensuring resource requirements at both Web server and database server are kept to a minimum.

What should you do?

- A. You should make use of a local SQL Server Express instance.
- B. You should make use of an ADO.NET DataReader object.
- C. You should make use of a Transact-SQL cursor.
- D. You should make use of an ADO.NET DataSet object.

Answer: D

Explanation: You should remember in the scenario that making use of the ADO.NET DataSet object that this is used with ADO.NET in disconnected mode and data is read from the database server and written to the DataSet which can be cached locally on the Web server.

1. Bilco Engineering has decided that they will make use of replication on the local database servers for provide inventory updates to the Web servers whilst ensuring that any orders generated on the Web servers are submitted to the local database servers using the locally deployed Web services configured at each local network.
2. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order.

Incorrect Answers:

- A: You should not consider taking this particular action in the scenario because we are required to minimize resource requirements on both the database server and Web server in the scenario and this option places additional resource requirements.
- B: You should not consider making use of this action in the scenario because this solution retrieves information incrementally in connected mode and this is the reason why resource requirements are not minimized.
- C: You should definitely not consider making use of this option in the scenario because this places additional strain on the server's resources and the solution does not meet the caching requirements required in the scenario.

QUESTION 7

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing performance tuning a database and database application.

You will be responsible optimizing the indexing strategies for Bilco Engineering.

You recently planned to deploy a new version of the application used to post and

manage weapons records. The application runs on the client computer and accesses data from a database named accounting on BILCO-SR02.

Bilco Engineering has the application as a commercial program and does not customize the application for individual civilians. The Bilco Engineering network users recently started complaining that the new version of the application does not run well as the previous version. You decide to make use of the Windows System Monitor to view resource use and find that SQL Server:Buffer Manager:Page reads/sec and SQL Server:Buffer Manager:Page writes/sec values have increased significantly.

You are required to improve the application performance whilst remembering that you can not revert to the earlier version because of changes made to database metadata when you upgraded to the new version.

What should you do?

- A. You should increase the amount of memory installed on the database server.
- B. You should use SQL Server Profiler to capture a trace to identify the worst-performing queries.
- C. You should capture a counter log and SQL Server Profiler trace and correlate the results to identify queries running during peak use.
- D. You should increase the amount of memory installed on each of the client computers.
- E. You should make use of SQL Server Profiler to capture trace data while users are running the application and use the Database Engine Tuning Advisor to recommend new indexes.

Answer: E

Explanation: You should remember in the scenario that higher values in SQL Server:Buffer Manager:page reads/sec and SQL Server:Buffer Manager:Page writes/sec indicates that the database is not well tuned causing excessive disk I/O activity which can be corrected taking the actions the answers provide.

1. Bilco Engineering has recently announced that they require having all the company local database servers to automatically report serious errors to the Atlanta main office and immediately paging all IT staff members is imperative if and when errors occur. Bilco Engineering has also expressed concern that the IT staff requires periodic information about selected resource use.

Incorrect Answers:

A, D: You should not consider taking these actions in the scenario because there are no symptoms that indicate there is a problem related to insufficient memory at neither server not client computers in the scenario.

B, C: You should not consider taking these actions in the scenario because these actions even though identifying the queries do not indicate any way to correct the problems.

QUESTION 8

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server

2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing efficient access to a SQL Server service. You will be responsible for designing appropriate data access technologies. You are in the process of designing data access for online weapons training courses. Bilco Engineering requires having data access occurring at the security context of the custom Web application used by online civilians.

Bilco Engineering wants the user's connection time to the database server minimized whilst the load on the database server resources should also be minimized. You are required to select the appropriate data access method ensuring your solution presents minimal security risk.

What should you do?

- A. You should make use of a parameterized query in connected ADO.NET.
- B. You should make use of an ad hoc query to create a Transact-SQL cursor.
- C. You should make use of a stored procedure to create a Transact-SQL cursor.
- D. You should make use of a parameterized query in disconnected ADO.NET.

Answer: D

Explanation: You should remember in the scenario by making use of the parameterized query in disconnected ADO.NET that we are able to help secure data access by preventing attacks such as SQL injection attacks as disconnected ADO.NET minimizes server resource use and connection time.

1. Bilco Engineering has decided that they will make use of replication on the local database servers for provide inventory updates to the Web servers whilst ensuring that any orders generated on the Web servers are submitted to the local database servers using the locally deployed Web services configured at each local network.

2. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order.

Incorrect Answers:

A: You should not consider taking this action in the scenario because connected ADO.NET maintains a connection with the server and does not meet the security requirements of Bilco Engineering.

B: You should not consider making use of an ad hoc query in the scenario because ad hoc queries are especially susceptible to may different types of attacks against a database and is vulnerable to SQL injection attacks.

C: You should not consider taking this action in the scenario because the cursor type in question is a server-side cursor and does not minimize the resource requirements as required in the scenario.

QUESTION 9

You work as the network database administrator at Bilco Engineering. The Bilco

Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing error-handling techniques for Bilco Engineering. You will be responsible for designing code that detects and reacts to errors. Bilco Engineering has a series of stored procedures used at each location to generate tables used for the monthly reports. You are in the process of creating a batch to run the stored procedures.

Bilco Engineering wants to ensure that when errors occur while a stored procedure is running that you are able to identify the failed procedure. You decided to run the stored procedure in context of a TRY clause. You are required to select which of the following could be run if an error should occur.

What should you do?

- A. You should make use of the SELECT ERROR_PROCEDURE().
- B. You should make use of the SELECT ERROR_MESSAGE().
- C. You should make use of the SELECT ERROR_STATE().
- D. You should make use of the SELECT ERROR_NUMBER().

Answer: A

Explanation:

1. Bilco Engineering has decided that they will make use of replication on the local database servers for provide inventory updates to the Web servers whilst ensuring that any orders generated on the Web servers are submitted to the local database servers using the locally deployed Web services configured at each local network.
2. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order.

Incorrect Answers:

B: You should not decide to run the statement in the scenario because this statement would return the textual message associated with the error not the stored procedure that caused the error.

C: You should not consider taking this action in the scenario because this statement would only return the numeric state value for the error.

D: You should not consider taking this action in the scenario because this particular statement would only return the error number of the error.

QUESTION 10

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing a database query strategy for Bilco Engineering. You will be responsible for designing a cursor strategy for Bilco Engineering. Bilco Engineering makes use of a WeaponsOps.CivilianData view which contains information needed for producing civilian summaries. The data is used in the reports and should be as accurate and up-to-date as possible. You are in the process of creating the application to retrieve the data from a Transact-SQL cursor to create the summary. Bilco Engineering wants the cursor used to meet the requirements below:

1. The row orders remain fixed after the cursor is declared.
2. The changes to non-key columns are visible as the application fetches rows from the cursor.
3. The cursor cannot be used to update the source data.
4. The application must be able to retrieve rows randomly.

You are required to declare a cursor that retrieves all the columns from WeaponOPS.CivilianData.

What Transact-SQL statement should you run?

- A. CREATE stu_cursor FAST_FORWARD
FOR SELECT *FROM WeaponOps.CivilianData
- B. CREATE stu_cursor INTENSIVE CURSOR
FOR SELECT *FROM WepaonOps.CivilianData
- C. CREATE stu_cursor SCROLL DYNAMIC READ_ONLY
FOR SELECT *FROM WeaponOps.CivilianData
- D. CREATE stu_cursor CURSOR SCROLL KEYSET READ_ONLY
FOR SELECT *FROM WeaponOps.CivilianData

Answer: D

Explanation: You should remember in the scenario that a key feature of a keyset-driven cursor is that the order and membership does not change and the changes to non-key columns are visible as the rows are fetched which allows us to randomly retrieve information.

1. Bilco Engineering has decided that they will make use of replication on the local database servers for provide inventory updates to the Web servers whilst ensuring that any orders generated on the Web servers are submitted to the local database servers using the locally deployed Web services configured at each local network.
2. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order.

Incorrect Answers:

A: You should not consider making use of this option in the scenario as this cursor is intended for use when required to optimize access to data and plan to retrieve the data using NEXT fetches.

B: You should not consider making use of this particular cursor in the scenario because

this only takes a snapshot of the data and changes to the data will not be seen.

C: You should not consider making use of a dynamic cursor in the scenario because this allows all changes to the data including those to columns to be visible as rows are fetched meaning the row order can change.

QUESTION 11

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing efficient access to a SQL Server service. You will be responsible for designing client libraries to write applications that administer a SQL Server service for Bilco Engineering. You are in the process of designing a new management solution for database servers deployed to civilians. Bilco Engineering company policy does not allow users not authorized to access SQL Server Management Studio.

You decided to write a Microsoft .NET Framework Windows application using SQL Management Objects (SMO). Bilco Engineering requires the application to implement transaction processing. You are required to ensure that the client computers running the management application supports SMO whilst ensuring your solution minimizes resource requirement on the client computers.

What should you do?

- A. You should install the SQL Server Express database engine on each client computer.
- B. You should install SQL Server Client Tools on each of the client computers.
- C. You should install SQL Distributed Management Objects.
- D. You should install SQL Server Native Client on each of the client computers.

Answer: D

Explanation: You should always remember in the scenario that the SQL Server Native Client is a requirement for an SMO application to connect to and manage a SQL Server database server and requires nothing additional for the application to support transaction processing.

1. Bilco Engineering wants to have the replication administered locally from the archival server whilst requiring the local offices to have control over specific aging requirements for order archives. The Bilco Engineering local offices sometimes require changing the length of time that they hold purchase orders in cases when there is a dispute with a vendor over charges.

2. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order. Bilco Engineering has recently expressed concern about orders for the current quarter

being kept separate from orders for previous quarters.

Incorrect Answers:

A: You should not consider taking this action in the scenario since there is no reason for you to install an instance of SQL Server on each of the client computers in the scenario.

B: You should not consider taking this action in the scenario because the SQL Server Client Tools contains the SQL Server Management Studio which company policy states users should not be able to access.

C: You should not consider taking this action in the scenario as this was used in earlier versions of SQL Server and has been replaced by SMO.

QUESTION 12

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing performance tuning a database and database application.

You will be responsible optimizing the indexing strategies for Bilco Engineering.

You are in the process of creating the application which will retrieve selected portions of existing weapons courses to reuse and repurpose in new weapons courses. Bilco Engineering retrieves the content from the existing weapons courses as XML document fragments.

You decided to make use of the XML type query() method to retrieve the XML fragments. You have recently discovered that performance of the queries does not meet company expectations. You are required to improve the performance when querying the column containing the course content (XMLContent).

What should you do? (Choose TWO.)

- A. You should create a nonclustered index on the content table with XMLContent as a nonkey column.
- B. You should create an XML Property secondary index on XMLContent.
- C. You should create an XML Value secondary index on XMLContent.
- D. You should create a nonclustered index on the content table with XMLContent as a key column.
- E. You should create an XML primary index on XMLContent.
- F. You should create an XML Path secondary index on XMLContent.

Answer: E, F

Explanation:

1. Bilco Engineering wants to have the replication administered locally from the archival server whilst requiring the local offices to have control over specific aging requirements for order archives. The Bilco Engineering local offices sometimes require changing the length of time that they hold purchase orders in cases when there is a dispute with a vendor over charges.
2. Bilco Engineering always makes use of the Web Server which generates orders

making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order. Bilco Engineering has recently expressed concern about orders for the current quarter being kept separate from orders for previous quarters.

Incorrect Answers:

A, D: You should not consider taking this action in the scenario because these actions cannot be used as a key column and the performance would only improve if retrieving an XML instance.

B: You should not consider taking this option in the scenario as this type of query would improve performance when retrieving values in this case multiple values in one query, we are retrieving fragments

C: You should not consider taking this option in the scenario as this type of query would improve performance when retrieving element and attribute values and we are retrieving fragments.

QUESTION 13

You work as the network database administrator at Bilco Engineering. The Bilco Engineering network consists of a single Active Directory domain named Bilco Engineering. All servers on the Bilco Engineering network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Bilco Engineering.com network CIO to start designing a database query strategy for Bilco Engineering. You will be responsible for designing a cursor strategy for Bilco Engineering. You are in the process of designing a stored procedure that uses a cursor based on a join between tables on the server to process civilian records. Bilco Engineering wants the cursor to be used within a transaction.

Bilco Engineering wants to have SQL Server to acquire a shared lock when it reads a row in the cursor but the lock should be free after reading the row. You are required to identify the transaction isolation level needed to support the required functionality.

What should you do?

- A. You should make use of the Serializable transaction isolation level.
- B. You should make use of the Read Uncommitted transaction isolation level.
- C. You should make use of the Repeatable Read transaction isolation level.
- D. You should make use of the Read Committed transaction isolation level.

Answer: D

Explanation: You should remember in the scenario that making use of this particular isolation level in the scenario ensures that the lock is released immediately after the row is read in the scenario.

1. Bilco Engineering has recently decided to deploy local Web servers at each location and is considered a priority project. Bilco Engineering wants to ensure when this happens

that each office will be responsible for posting its own orders. Bilco Engineering additionally wants to have the orders delivered to Atlanta nightly when this happens.

2. Bilco Engineering has also recently stated that their company policy specifies that purchase orders can be kept locally for up to 10 days, but must be kept until the order is either received (and given a status of Closed) or canceled. Bilco Engineering wants to ensure that any order not received within 10 days should be canceled whilst the archival purchase order records are used to update BILCO-SR02 SQL server.

3. Bilco Engineering always makes use of the Web Server which generates orders making use of a stored procedure. Bilco Engineering has additionally stated that all the orders be posted to the database server in each office's internal network from the Web servers. Bilco Engineering will also be responsible for sending orders to Atlanta for consolidation whilst making it possible to easily identify the source office for each order. Bilco Engineering has recently expressed concern about orders for the current quarter being kept separate from orders for previous quarters.

Incorrect Answers:

A, C: You should not consider making use of these isolation levels in the scenario because neither of the isolation levels used within a transaction will have the locks released as required in the scenario.

B: You should not consider taking this action in the scenario as you are required to acquire a shared lock and releasing the lock when the row is read, this answer does not request a lock.

Topic 4, Courseware Publishers, Scenario

BACKGROUND:

Company overview

Courseware Publishers is a major publisher of self-study courseware material for students in the information and technology (IT) field. The company has its office in New York. Courseware Publishers currently sells self-study books through third-party bookstores. Courseware Publishers has a public Web site that is used for marketing purposes.

Planned Changes

The company is in the process of redeveloping its Web site to allow registered members to purchase courseware online and to allow users to view sample chapters of its books online. The company also wants to allow authors to access the Books database from a Web application.

Courseware Publishers wants to use SQL Server 2005 to manage the registered membership for the redesigned Web site.

The company also plans to launch a branch office in London, England. The London office will be connected to the New York office by means of a demand-dial router over a dial-up connection.

EXISTING ENVIRONMENT:

Courseware Publishers has a single Windows Server 2003 Active Directory domain named courseware.com. All servers on the Courseware Publishers network run Windows Server 2003 and all client computers run Windows XP Professional. The Courseware Publishers network contains a stand-alone server named CP-SR09 that hosts the intranet Web site for the company and a SQL Server 2005 database server named CP-DB04.

CP-SR09 is located on a perimeter network while CP-DB04 is located on the internal network at the New York office. An external firewall is located between the perimeter network and the Internet and allows only incoming traffic on ports 80 and 443. An internal firewall is located between the perimeter network and the internal network and allows only incoming traffic from CP-SR09 to CP-DB04 on ports 80 and 443.

The company currently has a Web master who converts sample chapters to HTML and publishes them on the Web server. The Web master also uses an application to copy the sample chapter text from Microsoft Word to the Samples table.

BUSINESS REQUIREMENTS:

General Requirements

Courseware Publishers wants Internet users to register with the courseware.com Web site before they are able to view sample chapters. Registered members should be able to subscribe to a weekly e-mail newsletter that informs them of new books, upcoming books related to sample chapters that they have viewed, or books on a particular topic.

Registered members should not receive more than one e-mail a week.

Registered members will access the sample chapters via a secure Web application on the courseware.com Web site. The Web application will display the sample chapter and allow the user to navigate through it by clicking on hyperlinks between sections of the sample chapter. The author of the book will determine which chapter will become the sample chapter, and will define the hyperlinks using Extensible Markup Language (XML) tags.

Registered members must be able to search for sample chapters based on the keywords. They must be able to search for synonyms and inflective forms within the text of each sample chapter. The title, authors or editors, a short abstract, and the average rating of the book from which the sample chapter that matches the search criteria must be displayed in the search results. The short abstract will be stored in the XML documents. The schema for the Samples table is shown in the Samples Table exhibit.

Samples Table Exhibit:

Samples	
PK	SampleID
FK1	AuthorID SampleText
FK2	SubcategoryID Date

The user accounts of registered members will be retained for three months after the user's last visit to the web site. Registered members who do not use the Web site for more than three months will have their user accounts deleted from the system. The billing address, credit card number, expiration date, and social security number of registered members who purchase books through the Web site must be accessible only by Courseware Publishers employees in the Accounting department must only be able to retrieve this information using a stored procedure named getPurchaseDetails.

Registered members must only provide credit card information when they become customers. This requires that they register as customers. Credit card numbers, expiration dates, e-mail addresses and phone numbers must be in a valid format. The proposed schema for the customer and author data is shown in the Member Data exhibit.

Member Data Exhibit:

Members		Customers	
PK	MemberID		
	FirstName	FK	CustomerID
	LastName		CreditCard
	Address		ExpireDate
	Email		BillingAddress
	Phone		

The Web master cannot keep up with the work and a new solution must be developed to make it easier for her to add sample chapters to the Samples database.

Each sample chapter published through the Web application will be stored the Samples.sampleText column that is defined as an xml data type column. The XML schema for the Samples.sampleText column is shown below:

```
<SampleChapter>
<BookTitle>title of book</BookTitle>
<Author>author of book</Author>
<ChapterText>
<ChapterTitle>title of chapter</ ChapterTitle >
<Intro>introduction or synopsis</Intro>
<Topic Title="topic title">topic text</Topic>
<Topic Title="next topic title">next topic text</Topic>
</ChapterText>
</SampleChapter>
```

The Sample.Date column should be updated each time the sample chapter is modified.

TECHNICAL REQUIREMENTS:**Security**

Some personal information about the customers, registered members, and prospective authors are stored in the database. Only Courseware Publishers employees in the Accounting department should be able to access this information and registered members and prospective authors should be able to modify their own personal information. The personal information must be encrypted at all times, including when they are stored and when they are transferred.

Performance

Courseware Publishers wants the courseware.com Web site to offer optimal performance even if a large number of users connect to the Web site.

The search function must be able to return search results to the user within 15 seconds.

These searches will be performed on only the relational data.

Performance for the search function must not be affected by database reporting requirements. The highest performance priorities are:

1. New registrations.
2. Payment processing.
3. Searching for books based on topic.
4. Searching for sample chapters by SampleID.
5. Searching for sample chapters based on concepts and keywords.

Availability

The courseware.com Web site must be available 24 hours a day, seven days a week but downtime of one hour per month is permitted only for scheduled maintenance.

Topic 4, Courseware Publishers (13 Questions)

QUESTION 1

You need to design the index on the Authors table that is very rarely updated. The Authors table is usually used to return Author names based on the AuthorID column. You need to ensure that your solution minimizes overheads for lock management when accessing the Authors table. What should you do?

- A. Set ALLOW_ROW_LOCKS to OFF and set ALLOW_PAGE_LOCKS to ON.
- B. Set ALLOW_ROW_LOCKS to ON and set ALLOW_PAGE_LOCKS to OFF.
- C. Set ALLOW_ROW_LOCKS to ON and set ALLOW_PAGE_LOCKS to ON.
- D. Set ALLOW_ROW_LOCKS to OFF and set ALLOW_PAGE_LOCKS to OFF.

Answer: D

Explanation: Setting both ALLOW_ROW_LOCKS and ALLOW_PAGE_LOCKS to OFF will require a single table lock each time the table is accessed. This will have less lock management overhead than multiple row locks.

Incorrect Answers:

A, B, C: Setting either ALLOW_ROW_LOCKS or ALLOW_PAGE_LOCKS to ON will require a large number of shared locks when a large result set is returned.

QUESTION 2

You need to design the indexes that will optimize performance for search functionality. You need to ensure best performance for a query that return the book title, author, and introduction of a specific sample chapter. The search query is shown below:

```
SELECT SampleText.value("\SampleChapter\BookTitle[1]', 'varchar'),  
SampleText.value("\SampleChapter\Author[1]', 'varchar'),  
SampleText.value("\SampleChapter\ChapterText\ChapterTitle[1]',  
'varchar'),  
SampleText.value("\SampleChapter\ChapterText\Intro[1]', 'varchar')  
FROM Samples  
WHERE SampleID = @sampleID
```

You need to ensure that your solution meets Courseware Publisher's technical and business requirements.

What should you do? (Choose all that apply.)

- A. Create a PROPERTY secondary XML index on the SampleText column.
- B. Create a VALUE secondary XML index on the SampleText column.
- C. Create a nonclustered index on the AuthorID column.
- D. Create a clustered index on the SampleID column.

- E. Create a nonclustered index on the SampleID column.
- F. Create a primary XML index on the SampleText column.

Answer: A, D, F

Explanation: The SampleID column is the primary key column and will be used to locate chapter samples. You should create a clustered index on this column to improve search performance.

Keywords and concepts are stored in SampleText column which must hold XML data as the chapter samples must adhere to an XML schema. You need to use a PROPERTY secondary XML index which uses the SampleID and xml path and xml node values to create an index of the data. However, before you can create a secondary XML index, you must first create a primary XML index.

Incorrect Answers:

B: You need to use a VALUE secondary XML index as users will search for an attribute with a value such as <BookTitle>.

C: No searches are run against the AuthorID column; therefore you don't need to index the AuthorID column.

E: You should create a clustered index on the SampleID column as this provides better performance.

QUESTION 3

You need to design a stored procedure that will return a list of sample chapters from the Samples table. The returned data must include the book title and the sample chapter title but must not return any XML tags. You need to ensure that your solution meets Courseware Publisher's technical and business requirements. What stored procedure you create?

- A. `SELECT SampleText.nodes('\SampleChapter\BookTitle', 'varchar'),
SampleText.nodes('\SampleChapter\ChapterText\ChapterTitle', 'varchar')
FROM Samples`
- B. `SELECT SampleText.query('\SampleChapter\BookTitle'),
SampleText.query('\SampleChapter\ChapterText\ChapterTitle')
FROM Samples`
- C. `SELECT SampleText.value('\SampleChapter\BookTitle[1]', 'varchar'),
SampleText.value('\SampleChapter\ChapterText\ChapterTitle[1]', 'varchar')
FROM Samples`
- D. `SELECT SampleText.query ('child:BookTitle'),
SampleText.query ('child:child:ChapterTitle')
FROM Samples`

Answer: C

Explanation: You must use the value method to query XML data in an xml data type column. The correct syntax for the value method is `column_name.value('XQuery', 'data_type')`. The '[1]' in the XQuery indicates that a

single value is being returned.

Incorrect Answers:

A: The nodes method is used to create relational data from an xml column. It is not used to return scalar values.

B, D: The query method returns the XML tags together with the requested data.

QUESTION 4

You need to redesign the Samples table. You want to implement full-text search to allow authors to search for chapters that have specific information. You need to ensure that the Samples table supports Courseware Publisher's data storage and retrieval requirements.

What Transact-SQL statement should you use?

A. CREATE TABLE Samples

(
SampleID bigint UNIQUE,
AuthorID int REFERENCES Authors.AuthorID,
SampleText varchar(max),
Date smalldatetime NOT NULL
)

B. CREATE TABLE Samples

(
SampleID bigint PRIMARY KEY CLUSTERED,
AuthorID int REFERENCES Authors.AuthorID,
SampleText xml,
Date smalldatetime NOT NULL
)

C. CREATE TABLE Samples

(
SampleID bigint PRIMARY KEY CLUSTERED,
AuthorID int REFERENCES Authors.AuthorID,
SampleText varbinary(max),
FileExtension varchar(8),
Date smalldatetime NOT NULL
)

D. CREATE TABLE Samples

(
SampleID bigint UNIQUE,
AuthorID int REFERENCES Authors.AuthorID,
SampleText image,
FileExtension varchar(8),
Date smalldatetime NOT NULL
)

Answer: C

Explanation: You should create the SampleID column as the primary key so that it can be used as the key column for the full-text search. Because the Web master uses an application to copy the sample chapter text from Microsoft Word to the Samples table, you need to define the SampleText column as either varbinary(max) or image. To enable Full-text indexing on a Microsoft Word document, you also need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

Incorrect Answers:

A: Full-text search requires a key column that contains unique, non-null values. A UNIQUE constraint does not ensure that the column has no null values; it limits null values to one row. Because the Web master uses an application to copy the sample chapter text from Microsoft Word to the Samples table, you need to define the SampleText column as either varbinary(max) or image and not varchar(max). Also, to enable Full-text indexing on a Microsoft Word document, you need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

B: Because the Web master uses an application to copy the sample chapter text from Microsoft Word to the Samples table, you need to define the SampleText column as either varbinary(max) or image and not xml. Also, to enable Full-text indexing on a Microsoft Word document, you need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

D: Full-text search requires a key column that contains unique, non-null values. A UNIQUE constraint does not ensure that the column has no null values; it limits null values to one row.

QUESTION 5

You need to design a stored procedure that will be used by a Web application to retrieve the number of sample chapters in each category and subcategory that have been added in the last 7 days. Your stored procedure must also return the name of the subcategory and the category to which it belongs.

What Transact-SQL statement should you use?

A. `SELECT c.CategoryName, c.SubcategoryName, COUNT(s.SubcategoryID)
FROM Samples AS s INNER JOIN Categories AS c
ON s.SubcategoryID = c.SubcategoryID
WHERE DateDiff("day", s.Date, GETDATE()) <= 7
GROUP BY c.CategoryName, c.SubcategoryName
ORDER BY c.CategoryName, c.SubcategoryName`

B. `SELECT c.CategoryName, c.SubcategoryName, COUNT(s.SubcategoryID)
FROM Samples AS s INNER JOIN Categories AS c
ON s.SubcategoryID = c.SubcategoryID
WHERE DateDiff("week", s.Date, GETDATE()) <= 1
GROUP BY c.SubcategoryName
ORDER BY c.SubcategoryName, c.CategoryName`

C. `SELECT c.CategoryName, c.SubcategoryName, COUNT(s.SubcategoryID)
FROM Samples AS s INNER JOIN Categories AS c
ON s.SubcategoryID = c.SubcategoryID`

```
WHERE DateDiff("week", s.Date, GETDATE()) <= 1
GROUP BY c.CategoryName, c.SubcategoryName
ORDER BY c.CategoryName, c.SubcategoryName
D. SELECT c.CategoryName, c.SubcategoryName, COUNT(s.SubcategoryID)
FROM Samples AS s INNER JOIN Categories AS c
ON s.SubcategoryID = c.SubcategoryID
WHERE DateDiff("day", s.Date, GETDATE()) <= 7
ORDER BY c.CategoryName, c.SubcategoryName
```

Answer: A

Explanation: The COUNT aggregate function will return the number of non-null items within the result set based on the GROUP BY clause. To only retrieve the sample chapters that have been added in the last 7 days, you should use the DateDiff function to limit the result set to 7 days.

Incorrect Answers:

B: The DateDiff function will return sample chapters that have been added either this week or last week rather than the last 7 days. Also, when the COUNT aggregate function is used, the columns referenced in the SELECT statement must appear in the GROUP BY clause or a syntax error will occur.

C: The DateDiff function will return sample chapters that have been added either this week or last week rather than the last 7 days.

D: When the COUNT aggregate function is used, the columns referenced in the SELECT statement must appear in the GROUP BY clause or a syntax error will occur.

QUESTION 6

You need to design a strategy for making data available at the new office in London. A database server named CW-DB05 is installed in the London office. You need to ensure that your solution meets Courseware Publisher's technical and business requirements.

What should you do?

- A. Configure federated database servers.
- B. Configure merge replication.
- C. Configure snapshot replication.
- D. Configure log shipping.

Answer: B

Explanation: Merge replication can be used to merge the data from two autonomous sites.

Incorrect Answers:

A: Federated database servers are not used to distribute data between locations. You can use federated database servers to partition the data processing load across multiple well-connected database servers.

C: Snapshot replication is best used in small databases that do not get updated on a

regular basis.

D: Log shipping is used for fault tolerance. It is not used to distribute data between locations.

QUESTION 7

You need to design a stored procedure that will be used by a Web application to retrieve the topics for a single sample chapter. You need to ensure that your solution meets Courseware Publisher's technical requirements.

What should you do?

- A. Use the SampleText.openXML method.
- B. Use the SampleText.value method.
- C. Use the SampleText.nodes method.
- D. Use the SampleText.query method.

Answer: C

Explanation: The nodes method is used to create relational data from an xml column. It is used to parse a series of XML elements into rowsets.

Incorrect Answers:

- A: OPENXML is used with the SELECT statement to extract relational data from an XML document. It cannot be used as a method of SampleText.
 - B: The value method returns a scalar value rather than a relational result set.
 - D: The query method returns the XML element and its child elements. It does not return a relational result set.
-

QUESTION 8

You need to design the functionality that allows authors to update sample chapters that are published on the Web site. You need to ensure that your solution ensures maximum concurrency. You also need to ensure that your solution meets Courseware Publisher's technical and business requirements.

What should you do?

- A. Use a SERIALIZABLE transaction isolation level.
- B. Use a READ COMMITTED SNAPSHOT transaction isolation level.
- C. Use the TABLOCK query hint in the UPDATE statement.
- D. Use a READ UNCOMMITTED transaction isolation level.
- E. Use a REPEATABLE READ transaction isolation level.

Answer: B

Explanation: You should use the snapshot isolation level. Snapshot isolation prevents both uncommitted dependencies (dirty reads) and nonrepeatable reads. It uses row versioning by reading data into tempdb and accessing data from there. It is also a lower isolation level than Serializable and therefore less likely to result in blocking locks and deadlock conditions.

Incorrect Answers:

A: Serializable isolation prevents both uncommitted dependencies (dirty reads) and nonrepeatable reads but it serializable is also the highest isolation level and is the most likely to result in deadlocks.

C: A TABLOCK query hint will place a lock on the entire table. There is not need to lock the entire table when an author updates a single row.

D: Read uncommitted isolation is the least likely to cause deadlock conditions, but prevents neither uncommitted dependencies (dirty reads) nor nonrepeatable reads. This may result in data inconsistencies.

E: Repeatable read isolation will take shared locks; therefore the transaction will be blocked if the Investments table is locked by another session.

QUESTION 9

You need to design the security for the Samples table. You need to ensure that the principle of least privilege is implemented when you assign permissions. You create two database roles named Authors and Members.

What should you do? (Choose all that apply.)

A. Assign the SELECT and UPDATE permissions to the Authors database role.

B. Assign the SELECT and INSERT permissions to the Authors database role.

C. Assign the SELECT permission to the Members database role.

D. Assign the UPDATE permission to the Members database role.

E. Assign the INSERT permission to the Members database role.

Answer: B, C

Explanation: The authors need access to the sample chapter. Therefore they need SELECT permissions. They must also be able to apply hyperlinks to the sample chapters. Therefore they need INSERT permissions.

Incorrect Answers:

A: Authors do not need UPDATE permissions as updated documents must be treated as new versions.

D, E: Registered members need only SELECT permissions to view the sample chapters. They are not required to insert keywords and will not update the documents.

QUESTION 10

Six month after the new database system was deployed; users at Courseware Publishers complain that the database intermittently performs poorly. You need to identify which user sessions are running a large number of processes or holding a large number of open locks. You want to collect the data in real-time.

What should you do?

A. Use the SQL Server Management Studio Activity Monitor.

B. User Windows Performance with the SQL Server General Statistics counter.

C. Use the Database Engine Tuning Advisor.

D. User Windows Performance with the SQL Server Locks counter.

Answer: A

Explanation: The SQL Server Management Studio Activity Monitor allows you to view locks and process information by connection and in real-time.

Incorrect Answers:

B: The SQL Server General Statistics counter can be used by Windows Performance to view general information but not information about locks and processes.

C: The Database Engine Tuning Advisor is used to obtain index recommendations. It does not provide information on locks and processes.

D: The SQL Server Locks counter can be used by Windows Performance to view information about locks but not processes. The SQL Server Locks counter also does not provide information by connection not does it provide information in real-time.

QUESTION 11

You need to design a system that will allow authors to submit new course material. The submitted material must be placed into a folder that can be accessed by Courseware Publishers editors. You need to ensure that your solution meets Courseware Publisher's technical and business requirements. What should you do? (Choose all that apply.)

- A. Create an event-driven subscription.
- B. Configure a Server Broker queue.
- C. Use the File protocol as the delivery channel.
- D. Configure Event Notification.
- E. Create a custom event provider.

Answer: A, C

Explanation: You need to use Notification Services to accomplish this task. To use Notification Services, you must create an event-driven subscription and use the File protocol as the delivery channel.

Incorrect Answers:

B: Service Broker is used to create applications that execute synchronously.

D: Event Notification is used to respond to data definition language (DDL) events such as creating tables. It cannot be used to respond to data manipulation language (DML) events, such as data insertions.

E: You do not need to create a custom event provider.

QUESTION 12

Six month after the online investment system had been deployed; Web site users complain about poor response times from the Web site. You need to determine if the problem is caused by deadlocks. If the problem is caused by deadlocks, you want to determine which objects are involved in the deadlocks.

What should you do?

- A. Use SQL Server Profiler to run a trace.
- B. Run the sp_lock stored procedure.
- C. Query the sys.dm_tran_database_transactions dynamic management view.
- D. Query the sys.dm_tranlocks dynamic management view.

Answer: A

Explanation: You can use SQL Server Profiler to monitor locks and deadlocks.

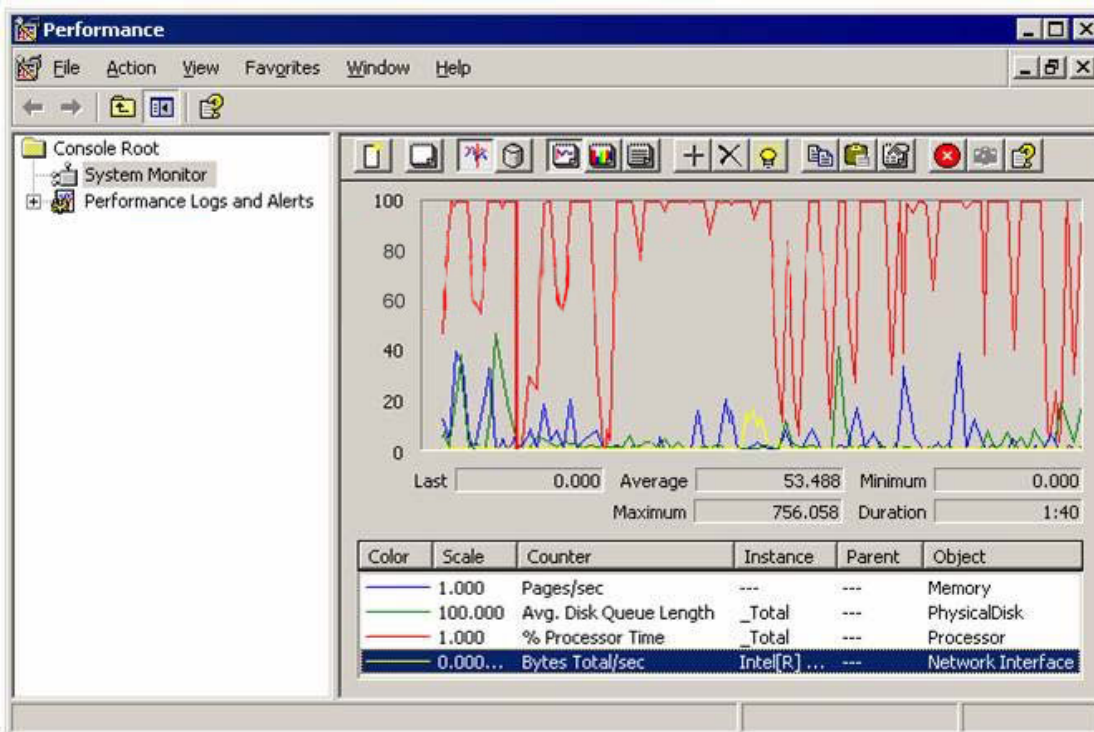
Incorrect Answers:

B, D: The sp_lock stored procedure has been deprecated in SQL Server 2005 in favor of the sys.dm_tranlocks dynamic management view. However, the sys.dm_tranlocks dynamic management view allows you to view current locks and deadlocks. It does not allow you to view a history of deadlocks

C: The sys.dm_tran_database_transactions dynamic management view provides information about transactions rather than locks and deadlocks.

QUESTION 13

Six month after the new database system was deployed; users at Courseware Publishers complain that the performance of the database has deteriorated. You use System Monitor to monitor the performance of Certkiller -DB01 and receive the output as shown in the exhibit.



You also notice that the ratio of SQL Recompilations/sec to Batch Requests/sec is unusually high. You suspect that the number of recompiles is causing the high processor time value. You need to improve the performance of the database. What should you do?

- A. Turn off automatic updates of statistics for all tables in the database.
- B. Install an additional processor on the database server.
- C. Run the Database Engine Tuning Advisor.
- D. Use SQL Server Profiler to identify the stored procedures being recompiled.

Answer: D

Explanation: You can use SQL Server Profiler to identify the stored procedures that are being recompiled. SQL Profiler will indicate which stored procedures are being recompiled and why each recompilation is occurring.

Incorrect Answers:

A: Turning off automatic updates of statistics for the tables may improve performance but it does not address the problem caused by the recompiles.

B: The high processor usage could be caused by the high number of recompiles. You should first reduce the number of recompiles before installing extra hardware on the server.

C: You can use the Database Engine Tuning Advisor to determine whether indexing and partitioning of a table would improve database performance but you cannot use it to identify which stored procedures are being recompiled.

Topic 5, Mondo Transport Inc, Scenario

BACKGROUND:

Company Overview

Mondo Transport is the worlds leading specialists who specialize in vehicle, motorcycle and aviation forced induction systems as well as prototype induction systems and accessories. Mondo Transport currently consists of a main office in Houston and four branch offices located in Seattle, Atlanta, Phoenix and Dallas.

Mondo Transport has recently decided to have all the records kept at each branch office. Mondo Transport has additionally started sending a hard copy of the data about the forced induction systems, prototype induction systems and accessories to the main office in Houston. Mondo Transport has recently revealed their plans for the business operations by requiring the operations to be computerized. Mondo Transport in addition has also acquired a Microsoft SQL Server 2005 computer running Microsoft Windows Server 2003 and your responsibility is to design the database and the application that will be used to enter and retrieve data.

Infrastructure

Mondo Transport currently makes use of their main office network which consists of an Active Directory domain which is configured to make use of Internet Security and Acceleration (ISA) Server for the purpose of sharing the Internet connection. Mondo Transport has recently informed you that the current network has no virtual private network (VPN) endpoint configured and has no plans to add one.

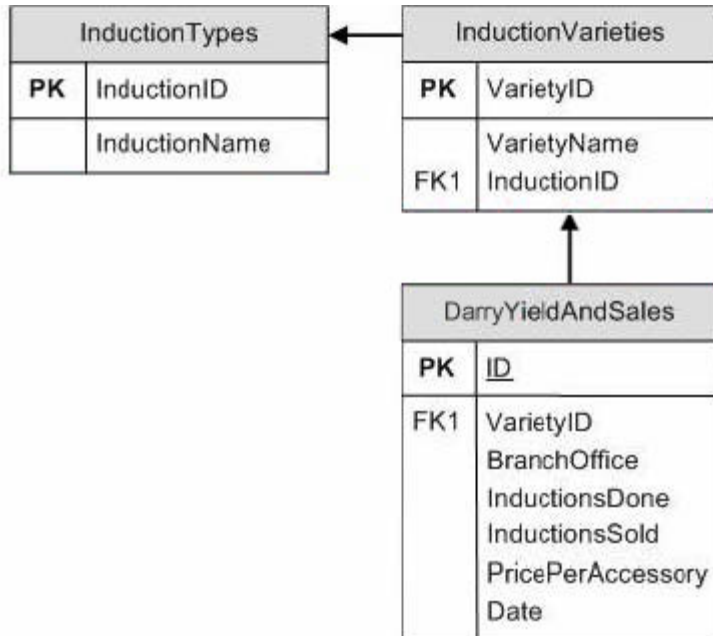
The Mondo Transport branch offices currently only have between one and three computers which are connected to the Internet via the ISA server and are primarily used to place orders and perform required research.

Existing Environment

Mondo Transport has also stated that the network branch office located in Houston and Phoenix are currently used to track yields and sales in a Microsoft Excel spreadsheet whilst the other network branch office were instructed to make use of a paper-based tracking system.

BUSINESS REQUIREMENTS:

Mondo Transport has recently decided that the network should consider it imperative that we are able to track daily yield for each branch office by Induction variety, track daily sales for each branch office. Mondo Transport has recently decided to make use of the proposed schema for tracking this information shown in the Yield and Sales Tracking exhibit below:



1. Mondo Transport has recently decided to authorize the branch offices the capability to set their own prices based on regional differences as the prices can change on a daily basis.

1. Mondo Transport has recently decided to have all the branch offices donate their excess induction accessories at the end of each week. Mondo Transport has recently decided to ensure that one-half of the prototype forced induction accessories should be donated to the ultimate high performance vehicle, motorcycle and aviation industry racing teams to evaluate the forced induction prototypes.

1. Mondo Transport has recently informed you that they require the ability to retrieve the amount of accessories produced in the past year.

1. Mondo Transport additionally requires the induction developers in the Houston main office to be able to analyze the data about the daily yield in each location and the impact of wear and tear data.

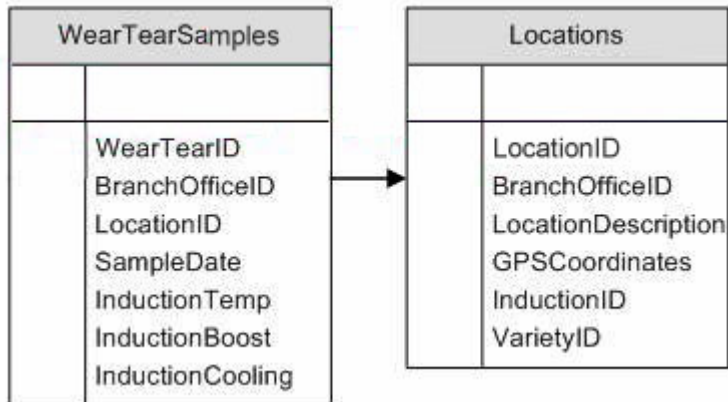
1. Mondo Transport will have the induction developers typically analyzing the data from a single branch office whilst taking samples from approximately 20 locations at each branch office.

TECHNICAL REQUIREMENTS:

Performance

Mondo Transport is currently aware that the WearTear table will be expected to become

quite large. The proposed schema of the WearTear table and proposed schema of the Locations table are shown in the WearTear exhibit below:



Availability

Mondo Transport has recently informed you that the data should be sent from each location at the end of the day.

Security

Mondo Transport has recently announced that it considers the data that is sent from each location as sensitive and it is imperative the data should be encrypted.

Topic 5, Mondo Transport (10 Questions)

QUESTION 1

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start designing a database query strategy. You are currently busy writing and modifying queries for Mondo Transport. You have recently received additional instruction to develop a stored procedure which will be used by the induction developers.

You are required to create the stored procedure by selecting the proper code to use that will have the capability of accepting a range of dates and returns the number of forced induction systems produced and forced induction prototype accessories at each of the branch offices during that period of time.

What should you do?

- A. SELECT f.Induction, v.InductionName, BranchOffice, SUM(ys.InductionsProduced)
FROM BranchOffice.DailyYeildAndSales ys
JOIN InductionVarieties v
ON v.VarietyID = ys.VarietyID
JOIN Induction f
ON f.InductionID = v.InductionID
WHERE ys.[Date] >= @startDate AND ys.[Date] <= @endDate
GROUP BY v.VarietyName
- B. SELECT f.Induction, v.VarietyName, BranchOffice, SUM(ys.InductionsProduced)


```
FROM BranchOffice.DailyYieldAndSales ys
FULL JOIN InductionVarieties v, Induction f
ON v.VarietyID = ys.VarietyID AND f.InductionID = v.InductionID
WHERE ys.[Date] >= @startDate AND ys.[Date] <= @endDate
GROUP BY BranchOffice, f.Induction, v.VarietyName
C. SELECT f.Induction, v.VarietyName, BranchOffice, SUM(ys.InductionsProduced)
FROM BranchOffice.DailyYieldAndSales ys
INNER JOIN InductionVarieties v
ON v.VarietyID = ys.VarietyID
INNER JOIN Induction f
ON f.InductionID = v.InductionID
WHERE ys.[Date] BETWEEN @startDate AND @endDate
GROUP BY f.Induction, v.VarietyName
D. SELECT f.Induction, v.VarietyName, BranchOffice SUM(ys.InductionsProduced)
FROM BranchOffice.DailyYieldAndSales ys
LEFT JOIN InductionVarieties v
ON v.VarietyID = ys.VarietyID
LEFT JOIN Induction f
ON f.InductionID = v.InductionID
WHERE ys.[Date] BETWEEN @startDate AND @endDate
GROUP BY BranchOffice, f.Induction, v.VarietyName
```

Answer: D

Explanation: You should always remember in the scenario that making use of the code in the answer that you are capable of combining multiple JOINS into a single SELECT statement because a LEFT join will cause all the rows from the left table matching rows from the right table to be returned.

1. Mondo Transport has recently decided that the network should consider it imperative that we are able to track daily yield for each branch office by Induction variety, track daily sales for each branch office.
2. Mondo Transport will have the induction developers typically analyzing the data from a single branch office whilst taking samples from approximately 20 locations at each branch office.
3. Mondo Transport additionally requires the induction developers in the Houston main office to be able to analyze the data about the daily yield in each location and the impact of wear and tear data.

Incorrect Answers:

A: You should not consider using this code as your query in the scenario because when you do not specify INNER, LEFT, RIGHT or FULL join SQL Server will automatically assume that it is an INNER join and the BranchOffice and Induction columns are not included in the GROUP BY clause or in aggregate.

B: You should not consider making use of the code for your query in the scenario because in order to execute a statement that has multiple joins you should join them one table at a time and you should not make use of a FULL JOIN as it returns all rows from all the tables.

C: You should not consider making use of this code as your query in the scenario because an INNER join only returns rows that match the join condition rendering the induction systems that were not yet used and will not be listed in addition the BranchOffice column is included in the SELECT list but not in the GROUP BY clause.

QUESTION 2

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start designing and performance tuning a database and database application. You have recently started with optimizing indexing strategies. You are required to create a unique index on the Locations table which has been partitioned on the BranchOfficeID column to meet query optimization requirements.

Mondo Transport is aware that there is not any index defined on the table currently. You are required to ensure that the index allows you to add a new partition if a Branch office is added or subdivided.

What should you do?

- A. You should create a nonclustered index using the LocationID column as the key column and include the BranchOfficeID column.
- B. You should create a nonclustered index using the BranchOfficeID column as the key column and include the LocationID column.
- C. You should create a clustered index using the using the LocationID column as the key column.
- D. You should create a clustered index using the LocationID and BranchOfficeID as the key column.

Answer: D

Explanation: You should remember in the scenario in order to create a unique index that the unique index must include the partitioning column as a key column, as for optimal queries that limit the result set by Branch office that is why we should partition the table on the BranchOfficeID column.

1. Mondo Transport has recently decided that the network should consider it imperative that we are able to track daily yield for each branch office by Induction variety, track daily sales for each branch office.
2. Mondo Transport will have the induction developers typically analyzing the data from a single branch office whilst taking samples from approximately 20 locations at each branch office.
3. Mondo Transport additionally requires the induction developers in the Houston main office to be able to analyze the data about the daily yield in each location and the impact of wear and tear data.

Incorrect Answers:

A: You should remember in the scenario in order to create a unique index that the unique index must include the partitioning column as a key column and this solution does only include the column and there for should not be used.

B: You should remember in the scenario in order to create a unique index that the unique index must include the partitioning column as a key column and this solution suggest using the BranchOfficeID which does not have any unique values.

C: You should not consider taking this action in the scenario because in order for the index to be aligned with the partitions the partitioning column should be added as a key column.

QUESTION 3

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start designing efficient access to a SQL Server Service. You are in the process of designing appropriate data access technologies. You have recently received additional instruction to create the application which will allow the branch offices to retrieve the weekly donation list. You are required to design a data access method for the data.

What should you do?

- A. You should create a stored procedure and expose it as a Web method.
- B. You should create an ad hoc query and issue it from the ADO.NET connection.
- C. You should create a table-valued function and expose it as a Web method.
- D. You should create a table-valued function and call it through an ADO.NET connection.

Answer: A

Explanation: You should remember in the scenario that by making use of the stored procedure and exposing it as a Web method you can create an HTTP endpoint on the Secure Sockets Layer (SSL)

port to allow encrypted access over the Internet and the firewall will have to be configured to forward HTTPS requests to the SQL Server.

1. Mondo Transport has recently decided to have all the branch offices donate their excess induction accessories at the end of each week.
2. Mondo Transport has recently announced that it considers the data that is sent from each location as sensitive and it is imperative the data should be encrypted.

Incorrect Answers:

B, D: You should not consider making use of an ADO.NET connection in the scenario because this particular implementation requires changes to be made to the network infrastructure to be done securely by adding and configuring a Virtual Private Network (VPN).

C: You should remember in that scenario that you are only capable of exposing

scalar functions and stored procedures as Web methods so making use of this answer is incorrect.

QUESTION 4

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start performance tuning a database and database application. You will be responsible for optimizing and tuning queries for performance. You have recently received additional information concerning one of the induction developers who knows a little of SQL Server frequently runs ad hoc queries against the database. This has sometimes caused negative performance impacts on the other running queries. You are required to reduce the impact of the induction developer's queries.

What should you do?

- A. You should execute the `sp_configure 'query governor cost limit', 5`
- B. You should execute the `sp_configure 'query governor cost limit', 0`.
- C. You should execute the `SET STATISTICS TIME = ON`.
- D. You should `SET QUERY_GOVERNOR_COST_LIMIT = 1`.
- E. You should execute the `SET STATISTICS TIME = OFF`.

Answer: A

Explanation: You should make the configurations suggested in the answer in the scenario because following these instructions you would cause the query governor to terminate a query if its estimated execution time exceeds 5 seconds.

1. Mondo Transport has recently revealed their plans for the business operations by requiring the operations to be computerized.

Incorrect Answers:

B: You should not consider making this configuration in the scenario because this particular configuration will turn the query governor off and long running queries will not be stopped.

C, D: You should not consider taking this action in the scenario because setting the `STATISTICS TIME = ON` causes SQL Server to report the amount of time it takes to parse, compile and execute the query and is the default setting in SQL Server additionally setting `STATISTICS TIME = OFF` is not a good idea either.

QUESTION 5

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start

performance tuning a database and database application. You will be responsible for optimizing and tuning queries for performance. You have recently started the creation of a stored procedure which will be used to accept a data range and return the Average total sales at each branch office during that date range.

The code you wrote for the creation of the stored procedure is shown below:

```
SELECT BranchOffice, VarietyID, AVG (InductionsSold *  
PricePerAccessory)
```

```
FROM BranchOffice.DailyYieldAndSales
```

```
WHERE [Date] BETWEEN @startDate AND @endDate
```

```
GROUP BY BranchOffice, VarietyID
```

You have recently received additional information about the tables cluster index on the ID column. The Mondo Transport network users have recently started complaining that the query runs very slow. You are required to make the required configuration changes to optimize the query.

What should you do?

A. You should add a computed column named DailySales to the BranchOffice.DailyYieldAndSales table.

You should then create a non-clustered index with BranchOffice, VarietyID, Date and DailySales as key values.

You should finally modify the query to reference the computed column.

B. You should add a computed column named DailySales to the BranchOffice.DailyYieldAndSales table.

You should then add the InductionsSold and PricePerAccessory columns to the clustered index.

You should finally modify the query to reference the computed column.

C. You should create a table named BranchOffice.

You should then create a foreign key constraint on the BranchOffice column.

You should finally modify the query to use an inner join.

D. You should create a table named TotalSales which includes BranchOffice, VarietyID, Date and Sales columns.

You should then create an INSERT trigger on the DailyYieldAndSales table to update the TotalSales table.

You should finally modify the query to use an inner join.

Answer: A

Explanation: You should remember in the scenario that by adding a computed column named DailySales to the BranchOffice.DailyYieldAndSales table and the covering index and since the table is updated once daily you could persist computed column to prevent calculations from being made during the query.

1. Mondo Transport has recently decided that the network should consider it imperative that we are able to track daily yield for each branch office by Induction variety, track daily sales for each branch office.

2. Mondo Transport has recently revealed their plans for the business operations by requiring the operations to be computerized.

Incorrect Answers:

B: You should not consider taking the action in the scenario because you are required to improve the performance of the query and this option would cause the query optimizer not to choose the clustered index as the ID column is not used in the query.

C: You should not consider taking the action in the scenario because you are required to improve the performance of the query and this option would only degrade the performance further.

D: You should not consider taking the action in the scenario because you are required to improve the performance of the query and by adding a TotalSales column which is populated with a trigger would only degrade the performance further.

QUESTION 6

DRAG DROP

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start performance tuning a database and database application. You will be responsible for the scaling of database application of Mondo Transport. You have recently started creating the tables that the induction developers require. The induction developers will perform frequent joins between the Locations and WearTearSamples tables.

You are required to design a partitioning scheme that would optimize performance. What should you do? (To answer, drag the appropriate column from the pane on the left to the correct location in the pane on the right.)

Partitioning columns, select from these	Tables, place here
LocationID	Location Place here.
WearTearID	Place here, if any.
BranchOfficeID	WearTearSamples Place here.
SampleDate	Place here, if any.

Answer:



Explanation:

You should remember in the scenario that we are required to optimize the performance of the query, and by taking the actions considered in the answer you would effectively ensure that the induction developers will be making use of an optimized query that joins the two tables as partitions can be joined.

1. Mondo Transport additionally requires the induction developers in the Houston main office to be able to analyze the data about the daily yield in each location and the impact of wear and tear data
2. Mondo Transport will have the induction developers typically analyzing the data from a single branch office whilst taking samples from approximately 20 locations at each branch office

Incorrect Answers:

You should not consider making the partition on the LocationID column in the scenario because the branch offices approximately contain about 20 locations. You should additionally not consider making use of the WearTearSamples table partitioned on the WearTearID columns as the queries will not be performed based on the WearTearID.

QUESTION 7

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start performance tuning a database and database application. You will be responsible for optimizing indexing strategies. You are currently busy designing the functionality required to report the amount of forced induction systems donated by each branch office at the end of the week.

You just finished creating the query design which is shown below:

```
SELECT BranchOffice, f.Induction, v.VarietyName,
(SUM (ys.InductionsProduced-ys.InductionsSold) ) /2 As Excess
FROM BranchOffice.DailyYieldAndSales ys
INNER JOIN InductionVarieties v ON v.VarietyID = ys.VarietyID
INNER JOIN Induction f ON f.InductionID = v.InductionID
WHERE ys. [Date] BETWEEN DATEDIFF (days, GETDATE (), 7) AND GETDATE ()
GROUP BY BranchOffice, f.Induction, v.VarietyName
```

You are required to optimize the performance for the query by creating an indexed view whilst ensuring the optimization also optimizes the query that determines the amount of donated forced induction systems.

What code should you use?

A. CREATE VIEW vDonations WITH SCHEMABINDING
AS
SELECT BranchOffice, f.Induction AS Induction, v.VarietyName AS
Variety, count_big(*),
SUM(ys.InductionsProduced -ys.InductionsSold) As Excess
FROM BranchOffice.DailyYieldAndSales ys
INNER JOIN InductionVarieties v
ON v.VarietyID = ys.VarietyID
INNER JOIN Induction f
ON f.InductionID = v.InductionID
GROUP BY BranchOffice, f.Induction, v.VarietyName
CREATE UNIQUE CLUSTERED INDEX ix_Donated ON vDonations(BranchOffice,
Induction, Variety)

B. CREATE VIEW vDonations
AS
SELECT BranchOffice, f.Induction AS Induction, v.VarietyName AS Variety,
ys.InductionsProduced -ys.InductionsSold As Excess
FROM BranchOffice.DailyYieldAndSales ys
INNER JOIN InductionVarieties v
ON v.VarietyID = ys.VarietyID
INNER JOIN Induction f
ON f.InductionID = v.InductionID
GROUP BY BranchOffice, f.Induction, v.VarietyName
CREATE UNIQUE CLUSTERED INDEX ix_Donated ON vDonations(BranchOffice,
Induction, Variety)

C. CREATE VIEW vDonations
AS
SELECT BranchOffice, f.Induction AS Induction, v.VarietyName,
ys.InductionsProduced -ys InductionsSold As Excess
FROM BranchOffice.DailyYieldAndSales ys
INNER JOIN InductionVarieties v
ON v.VarietyID = ys.VarietyID
INNER JOIN Induction d
ON f.InductionID = v.InductionID
GROUP BY BranchOffice, f.Induction, v.VarietyName
CREATE UNIQUE INDEX ix_Donated ON vDonations(Date, BranchOffice, Induction,
Variety)

D. CREATE VIEW vDonations WITH SCHEMABINDING
AS
SELECT BranchOffice, f. Induction, v.VarietyName AS Variety, count_big(*),
(SUM(ys.InductionsProduced -ys.InductionsSold))/2 As Excess


```
FROM BranchOffice.DailyYieldAndSales ys
INNER JOIN InductionVarieties v
ON v.VarietyID = ys.VarietyID
INNER JOIN Induction f
ON f.InductionID = v.InductionID
GROUP BY BranchOffice, f.Induction, v.VarietyName
CREATE CLUSTERED INDEX ix_Donated ON vDonations(Induction, Variety)
```

Answer: A

Explanation: You should remember in the scenario that you should create the view with the WITH SCHEMABINDING when planning a index and additionally the SELECT statement includes an aggregate and requires containing the count_big(*) function. The first index created should be a unique clustered index since the view creates aggregated tables and the BranchOffice, Induction and Variety columns provide the unique values required.

1. Mondo Transport has recently decided to have all the branch offices donate their excess induction accessories at the end of each week. Mondo Transport has recently decided to ensure that one-half of the prototype forced induction accessories should be donated to the ultimate high performance vehicle, motorcycle and aviation industry racing teams to evaluate the forced induction prototypes

Incorrect Answers:

B: You should not consider making use of this particular code in the scenario because the view should be created with the WITH SCHEMABINDING and additionally the view will not provide better performance than the view that has already aggregated the data.

C: You should not consider making use of this particular code in the scenario because the view should be created with the WITH SCHEMABINDING and additionally the data is not aggregated in the view, the BranchOffice, Induction and Variety columns provide unique values required.

D: You should not consider using this code to create the view in the scenario because you are not able to use expressions on an aggregated value and additionally the first index created should be a unique clustered index.

QUESTION 8

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start Designing error-handling routines for applications. The solution your design will use user-defined messages to communicate application events. You have recently finished coding a stored procedure that will be used by the induction developers allowing them to retrieve the weekly donation list.

The stored procedure you coded will be exposed as a method from a native XML Web service and makes use of a TRY CATCH block to capture errors should they occur during execution. You have recently received additional instruction requiring

you to document any errors at the client using the client's Application log. You are required to capture the error number as well as error message whilst ensuring that you solution requires minimal effort.

What should you do? (Choose TWO.)

- A. You should return the `ERROR_NUMBER()` and `ERROR_MESSAGE()` values from the TRY block as output parameters and return NULL values if no errors occur.
- B. You should return the value from `@@ERROR` when any errors occur as output parameters or a NULL value if no errors occur. You should then create a lookup table on the client which associates error numbers with SQL Server error messages.
- C. You should code the application to generate an exception when it receives an error from the Web service.
- D. You should code the application to write the result to the Application log.
- E. You should return the `ERROR_NUMBER()` and `ERROR_MESSAGE()` values from the CATCH block as output parameters and return NULL values if no errors occur.

Answer: D, E

Explanation: You should remember in the scenario that we are required to retrieve error information which includes returning the `ERROR_NUMBER()` and `ERROR_MESSAGE()` values from the CATCH block as output parameters and returning NULL values if no errors occur. We can then successfully have the client application write the required information as an entry into the Windows Application log.

1. Mondo Transport has recently announced that it considers the data that is sent from each location as sensitive and it is imperative the data should be encrypted.
2. Mondo Transport has recently revealed their plans for the business operations by requiring the operations to be computerized.
3. Mondo Transport has recently decided to have all the branch offices donate their excess induction accessories at the end of each week.

Incorrect Answers:

- A: You should not consider making use of the TRY block in the scenario because the required information is only available in the CATCH block and would generate an error used in the TRY block.
- B: You should not consider making use of the `@@ERROR` when the error occurs as an output parameter in the scenario because the `@@ERROR` is used with legacy error handling in SQL Server and not with the TRY CATCH structures.
- C: You should not consider coding the application to raise an exception in the scenario because we are required to use minimal effort and this modification requires adding additional exception handling in the client application.

QUESTION 9

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start performance tuning a database and database application. You will be responsible for optimizing indexing strategies. You have recently started optimizing a database to support a redesigned query used in the induction developer's application. You are required to add a key column to a clustered index whilst ensuring that the base table remains available during the required changes.

What should you do?

- A. You should execute ALTER INDEX REBUILD with ONLINE = ON.
- B. You should execute ALTER INDEX REORGANIZE with ONLINE = ON.
- C. You should disable the index and execute CREATE INDEX with DROP EXISTING = ON and ONLINE = ON.
- D. You should execute CREATE INDEX with DROP EXISTING = ON and ONLINE = ON.

Answer: D

Explanation: You should remember in the scenario that when you execute the CREATE INDEX with DROP EXISTING = ON and ONLINE = ON statement the user will still have the ability to access the table while the index is being built.

1. Mondo Transport additionally requires the induction developers in the Houston main office to be able to analyze the data about the daily yield in each location and the impact of wear and tear data.
2. Mondo Transport will have the induction developers typically analyzing the data from a single branch office whilst taking samples from approximately 20 locations at each branch office.

Incorrect Answers:

- A, B: You should not consider making use of the ALTER INDEX statement in the scenario because you are not able to add a column by making use of the ALTER INDEX statement.
- C: You should not consider disabling the index and executing the statement in the scenario because a clustered index cannot be rebuilt online if it was disabled when CREATE INDEX was executed.

QUESTION 10

You work as the network database administrator at Mondo Transport. The Mondo Transport network consists of a single Active Directory domain named Mondo Transport. All servers on the Mondo Transport network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the Mondo.com network CIO to start performance tuning a database and database application. You will be responsible for optimizing indexing strategies. The induction developers of Mondo Transport make use of an application which accesses a database hosted on a SQL Server 2005 computer. The application is used to analyze the data about the daily yield in each location and the impact of wear and tear data and adding the information into a large number of rows to tables daily.

You have discovered that overtime the result of the information added has resulted in performance loss due to severely fragmented clustered and nonclustered indexes. You are required to design a solution which will correct the problem whilst ensuring the solution minimizes the reliance of manual operation.

What should you do?

- A. You should create a scheduled maintenance plan to periodically rebuild the database indexes.
- B. You should create a scheduled maintenance plan to periodically reorganize the database indexes.
- C. You should drop all the indexes and create a scheduled maintenance plan to periodically update distribution statistics.
- D. You should drop all nonclustered indexes and add the key columns from the nonclustered indexes as key columns to the clustered indexes.

Answer: A

Explanation: You should remember in the scenario that adding rows and deleting rows to database tables can cause severe fragmentation overtime regarding the amount of information entered and deleted. The best solution would effectively be avoiding adding and deleting rows to tables or periodically rebuilding the indexes and tables.

1. Mondo Transport will have the induction developers typically analyzing the data from a single branch office whilst taking samples from approximately 20 locations at each branch office.
2. Mondo Transport additionally requires the induction developers in the Houston main office to be able to analyze the data about the daily yield in each location and the impact of wear and tear data.
3. Mondo Transport is currently aware that the WearTear table will be expected to become quite large.

Incorrect Answers:

- B: You should not consider simply reorganizing the indexes in the scenario because the reorganizing will only reorganize the leaf nodes and severely fragmented indexes require rebuilding which involves reorganizing nonleaf and leaf nodes.
- C: You should not consider following the instructions in this answer in the scenario because dropping all indexes would only further decrease the system performance and will not prevent the fragmentation.
- D: You should not consider following the instructions in this answer in the scenario because dropping all the nonclustered indexes and adding their key columns to the clustered indexes will not resolve the problem and would only further degrade performance.

Topic 6, Culinary-Art, Scenario

Company overview

Culinary Art, Inc. has twenty-five locations around the United States: Chicago (head

office), and various distribution centers.

Products and Services on offer:

The company operates as an Internet-based department store chain. Culinary-Art, Inc. service clients on a country-wide basis and thus operates online.

Home delivery in selected areas is part of the service offered. All customers are to provide their ZIP code during registration. The nearest distribution centre for order delivery is identified when the user logs in.

Service is limited to a 50-mile radius of each distribution centre. Web site visitors outside of the distribution radius can explore the Web site and the inventory available from any distribution centre, but cannot place orders.

Future Plans

Culinary-Art, Inc plans to expand operations in future. To this end support must be provided for sales via other third party companies. These companies are to have their own Web sites. Their respective customer's orders will be packed for delivery at Culinary-Art, Inc. and the third party delivery personnel will then pick up finished orders to deliver to their customers.

EXISTING ENVIRONMENT

Existing Application and Supporting Environment

Culinary-Art, Inc at presents consists of 25 distribution centers. Each of these distribution centers has as internal network that includes a database server running SQL Server 2005. A Perimeter network isolates the internal network from the Internet.

The head office in Chicago currently hosts the Web servers. These Web servers are the only servers located in a private perimeter network.

Each of the distribution centers is responsible for the maintenance of its own inventory. Their database servers are configured as federated database servers. Local copies of all distribution centers' inventories in a custom designed database in a local SQL Server instance and customer registration information are hosted on the Web servers.

Registration information is also written to the local database servers.

The Chicago office supports the company's operations servers and two distribution warehouse servers. All financial accounting information and a master registered visitor list forms part of the information that is maintained by the Chicago office. Registered visitors are identified as either customers or visitors-only. The Chicago office also hosts one database server used as an archival records server. The archival records server is also used as a data source for data transform and load to a business intelligence server.

The servers in Chicago are configured for Windows authentication only.

All internal network database servers enjoy Active Directory membership.

The Web servers are configured as stand-alone servers.

The SQL server on the Web servers is configured for mixed mode authentication.

BUSINESS REQUIREMENTS

Each of the twenty-five distribution centers will be responsible for their own maintenance of detailed customer information. The Chicago office will maintain summary customer information and sensitive information such as credit card information. Each customer should be registered with and appear in only one local database. Culinary-Art, Inc. employees can access combined detailed customer information through a distributed view named Sales.v_All_Cust.

The Web server runs a stored procedure that generates orders. These orders are posted to

the database server in each Distribution centre internal network from the Web servers. Consolidation of orders is done at the Chicago office. To this end it must be possible to easily identify the source office for each order. Orders for the current quarter should be kept separate from orders for previous quarters.

Inventory maintenance at the distribution centers is maintained through just-in-time ordering with most inventories coming from local vendors. Purchase order stocking points are determined automatically, and vendor purchase orders are placed from the Chicago office. Local offices are authorized to order to fill spot shortages and for specialty items that require pick-up from the vendor.

Culinary-Art, Inc. policy dictates that purchase orders can be kept locally for up to 10 days, but must be kept until the order is either received (and given a status of Closed) or cancelled. In the event of an order not received within 10 days it should be cancelled.

Archival purchase order records are used to update the financial accounting server.

Replication should be administered locally from the archival server, but distribution centers need to have control over specific aging requirements for order archives.

Distribution centers sometimes need to temporarily change the length of time that they hold purchase orders, such as when there is a dispute with a vendor over charges.

All periodic reporting requirements rely on data maintained in the Chicago office.

The company needs to deploy local Web servers at each distribution centre. This is a priority project and all distribution centers will be held responsible for posting its own orders. Also, when this happens, orders will be delivered to Chicago at 19:00 hours on a daily basis.

TECHNICAL REQUIREMENTS

The company must identify the requirements for local office Web servers. After local deployment, each Web server must maintain a current copy of the company's inventory without cost information. Web server database server instances are configured for the simple recovery model.

All servers deployed on any local network must use Active Directory user accounts for server access. User credentials must be protected against unauthorized disclosure.

Following is the hard disk configuration of each of the local database servers:

1. 180 GB IDE hard disk as drive C:
2. 180 GB IDE hard disk as drive D:
3. 900 GB SCSI RAID 5 storage unit as drive E:

Local database servers should automatically report serious errors to the Chicago office.

All IT staff members need to be paged immediately, if possible, when errors occur. The IT staff also needs periodic information about selected resource use.

Local database servers must provide inventory updates to the Web servers through replication. Orders generated on the Web servers are submitted to the local database servers through locally deployed Web services configured at each local network.

Culinary-Art, Inc management wants to move the Web servers to the local networks as soon as possible. After you move the Web servers to the distribution centers, they will use the same network configuration currently used in the Chicago office.

The following database servers are hosted in the Chicago office:

1. CART-DB01 - The CART-DB01 server contains summary customer records, purchase order records, consolidated customer orders, and other information needed for daily operations. The CART-DB01 server has one database named BusOps. This database


server is configured as shown in the CART-DB01 Disk exhibit. Each filegroup contains a single read-write file.

1. CART-DBSecure - The CART-DBSecure server contains accounting records, secure employee records, customer credit information, and other sensitive data.

2. CART-DBDist1

1. CART-DBDist2 - CART-DBDist1 and CART-DBDist2 are distribution center servers and are configured like the servers for the other local office database servers.

CART-DB01 Disks



The BusOps database is configured for the full recovery model. Following is the backup schedule:

1. A monthly full database backup on the first Saturday of each month.
 2. A weekly full backup of the Primary filegroup on Sunday nights.
 3. A weekly full backup of the Secondary filegroup on Wednesday nights.
 4. A daily differential backup of each filegroup before start of business at 05:00 A.M.
- Availability of the database server in the Chicago office at all times must be guaranteed. There must be a minimization of potential single points of failure.

Topic 6, Culinary-Art, Inc. (16 Questions)

QUESTION 1

Culinary-Art, Inc. plans on starting a quarterly commission program for all its personnel working in the distribution centers. This program will include the managers as well. Commission will be based on on-time order preparation and will be calculated on a quarterly basis.

Each order record includes the employee who pulled the order and prepared it for shipment. The Employee table on each distribution centre database server includes a ManagerID column that has a foreign key constraint relating to the EmployeeID column. The same information is maintained on the consolidated Employee table at the Chicago office.

The commission calculating process involves the following steps:

1. Retrieval of all qualifying order for the distribution centre from the OrderHead and OrderDetail.
2. For each qualifying order, calculation of the employee and manager commission amount.
3. Generation of running totals by employee and by manager.
4. After record processing, update of the commissions table with employees, manager, current quarter and commission amount.

The calculation and storage of the commission amounts will be done through a

stored procedure. A custom application will execute the stored procedure from a client from part of a custom quarterly activities application. Following are the requirements that should be met by this application:

1. Interference with data access for other applications must be minimized.
2. Resource requirements at the server should be minimized.
3. Network traffic requirements should be minimized.

You thus need to identify the data access components that should be used and the required transaction isolation level. The stored procedure will be written based on your recommendations.

What should you do?

- A. Create a Fast Forward-only cursor based on OrderHead and OrderDetail and a Dynamic cursor based on the Commissions table. Make use of the Snapshot transaction isolation level.
- B. Create a Fast Forward-only cursor based on OrderHead and OrderDetail and a table variable based on the Commissions table. Make use of the Repeatable Read transaction isolation level.
- C. Create Dynamic cursors, one based on OrderHead and OrderDetail tables and the other based on the Commissions table. Make use of the Snapshot transaction isolation level.
- D. Create a Fast Forward-only cursor based on OrderHead and OrderDetail and a table variable based on the Commissions table. Make use of the Snapshot transaction isolation level.

Answer: D

Explanation: You should create a Fast Forward-only cursor based on OrderHead and OrderDetail and a table variable based on the Commissions table. The stored procedure will only require going through the order, in any order, making its calculations and processing the results. Data is read in from the tables, after that, there is no further need for access to the tables. Results can then be stored in a table variable that can be used to update the Commissions table. Table access, and thereby interference with other applications will then be minimized.

You should further make use of the Snapshot transaction isolation level because this level does not request or order locks. Thus there will be no interference with other applications. Since you are using source data from a time period that has already passed, there is no need to worry about source data changing.

Incorrect Answers:

A: This option is partly correct, but you should not create a Dynamic cursor that is based on the Commissions table. You could make updates through the Commissions table; however, it is not a requirement. Also a Dynamic cursor will interfere with other applications because a dynamic cursor will require interactive access to the Commissions table.

B: This option is only partly correct. You should not make use of the Repeatable Read transaction isolation level because it has a locking requirement that will interfere with the other tables and their activities. Since you will be reading static data there is no need for

interactive access to the data source tables.

C: Creating dynamic cursors as suggested in this option is not the solution. This option is only partly correct. The calculations required for sales commissions do not require dynamic access because the data should be static at this point in time. There is absolutely nothing that prompts you to update the commissions table as the amounts are calculated. Besides, you can maintain these in memory until after all of the calculations are completed, and make the updates to the commissions table when finished. Dynamic cursors will definitely interfere with other applications as well as require more resources than creating a Fast Forward-only cursor based on OrderHead and OrderDetail and a table variable based on the Commissions table and making use of the Snapshot transaction isolation level.

QUESTION 2

You need to perform some performance tuning for the Culinary-Art, Inc. database and applications. Following exhibit displays a query that is used to generate a list of items that is sold during a day:

```
SELECT sku, description
FROM sales.inventory i WHERE EXISTS
(SELECT *FROM sales.[invoice detail] d, sales.[invoice head] h
WHERE i.sku = d.sku
AND h.orderid = d.orderid
AND DATEPART (dy, h.orderdate) = DATEPART (dy, GETDATE())
AND DATEPART (yyyy, h.orderdate) = DATEPART (yyy, GETDATE())
```

You must rewrite this query converting it to a join that will return the same information.

What should you do?

A. Rewrite the query as follows:

```
SELECT i.sku, description
FROM sales.inventory i JOIN sales.[invoice detail] d
ON i.sku = d.sku JOIN
sales.[invoice head] h ON h.orderid = d.orderid
AND DATEPART (dy, h.orderdate) = DATEPART (dy, GETDATE())
AND DATEPART (yyyy, h.orderdate) = DATEPART (yyy, GETDATE())
```

B. Rewrite the query as follows:

```
SELECT DISTINCT i.sku, description
FROM sales.inventory i JOIN sales.[invoice detail] d
ON i.sku = d.sku JOIN
sales.[invoice head] h ON h.orderid = d.orderid
AND DATEPART (dy, h.orderdate) = DATEPART (dy, GETDATE())
AND DATEPART (yyyy, h.orderdate) = DATEPART (yyy, GETDATE())
```

C. Rewrite the query as follows:

```
SELECT sku, description
FROM sales.inventory i JOIN sales.[invoice detail] d
ON i.sku = d.sku JOIN
sales.[invoice head] h ON h.orderid = d.orderid
```

```
AND DATEPART (dy, h.orderdate) = DATEPART (dy, GETDATE())
AND DATEPART (yyyy, h.orderdate) = DATEPART (yyy, GETDATE())
D. Rewrite the query as follows:
SELECT DISTINCT sku, description
FROM sales.inventory i JOIN sales.[invoice detail] d
ON i.sku = d.sku JOIN
sales.[invoice head] h ON h.orderid = d.orderid
AND DATEPART (dy, h.orderdate) = DATEPART (dy, GETDATE())
AND DATEPART (yyyy, h.orderdate) = DATEPART (yyy, GETDATE())
```

Answer: B

Explanation: In the original subquery, one row is returned from the sales.inventory table for each SKU value found in sales.[invoice detail] for the current date. In the subquery following the EXISTS keyword, the following limits the result to the current day only:

```
AND DATEPART (dy, h.orderdate) = DATEPART (dy, GETDATE())
AND DATEPART (yyyy, h.orderdate) = DATEPART (yyy, GETDATE())
```

Also, d.order = h.orderid limits the result to orders for that specific date. Due to the outer SELECT statement being based on the sales.inventory table and passing those values to the inner query, the results are limited to distinct rows in sales.inventory.

In the join, you have the same limit on date, but you need to return the SKU values for each match with a SKU value in the sales.[invoice detail] table. In the event of a SKU appearing more than once in that table, it will appear more than once in the result if you did not include the DISTINCT keyword.

Incorrect Answers:

A: With this join statement the result will not be the same because the DISTINCT keyword is not included in the SELECT clause. This will result in the return of one or more rows with the same SKU; depending on how many times the SKU appears on that particular day's orders. You need to limit the result by using the DISTINCT keyword.

C: This statement will result in a syntax error due to the column name sku in the SELECT clause being ambiguous since this column is contained in both sales.inventory and sales.[invoice detail]. You need to identify which sku column is being used in the SELECT clause and then qualify it with the table name or the table alias.

D: This statement will result in a syntax error due to the column name sku in the SELECT clause being ambiguous since this column is contained in both sales.inventory and sales.[invoice detail]. You need to identify which sku column is being used in the SELECT clause and then qualify it with the table name or the table alias.

QUESTION 3

You are tasked with the responsibility to design the appropriate data access strategies for Culinary-Art, Inc. future ventures regarding third party company involvement. A development team is working with you on this assignment. In your solution do bear in mind that:

1. The data access solution should minimize the distribution servers' security exposure.

2. No third parties will be given any information regarding the underlying data structures at the database.
3. Provision must be made for third party developers to retrieve information regarding requirements for data access to use when building their own solutions.
4. All documentation is to be updated any time the requirements changes.
5. The time required for document updates should be minimized.
6. The data access solution must be both .NET Framework Windows applications and .NET ASP.NET application compatible without requiring any additional client software.

What should you do?

A. The distribution servers should be configured to support remote Web client access. Use stored procedures and user-defined functions (UDFs) to provide access to server data.

B. The distribution servers must have HTTP endpoints and XML Web Services employed.

Use stored procedures and user-defined functions (UDFs) to provide access to server data.

C. The distribution servers must have HTTP endpoints and XML Web Services employed

Use custom views that limit the data available for read or write access.

D. Use a custom ADO.NET SQL Web application that the third party companies can download.

Modify the source code, and redeploy on their Web servers with parameters for data access embedded as comments. Use stored procedures and user-defined functions (UDFs) to provide access to server data.

Answer: B

Explanation: Access to the server can be granted via implementing HTTP endpoints and XML Web services at the distribution servers. Access to data can be provided via stored procedures and user-defined functions. HTTP endpoints will minimize exposure and is compatible with .NET Framework Windows applications and .NET ASP.NET application compatible without requiring any additional client software. This will also minimize the amount of documentation required and the time needed to keep it up to date because Web services provides documentation for itself through Web Services Description Language.

1. Culinary-Art, Inc plans to expand operations in future. To this end support must be provided for sales via other third party companies. These companies are to have their own Web sites. Their respective customer's orders will be packed for delivery at Culinary-Art, Inc. and the third party delivery personnel will then pick up finished orders to deliver to their customers.

Incorrect Answers:

A: The distribution servers should not be configured to support remote Web client access. This will increase the servers' exposure. You will also require manually documentation of any access components. Furthermore Client applications that directly

access the database servers will also require SQL client communication components.

C: Creating custom views will expose more information regarding the underlying structures. It would be better to make use of stored procedures and UDFs.

D: Creating a custom ADO.NET SQL Web application as this option suggests will have the source code expose information regarding the Culinary-Art servers with parameters for data access embedded as comments. This option will also require that you find a way to document and distribute any changes that would impact access by the Web application.

QUESTION 4

You need to create the stored procedure to process completed customer orders and update customer account information for Culinary-Art, Inc. to meet performance specifications for the procedure it is imperative that the cursor opens as quickly as possible. The cursor is used to retrieve data from a single table. Your solution must ensure that the cursor opens as quickly as possible and aid in optimizing procedure performance.

You now need to make a decision as to which DECLARE CURSOR statement must be used to support the stored procedure in this case.

What should you do?

- A. Make use of a dynamic cursor.
- B. Make use of a synchronous static cursor.
- C. Make use of a read-only static cursor.
- D. Make use of a fast-forward only keyset-driven cursor.

Answer: A

Explanation: Of the given choices a dynamic cursor would best suit the stored procedure. A dynamic cursor opens faster than both static and keyset-driven cursors.

Incorrect Answers:

B: A synchronous static cursor will not meet all the requirements since it opens slower than a dynamic cursor.

C: A Read-Only Static cursor is not as quick to open as a dynamic cursor.

D: A fast-forward only keyset-driven cursor is not the solution since it does not open as quickly as a dynamic cursor. This type of cursor is means to support absolute fetches. This will not be optimizing the stored procedure if used.

QUESTION 5

Culinary-Art wants all data storage solutions to be optimized. You are currently modifying the Customer table in each distribution database to support a comment field. One requirement that should be kept in mind is that the field must be able to accommodate 2000 characters. All comments are stored in U.S. English.

You now need to change the table to add the Comment column. You must tune the database performance so that the table row size is minimized.

What should you do?

- A. Use a text column.
- B. Use a varchar(max) column.
- C. Use a varchar(2000) column
- D. Use a nvarchar(2000) column.

Answer: B

Explanation: SQL Server 2005 minimizes the in-table row space required for large object (LOB) data types. These data types include varchar(max). SQL Server stores a 24 byte pointer in the table row that points to the data page used to contain the column data.

Incorrect Answers:

A: A text data type is not the solution. Support for the text data type is provided for backward compatibility. But the text data type should not be used for any new development.

C: Actual space required for data storage is determined by the data contained in the column. In this case it could be as much as 2000 bytes for varchar(2000). Thus you should not create a varchar(2000) column.

D: you should not create a nvarchar(2000) column to store the comments because the actual space required for data storage is determined by the data contained in the column. And in this case this space could amount to 4000 bytes for nvarchar(2000).

QUESTION 6

Since there is a need to transfer inventory between distribution centers, you are designing a stored procedure that will facilitate this inventory transfers. These transfers will run in the context of a transaction that will ensure that both ends of the transfers are made. Part of this stored procedure that you are designing requires reading data from the source inventory file so as to ensure that the required stock is indeed available. You must ensure that the transaction will always obtain a shared lock as data is read. This lock should be released after the required data has been read regardless of database option settings.

What should you do?

- A. Specify the SERIALIZABLE option in the inventory data read - query by using the OPTION clause.
- B. Specify the READUNCOMMITTED option in the inventory data read - query by using the OPTION clause.
- C. Specify the READCOMMITTEDLOCK option in the inventory data read - query by using the OPTION clause.
- D. Specify the READCOMMITTED option in the inventory data read - query by using the OPTION clause.

Answer: C

Explanation: The READCOMMITTEDLOCK option will override database option settings and cause the database engine to obtain a shared lock when data is read and

release the lock when the read option is completed.

Incorrect Answers:

A: The SERIALIZABLE option holds locks until a transaction is complete and will not release the locks when a resource is no longer required.

B: The READUNCOMMITTED option does not obtain locks when reading data.

D: The READCOMMITTED option depends on the READ_COMMITTED_SNAPSHOT database option. When this option is OFF, it will result in the same effect as the READCOMMITTEDLOCK option. But when it is ON the database engine will not obtain locks rather, it will make use of row versioning. This option is not the solution.

QUESTION 7

The Culinary-Art, Inc management wants improvement of performance to be applied in all solutions. You are currently busy designing transaction strategies to accommodate the importing of inventory updates received from suppliers. In your solution you must ensure that the lock granularity is configured to allow multiple threads to bulk load data. This lock that is obtained must also prevent other processes (those processes that are not bulk load processes) from accessing the table during the bulk loading.

What should you do?

A. Set the table lock on bulk load option on the destination table using the sp_table option.

B. Set the insert row lock on bulk load option on the destination table using the sp_table option.

C. Run SET TRANSACTION ISOLATION LEVEL REPEATABLE READ prior to running BULK INSERT.

D. Configure the database to use the bulk-logged recovery model.

Answer: A

Explanation: The type of lock described in this scenario is a bulk update lock.

Setting the table lock on bulk load option on the destination table will cause the database engine to obtain a bulk update (BU) lock on the table when you start the bulk load procedure. And it will allow multiple connections for bulk load but blocking other types of access.

Incorrect Answers:

B: Setting the insert row lock option on the destination table will be ignored by and will not be used by SQL Server 2005. It is only supported as an option that only ensures compatibility with down-level batches and scripts. Thus this option will not have any effect on locking behavior when the procedure is called.

C: Running the SET TRANSACTION ISOLATION LEVEL REPEATABLE READ prior to running BULK INSERT will not make any difference as to whether a BU lock is obtained or not on the table.

D: There is no need to configure the database to use bulk-logged recovery model because

it has nothing to do with whether a BU lock is obtained or not when the procedure is called.

QUESTION 8

The Culinary-Art, Inc. personnel require the ability to view customer order information. To this end you are designing a query that will allow one to see CustomerID values for all customers as well as: the orders that are placed by individual customers, and customers that has not yet placed any orders. There is a foreign constraint on the CustomerID column of the Sales.[Order Head] table that references the Sales.Customer table CustomerID column. You want to run a single query, but you must ensure that the query return CustomerID and OrderID.

What should you do?

A. Run the following:

```
SELECT c.CustomerID, o.OrderID
FROM Sales.Customer c LEFT OUTER JOIN Sales.[Order Head] c
ON c.CustomerID = o.CustomerID
```

B. Run the following:

```
SELECT o.CustomerID, o.OrderID
FROM Sales.Customer c LEFT OUTER JOIN Sales.[Order Head] c
ON c.CustomerID = o.CustomerID
```

C. Run the following:

```
SELECT c.CustomerID, o.OrderID
FROM Sales.Customer c RIGHT OUTER JOIN Sales.[Order Head] c
ON c.CustomerID = o.CustomerID
```

D. Run the following:

```
SELECT c.CustomerID, o.OrderID
FROM Sales.Customer c JOIN Sales.[Order Head] c
ON c.CustomerID = o.CustomerID
```

Answer: A

Explanation: This option shows a LEFT OUTER JOIN. All rows are thus returned from the left table, Sales.Customer, whether there are matching rows in the right table or not. This means that all CustomerID values from Sales.Customer will be returned. In the event of the customer having an order, those OrderID values are also returned.

Incorrect Answers:

B: Although this option represents a LEFT OUTER JOIN, the result will only return values from the Sales.[Order Head] table because these are the only values that are specified in the SELECT clause.

C: This option shows that it is a RIGHT OUTER JOIN. This will result in all rows in Sales.[Order Head] being returned regardless of whether there are matching rows in Sales.Customer or not. Due to the foreign key relationship between Sales.[Order Head] and Sales.Customer, there shouldn't be CustomerID values in Sales.[Order Head] that

does not appear in Sales.Customer.

D: Although not explicitly identified, this option represents an INNER JOIN. Only rows meeting the qualifying condition are returned. Thus if a CustomerID does not appear in the Sales.[OrderHead] table, nothing will be returned for that particular customer.

QUESTION 9

You are designing the database query strategy for Culinary-Art. To this end you are creating a query that will return information regarding the customer orders for separate distribution centers. This information should include:

1. CustomerID (Sales.Customer and Sales.[Invoice Head] tables)
2. InvoiceNumber (Sales.[Invoice Head] and Sales.[Invoice Detail] tables)
3. Each ItemNumber (Sales.[Invoice Detail] table)
4. Partnumber (Sales.[Invoice Detail] and Sales.Inventory tables)
5. Description (Sales.Inventory Table)
6. Total cost of the item (Sales.[Invoice Detail] table - Quantity and SellPrice columns)

The result of the query should be limited to a specific date and the resources needed by the query should be minimized (use April 20, 2006 as the sample date for the sample query). You must ensure that your solution will result in a query that will return this information with columns in the order listed and in InvoiceNumber order.

What should you do?

A. Use the following query:

```
SELECT c.CustomerID, InvoiceNumber, d.ItemNumber, d.PartNumber, i.Description,  
(Quantity * SellPrices) as [Item Total]  
FROM Sales.[Invoice Head] h, Sales.[Invoice Detail] d, Sales.Inventory I,  
Sales.Customer c  
WHERE c.CustomerID = h.CustomerID, h.InvoiceNumber = d.InvoiceNumber AND  
d.PartNumber = i.PartNumber AND InvoiceDate = '4/20/2006'  
ORDER BY h.InvoiceNumber
```

B. Use the following query:

```
SELECT CustomerID, h.InvoiceNumber, d.ItemNumber, d.PartNumber, i.PartNumber,  
i.Description, SUM(Quantity * SellPrices) as [Item Total]  
FROM Sales.[Invoice Head] h, Sales.[Invoice Detail] d, Sales.Inventory i  
WHERE h.InvoiceNumber = d.InvoiceNumber AND d.PartNumber = i.PartNumber  
AND InvoiceDate = '4/20/2006'  
ORDER BY h.InvoiceNumber
```

C. Use the following query:

```
SELECT CustomerID, h.InvoiceNumber, d.ItemNumber, d.PartNumber, i.Description,  
(Quantity * SellPrices) as [Item Total]  
FROM Sales.[Invoice Head] h, Sales.[Invoice Detail] d, Sales.Inventory i  
WHERE h.InvoiceNumber = d.InvoiceNumber AND d.PartNumber = i.PartNumber  
AND InvoiceDate = '4/20/2006'  
ORDER BY h.InvoiceNumber
```

D. Use the following query:

```
SELECT CustomerID, InvoiceNumber, ItemNumber, d.PartNumber, Description,
```

```
(Quantity * SellPrices) as [Item Total]
FROM Sales.[Invoice Head] h, Sales.[Invoice Detail] d, Sales.Inventory i
WHERE h.InvoiceNumber = d.InvoiceNumber AND d.PartNumber = i.PartNumber
AND InvoiceDate = '4/20/2006'
ORDER BY InvoiceNumber
```

Answer: C

Explanation: This particular query will result in the return of the information as result set columns. The only joined columns are those needed to retrieve information, h.InvoiceNumber = d.InvoiceNumber and d.PartNumber = i.PartNumber. Implementation of the join operation is via the search conditions, and ordering is imposed by the ORDER clause.

Incorrect Answers:

A: This query does not require any additional information from the Sales.Customer table. Thus there is no reason that it should be included in the FROM list or even in the join logic c.CustomerID - h.CustomerID. This will just result in increased processing that requires extra disk resources.

B: This query will result in a syntax error. You are required to obtain a calculated value for the extended price and not an aggregate summation.

D: This query will generate a syntax error because there is an ambiguity in the Invoice Number in the SELECT and ORDER BY clauses when it is already included in the Sales.[Invoice Head] and Sales.[Invoice Detail] tables.

QUESTION 10

You are designing the data access components required for an application that will calculate stocking levels together with the Culinary-Art database programmers. This application will be used to calculate the suggested stocking levels for each of the twenty-five distribution centers. The sales history for each individual distribution centre will be used to do the calculations per distribution centre respectively. The application will process records sequentially and update the SuggestedStock column only. No other applications will be used to update this column.

You must ensure that the cursor opens as quickly as possible when declared so as to prevent other applications from accessing or changing other values on rows while it is running. You thus need to decide on the appropriate cursor type and cursor concurrency option to use.

What should you do?

- A. Make use of a dynamic cursor and use the SCROLL LOCKS concurrency option.
- B. Make use of a keyset-driven cursor and use the READ_ONLY concurrency option.
- C. Make use of a keyset-driven cursor and use the OPTIMISTIC WITH VALUES concurrency option.
- D. Make use of a dynamic cursor and use the OPTIMISTIC WITH VALUES concurrency option.

Answer: D

Explanation: A dynamic cursor opens faster than a keyset-driven cursor. This thus makes a dynamic cursor the appropriate choice in this case. The OPTIMISTIC WITH VALUES concurrency option does not place any locks on the underlying table. With this option, it is assumed that there is little if any chance of another application changing the same values that will be changed by this application.

Incorrect Answers:

A: This option is only partly correct, however, SCROLL LOCKS concurrency option does place locks on the underlying table and the requirements specify that you should avoid any impact on any other applications. Locking rows will definitely affect the other applications because it will prevent them from accessing the data that they require.

B: A keyset-driven cursor, used in conjunction with the READ_ONLY concurrency option will not work in this case since you are updating the inventory records through this cursor. You should rather make use of a dynamic cursor and the OPTIMISTIC WITH VALUES concurrency option.

C: This option is only partly correct except for the use of the key-set-driven cursor. This type of cursor takes longer to open than a dynamic cursor due to the fact that it must build an internal temporary work table. This overhead will thus result in slower opening performance.

QUESTION 11

You are designing the database query strategy for Culinary-Art. Culinary-Art needs a list of top selling items by quantity from Sales.[Invoice Detail]. The 20 top-selling items in descending order should be returned by this list. The items' unique SKU, description (from the Inventory table), and total number of units sold must be included in the query result.

The SKU column in the Sales.Inventory table has a unique constraint applied.

There is a foreign key constraint on the SKU in Sales.[Invoice Detail] that references Sales.Inventory.

You need to create a query that will retrieve this information.

What should you do?

A. You should run the following statement:

```
SELECT TOP 20 o.SKU, i.Description, SUM(o.Quantity)AS[Count]
FROM Sales.[Invoice Detail] o, Sales.Inventory i
WHERE o.partnumber = i.partnumber
GROUP BY o.SKU ORDER BY SUM (o.Quantity) DESC
```

B. You should run the following statement:

```
SELECT TOP 20 o.SKU, i.Description, SUM(o.Quantity)AS[Count]
FROM Sales.[Invoice Detail] o, Sales.Inventory i
WHERE o.partnumber = i.partnumber
GROUP BY o.SKU, i.Description ORDER BY SUM (o.Quantity) DESC
```

C. You should run the following statement:

```
SELECT TOP 20 o.SKU, i.Description, SUM(o.Quantity)AS[Count]
FROM Sales.[Invoice Detail] o, Sales.Inventory i
```

```
WHERE o.partnumber = i.partnumber
GROUP BY o.SKU, i.Description ORDER BY SUM (o.Quantity)
D. You should run the following statement:
SELECT TOP 20 o.SKU, i.Description, SUM(o.Quantity)AS[Count]
FROM Sales.[Invoice Detail] o, Sales.Inventory i
WHERE o.partnumber = i.partnumber
GROUP BY SUM (o.Quantity) GROUP BY o.SKU, i.Description
E. You should run the following statement:
SELECT TOP 20 o.SKU, i.Description, SUM(o.Quantity)AS[Count]
FROM Sales.[Invoice Detail] o, Sales.Inventory i
WHERE o.partnumber = i.partnumber
GROUP BY o.SKU, i.Description ORDER BY COUNT(o.Quantity) DESC
```

Answer: B

Explanation: The TOP 20 keyword limits the result to the top 20 items.
The SELECT clause includes the required columns (including the computed column named Count)
The SUM aggregate will add the Quantity values for each SKU.
The inner join implemented by means of the WHERE clause is required to retrieve the description from Sales.Inventory.
The GROUP BY clause must include all columns in the SELECT clause except aggregates.
The ORDER BY clause sorts the result by the total quantity, calculated as SUM(o.Quantity), in descending order.

Incorrect Answers:

A: This option would be correct except for this part: GROUP BY o.SKU ORDER BY SUM (o.Quantity) DESC. This will result in a syntax error because Description is not included in the GROUP BY clause. And the GROUP BY clause must include all non-aggregate columns from the SELECT clause.

C: This option would be correct except for this part: GROUP BY o.SKU, i.Description ORDER BY SUM (o.Quantity). This statement does not have the ORDER BY clause including the DESC keyword. This means the sums will be sorted in ascending order instead of descending order. In fact this statement will result in returning the 20 lowest selling items in the Sales.[Invoice Detail] column.

D: This option would be correct except for this part: GROUP BY SUM (o.Quantity) GROUP BY o.SKU, i.Description. When a SELECT statement uses both GROUP BY and ORDER BY, the GROUP BY clause must precede the ORDER BY clause.

E: This option would be correct except for this part: GROUP BY o.SKU, i.Description ORDER BY COUNT(o.Quantity) DESC. This statement makes use of COUNT instead of SUM. This will thus result in a count of the number of line items that used that SKU; and have a value in Quantity and not a sum of the quantity values.

QUESTION 12

The need to fill customer order sometimes necessitates the transfer of inventory between the Chicago distribution warehouses. Instead of physically moving the

inventory the order is transferred to the alternative warehouse if the initial warehouse cannot fill the order. In such a case, you have to drop the original order and create a new order for the other warehouse. You need to design the Culinary-Art transaction strategy.

In your solution you need to ensure that both these operations complete before the order is processed. If any part of the process can complete, none of the process should be completed.

What should you do?

- A. Create a stored procedure to drop the old order, create a new order and process it
- B. Create a batch to drop the old order, create a new order and process it.
- C. Drop the old order, create a new order. Then process it in the context of the same distributed transaction.
- D. Drop the old order and create the new order in the context of the same distributed transaction and process it as a separate local transaction.

Answer: C

Explanation: Dropping the old order, creating a new order and then processing it in the context of the same distributed transaction is the correct solution. This needs to be run as a distributed transaction because it affects different databases on different servers. If one completes all steps in one transaction, one can have the transaction roll back changes in case an error occurs.

Incorrect Answers:

A: There is no need to create a stored procedure to drop the old order, create a new order and process it as it will not cause the actions to occur in the context of the same transaction. This means that rollback of any changes might not be supported.

B: There is no need to create a batch to drop the old order, create a new order and process it as it will not cause the actions to occur in the context of the same transaction. This means that rollback of any changes might not be supported.

D: Do not drop the old order and create the new order in the context of the same distributed transaction and process it as a separate local transaction. In case an error occurs in the transaction in which the order is processed, the error will not cause the old and new orders to roll back.

QUESTION 13

You need to create the stored procedure to fulfill a number of requirements. The stored procedure must calculate suggested stocking values using local inventory information at each of the Culinary-Art distribution centers. Changes that are made to the inventory while the procedure is operational must also be taken into account because the cursor used must support absolute fetches. Your solution must ensure that the cursor opens as quickly as possible and aid in optimizing procedure performance.

You now need to make a decision as to which DECLARE CURSOR statement must be used to support the stored procedure.

What should you do?

- A. Make use of a dynamic cursor.
- B. Make use of a static cursor.
- C. Make use of a keyset-driven cursor.
- D. Make use of a fast-forward cursor.

Answer: C

Explanation: Of the given choices a keyset-driven cursor would best suit the stored procedure. An absolute fetch is a random fetch, directly fetching a specific row from the cursor. From the above choices only keyset-driven cursors display dynamic changes to the cursor rows. The cursor membership and row order do not change, but data values changes made by other applications are shown when fetching rows.

Incorrect Answers:

A: A dynamic cursor does show changes made to the cursor data when fetching rows, but it does not support absolute fetches. And even though a dynamic cursor loads quicker than a keyset-driven cursor, it does not meet all the requirements.

B: A Static cursor does support absolute fetches; however, a static cursor is also a snapshot of the data values when the cursor is loaded and this results in a situation where the changes are not shown when rows are fetched from the cursor.

D: A fast-forward cursor does not support absolute fetches and does not show data changes when rows are fetched. This will not be optimizing the stored procedure if used.

QUESTION 14

The Culinary-Art, Inc management wants improvement of performance to be applied in all solutions. You are currently busy designing transaction strategies to accommodate the importing of inventory updates received from suppliers. In your solution you must ensure that the lock granularity is configured to allow multiple threads to bulk load data. This lock that is obtained must also prevent other processes (those processes that are not bulk load processes) from accessing the table during the bulk loading.

What should you do?

- A. Run SET TRANSACTION ISOLATION LEVEL REPEATABLE READ prior to running BULK INSERT.
- B. Configure the database to use the bulk-logged recovery model.
- C. When running BULK INSERT to load the data, specify the TABLOCK hint.
- D. Set the insert row lock on bulk load option on the destination table using the sp_table option.

Answer: C

Explanation: The type of lock described in this scenario is a bulk update lock. Specifying the TABLOCK hint when running BULK INSERT to load the data will cause the database engine to obtain a bulk update (BU) lock on the table when you start the bulk load procedure. And it will allow multiple connections for bulk load

but blocking other types of access.

Incorrect Answers:

A: Running the SET TRANSACTION ISOLATION LEVEL REPEATABLE READ prior to running BULK INSERT will not make any difference as to whether a BU lock is obtained or not on the table.

B: There is no need to configure the database to use bulk-logged recovery model because it has nothing to do with whether a BU lock is obtained or not when the procedure is called.

D: Setting the insert row lock option on the destination table will be ignored by and will not be used by SQL Server 2005. It is only supported as an option that only ensures compatibility with down-level batches and scripts. Thus this option will not have any effect on locking behavior when the procedure is called.

QUESTION 15

The need to fill customer order sometimes necessitates the transfer of inventory between the Chicago distribution warehouses. Instead of physically moving the inventory the order is transferred to the alternative warehouse if the initial warehouse cannot fill the order. In such a case, you have to drop the original order and create a new order for the other warehouse. You need to design the Culinary-Art transaction strategy.

In your solution you need to ensure that both these operations complete before the order is processed. If any part of the process can complete, none of the process should be completed.

What should you do?

A. Create a trigger on the inventory table that will drop the old order, create a new order in the event of the inventory levels being insufficient.

B. Drop the old order and create the new order in the context of the same distributed transaction and process it as a separate local transaction.

C. Create a batch to drop the old order, create a new order and process it.

D. Drop the old order, create a new order. Then process it in the context of the same distributed transaction.

Answer: D

Explanation: Dropping the old order, creating a new order and then processing it in the context of the same distributed transaction is the correct solution. This needs to be run as a distributed transaction because it affects different databases on different servers. If one completes all steps in one transaction, one can have the transaction roll back changes in case an error occurs.

Incorrect Answers:

A: There is no need to create a trigger on the inventory table that will drop the old order, create a new order in the event of the inventory levels being insufficient. It will not implicitly cause all actions performed by the trigger to run in the context of the same transaction.

B: Do not drop the old order and create the new order in the context of the same

distributed transaction and process it as a separate local transaction. In case an error occurs in the transaction in which the order is processed, the error will not cause the old and new orders to roll back.

C: There is no need to create a batch to drop the old order, create a new order and process it as it will not cause the actions to occur in the context of the same transaction. This means that rollback of any changes might not be supported.

QUESTION 16

Culinary-Art, Inc. plans on starting a quarterly commission program for all its personnel working in the distribution centers. This program will include the managers as well. Commission will be based on on-time order preparation and will be calculated on a quarterly basis.

Each order record includes the employee who pulled the order and prepared it for shipment. The Employee table on each distribution centre database server includes a ManagerID column that has a foreign key constraint relating to the EmployeeID column. The same information is maintained on the consolidated Employee table at the Chicago office.

The commission calculating process involves the following steps:

1. Retrieval of all qualifying order for the distribution centre from the OrderHead and OrderDetail.
2. For each qualifying order, calculation of the employee and manager commission amount.
3. Generation of running totals by employee and by manager.
4. After record processing, update of the commissions table with employees, manager, current quarter and commission amount.

The calculation and storage of the commission amounts will be done through a stored procedure. A custom application will execute the stored procedure from a client from past of a custom quarterly activities application. Following are the requirements that should be met by this application:

1. Interference with data access for other applications must be minimized.
2. Resource requirements at the server should be minimized.
3. Network traffic requirements should be minimized.

You thus need to identify the data access components that should be used and the required transaction isolation level. The stored procedure will be written based on your recommendations.

What should you do?

- A. Create a Fast Forward-only cursor based on OrderHead and OrderDetail and a table variable based on the Commissions table. Make use of the Snapshot transaction isolation level.
- B. At the client load data from the OrderHead and OrderDetail tables and the Commissions table into DataSet objects. Make use of the Snapshot transaction isolation level.
- C. Create a Fast Forward-only cursor based on OrderHead and OrderDetail and a table variable based on the Commissions table. Make use of the Read Committed transaction isolation level.

D. Create dynamic cursors, one based on OrderHead and OrderDetail tables and the other based on the Commissions table. Make use of the Snapshot transaction isolation level.

Answer: A

Explanation:

You should create a Fast Forward-only cursor based on OrderHead and OrderDetail and a table variable based on the Commissions table. The stored procedure will only require going through the order, in any order, making its calculations and processing the results. Data is read in from the tables, after that, there is no further need for access to the tables. Results can then be stored in a table variable that can be used to update the Commissions table. Table access, and thereby interference with other applications will then be minimized.

You should further make use of the Snapshot transaction isolation level because this level does not request or order locks. Thus there will be no interference with other applications. Since you are using source data from a time period that has already passed, there is no need to worry about source data changing.

Incorrect Answers:

B: This option is partly correct, but you should not load data from the OrderHead and OrderDetail tables and the Commissions table into DataSet objects at the client. Data access requirements to the source tables are greater than the proposed solution due to reading that has to be done in the Commissions table. This will also result in increased network traffic because the data would have to be passed back to the server for processing in this option.

C: This option is only partly correct. You should not make use of the Read Committed transaction isolation level because it has a locking requirement that will interfere with the other tables and their activities. Since you will be reading static data there is no need for interactive access to the data source tables.

D: Creating dynamic cursors as suggested in this option is not the solution. This option is only partly correct. The calculations required for sales commissions do not require dynamic access because the data should be static at this point in time. There is absolutely nothing that prompts you to update the commissions table as the amounts are calculated. Besides, you can maintain these in memory until after all of the calculations are completed, and make the updates to the commissions table when finished. Dynamic cursors will definitely interfere with other applications as well as require more resources than creating a Fast Forward-only cursor based on OrderHead and OrderDetail and a table variable based on the Commissions table and making use of the Snapshot transaction isolation level.

Topic 7, WillowBridge, Scenario

BACKGROUND:

Company Overview

Willow Bridge is a company that specializes in oncology treatment. This expertise includes research, formulation, testing, distributing and sales of the oncology pharmaceutical products that are produced. The Willow Bridge sales offices are located

in Miami, Dallas and Los Angeles and the research facility in Chicago. The sales offices are connected to the Chicago office by demand-dial links.

Infrastructure and planned changes

Currently the Willow Bridge network is implemented as a single domain with four sites. Located in the Chicago research facility, there is a computer named WB-DB01.

WB-DB01 runs Microsoft SQL Server 2000. There is also a computer named WB-DB02 that runs Microsoft SQL Server 2005. The Willow Bridge management wants to upgrade WB-DB01 to SQL Server 2005.

The Willow Bridge management also wants to open another research facility in Berlin, Germany. When the Berlin research facility is brought online, they want it to be configured as a separate domain in a different forest. The Chicago and Berlin research facilities will make use of a virtual private network (VPN) connection for communication purposes.

The Willow Bridge management wants to expand business by selling their products and drugs on the company Web site. WB-DB02 will host a database to support the e-commerce application that will be for the benefit of medical practitioners world-wide. The e-commerce application will make use of Simple Object Access Protocol (SOAP) to retrieve product information and submit orders.

EXISTING ENVIRONMENT:

WB-DB01 supports the research initiatives. A database named Research hosts all information about drug manufacturing processes, laboratory trials, and clinical trials.

The Chicago research facility is divided into various laboratories where the experiments are run. A different table is used to store the data for each laboratory experiment. The schema varies somewhat based on the nature of the experiment. Before the experiment begins, the researcher provides the database development team with specifications of the values that need to be tracked and their tolerance range.

Regularly at the end of each month, reports are generated manually and e-mailed to the Chicago office. Monthly sales and invoice aging information is included in these reports.

The Chicago research facility also operates as the head quarters and as such handles debt collection. WB-DB02 supports the sales and accounting operations. All customers with an outstanding balance over 60 days are also sent to the Chicago office. The accounting department makes use of an application that imports XML data.

BUSINESS REQUIREMENTS:

Two applications are required to run Willow Bridge. These applications are named the Lab trial application and the Clinical trial application respectively.

Lab Trial Application

All laboratory trial data is accessed through a .NET application. This .NET application makes use of ADO.NET to access the data for each lab sample. On average each experiment consists of approximately 300 to 500 samples. At present Lab Trial application uses a keyset-driven cursor to fetch a single row at a time. The lab technician takes measurements and updates the row of data with the current sample data. Previous data is archived to a history table during the update. Lab technicians must be able to verify the changes they made to the sample data when they go back to check their work. In the past it has happened that some experimental samples have been rendered invalid due to data being entered incorrectly. One requirement for the Lab Trial application is thus that data entered by technicians must be validated to ensure that the values are within the

tolerance range for the experiment. In the event of data falling outside of the tolerance range being entered, an error message should be reported to the technician. Information such as the appropriate data range, and if required additional instructions should be contained in this report.

The same application is used to manage samples for all experiments. Lab technicians select the experiment name and the code makes use of a lookup table to discover the name of the table that stores the sample data.

Clinical Trial Application

Information regarding participants in clinical trials is stored in a separate table for each trial. The Willow Bridge researchers that run clinical trials often need to make follow-up telephone calls to participants. To this end they need to obtain a list of participants, who need a follow-up call; they need to select participants that have participated in a specific study or studies. They sometimes limit the query to those who have participated in one study, but not in another.

TECHNICAL REQUIREMENTS:

Performance

The lab technicians were complaining that they encounter serious performance lag when they attempt to access sample data. The current system is too slow. Perceived performance for accessing sample data must be improved in the updated system.

Availability

When the research lab is opened in Berlin, a database server running SQL Server 2005 will be installed there. This server will be named WB-DB03. Researchers at both the Chicago and Berlin research facilities will need to be able to query both WB-DB01 and WB-DB03 to share research data.

Topic 7, Willow Bridge (11 Questions)

QUESTION 1

You are busy designing the database query strategy for Willow Bridge that will retrieve the result set used to report on monthly sales trends. To this end you are designing the stored procedure. This stored procedure will be used by an application that makes use of Microsoft Visual C# .NET.

Of the requirements that you need to keep in mind is that the analysts need the ability to retrieve a result set with approximately 100 records and scroll through them to gather information. These analysts will need the ability to locate records based on either factory or product ID. They are not responsible for updating any data. You thus need to design the most appropriate cursor strategy for National Retailers.

What should you do?

- A. Create a server-side static cursor.
- B. Create a client-side forward-only cursor.
- C. Create a server-side dynamic cursor.
- D. Use a default result set and do not create a cursor.

Answer: D

Explanation: The Default Result Set caches all records in the result set to the client. This will thus reduce round trips across the network to one and prevent data from being stored in the tempdb.

1. Regularly at the end of each month, reports are generated manually and e-mailed to the Chicago office. Monthly sales and invoice aging information is included in these reports.

Incorrect Answers:

A: If you create a Server-side static cursor it will consume server resources and require a round-trip across the network each time the client fetches data. Furthermore the cursor will then be stored in tempdb.

B: Creating a Client-side forward-only cursor will not allow analysts to scroll through the data the way that they require to. These cursors only support the ability to fetch the next record and not to move to a specific record.

C: A Server-side dynamic cursor does require a round-trip across the network every time a client fetches data.

QUESTION 2

You are busy creating the objects required to support an experiment for Willow Bridge. This experiment is code-named Hyp120H. The lab sample data will be logged to a table named Hyp120Hsample. Historical sample data will be archived to Hyp120HsampleI. The following exhibit displays the query used to update the data in the Hyp120Hsample table:

```
UPDATE Hyp120Hsample
```

```
SET ppM = @ppMillion, temp = @temp
```

```
WHERE SampleID = @sampleID
```

You need to choose the appropriate code to create the query that will archive the data.

What should you do?

A. Archive the record using the following query:

```
INSERT INTO Hyp120HsampleI
```

```
SELECT * FROM deleted
```

B. Archive the record using the following query:

```
INSERT INTO Hyp120HsampleI
```

```
SELECT * FROM inserted
```

C. Use a stored procedure that includes the following:

```
INSERT INTO Hyp120HsampleI
```

```
ppM = @ppMillion, temp = @temp, SampleID = @sampleID
```

D. Use a stored procedure that includes the following:

```
INSERT INTO Hyp120HsampleI
```

```
SELECT * FROM Hyp120Hsample
```

```
WHERE SampleID = @sampleID
```

Answer: A

Explanation: To ensure that the data will always be archived when a record is

updates, you can place a query in the archive in a trigger. A trigger copies the data s it exists before the update to the deleted table in SQL Server 2005. Thus the unmodified data will be located in the deleted table.

1. At present Lab Trial application uses a keyset-driven cursor to fetch a single row at a time. The lab technician takes measurements and updates the row of data with the current sample data. Previous data is archived to a history table during the update.

Incorrect Answers:

B: In SQL Server 20055 the inserted table stores the data that will be inserted when the UPDATE occurs. Thus you will not be archiving previous data; instead you will be inserting the new data into the archive table if you make use of this trigger in the query.

C: Making use of a trigger is a way to ensure that the data is always archived. The stored procedure arguments are not accessible from within a trigger and furthermore, the arguments will contain new data and not the previous version of the data that should be archived.

D: Making use of a trigger is a way to ensure that the data is always archived. The stored procedure arguments are not accessible from within a trigger and furthermore when making use of an AFTER trigger, the data in Hyp120Hsample would already be changed when the trigger is executed.

QUESTION 3

You are designing the database query strategy that will be used to analyze the results of a clinical trial for an oncology drug. The upgraded WB-DB01 will be used by the analysts working on this oncology drug trial. The table that will be used to store the data is displayed in the following exhibit:

Oncology115	
PK	RecordID
	ParticipantID
	DrugRegimenID
	CellCount
	Date

The results must show the total cell count for each participant on each drug regimen. The data must be tabulated with a row for each drug regimen and a column for each participant. You need to design the query to achieve the necessary results whilst ensuring that your solution will be user friendly and easy for future developers.

What should you do?

- A. Create a SELECT statement with the GROUP BY clause.
- B. Create a SELECT CASE statement that populates a table variable.
- C. Create a nested IF ELSE statement that populate a table variable.
- D. Create a SELECT statement with the PIVOT clause.

Answer: D

Explanation: A SELECT statement with the PIVOT clause will allow you to pivot

the values in a specific column of a table to be individual columns in the result and then aggregate values to populate those columns. In this particular case you will pivot the ParticipantID column.

1. The Willow Bridge management wants to upgrade WB-DB01 to SQL Server 2005.
2. A different table is used to store the data for each laboratory experiment. The schema varies somewhat based on the nature of the experiment.
3. Information regarding participants in clinical trials is stored in a separate table for each trial.

Incorrect Answers:

A: A GROUP BY clause will allow one to aggregate the data but not to store it in the desired format. Instead you would have a row for each participant within each drug regiment or you would have a row for each drug regiment for each participant.

B: You will be able to achieve the required result set if you make use of a SELECT CASE statement, but it would require more coding and will prove to be more difficult to read and understand for the future developers.

C: This option will also work, but it will make it very difficult for future developers to read and understand because it is even more difficult than using a SELECT CASE statement.

QUESTION 4

You work as the database developer for Willow Bridge. You are currently designing the error-handling routines. The Willow Bridge lab trial application should be updated to meet the business requirements. You need to select the appropriate code that will add the necessary custom error messages.

What should you do?

A. Execute the following code:

```
EXEC sp_addmessage @msgnum = 70020, @severity = 12,  
@msgtext = N'The value must be between %d and %s.',  
@lang = 'us_english';  
EXEC sp_addmessage @msgnum = 70021, @severity = 12,  
@msgtext = N'L'Translated message%1! und %2!',  
@lang = 'German';
```

B. Execute the following code:

```
EXEC sp_addmessage @msgnum = 70020, @severity = 12,  
@msgtext = N'The value must be between %d and %s.',  
@lang = 'us_english';  
EXEC sp_addmessage @msgnum = 70021, @severity = 12,  
@msgtext = N'L'Translated message%d und %s.',  
@lang = 'German';
```

C. Execute the following code:

```
EXEC sp_addmessage @msgnum = 70020, @severity = 12,  
@msgtext = N'The value must be between %d and %d.',  
@lang = 'us_english';  
EXEC sp_addmessage @msgnum = 70020, @severity = 12,  
@msgtext = N'L'Translated message%1! und %2!';
```

```
@lang = 'German';  
D. Execute the following code:  
EXEC sp_addmessage @msgnum = 70020, @severity = 12,  
@msgtext = N'The value must be between %s and %s.',  
@lang = 'us_english';  
EXEC sp_addmessage @msgnum = 70020, @severity = 12,  
@msgtext = N'L'Translated message%s und %s',  
@lang = 'German';
```

Answer: C

Explanation: In SQL Server 2005 you are allowed to add error messages with the same error number for multiple languages. You identify the language by means of the @lang argument. The %d is a placeholder for a decimal value that will be passed as a parameter to RAISERROR. In the localized message, the %1! and %2! placeholders serve to identify the order in which the parameters should be displayed. In this scenario they are displayed in the same order.

1. The Willow Bridge management also wants to open another research facility in Berlin, Germany.
2. When the research lab is opened in Berlin, a database server running SQL Server 2005 will be installed there. This server will be named WB-DB03. Researchers at both the Chicago and Berlin research facilities will need to be able to query both WB-DB01 and WB-DB03 to share research data.

Incorrect Answers:

A: The localized version of the message should have the same @msgnum parameter setting. The RAISERROR function will raise a single error, regardless of the language. SQL Server will determine the language in which to display the message based on the session language.

B: There is no need to use a different @msgnum value for the localized message. You should also not be making use of %d and %s as placeholders in the localized message.

D: You need to make use of %1! and %2! as placeholders in the localized version of the message and not a placeholder that identifies the type of data. The %s placeholder indicates that a string could be inserted.

QUESTION 5

You work as the database developer for Willow Bridge. You are busy updating the laboratory trial application and need to implement data validation for sample data.

In your solution you must ensure that you meet the business requirements.

What should you do?

- A. Use a check constraint on the lab samples table for each experiment.
- B. Use a unique constraint on the lab samples table for each experiment.
- C. Use an AFTER trigger on the lab samples table for each experiment.
- D. Use an additional middle tier component that validates data prior to sending it to the database server.
- E. Data validation code should be added to the application.

Answer: C

Explanation: An AFTER trigger will allow you to check the values and restore the values from the deleted table if the new values are out of range. Then you can report a custom error message to the application.

1. The lab technician takes measurements and updates the row of data with the current sample data. Previous data is archived to a history table during the update. Lab technicians must be able to verify the changes they made to the sample data when they go back to check their work.
2. One requirement for the Lab Trial application is thus that data entered by technicians must be validated to ensure that the values are within the tolerance range for the experiment. In the event of data falling outside of the tolerance range being entered, an error message should be reported to the technician. Information such as the appropriate data range, and if required additional instructions should be contained in this report.
3. The same application is used to manage samples for all experiments. Lab technicians select the experiment name and the code makes use of a lookup table to discover the name of the table that stores the sample data.

Incorrect Answers:

A: A CHECK constraint will not allow you to send a custom error message to the application.

B: You cannot use a UNIQUE constraint in this scenario because it will not allow you to send a custom error message to the application.

D: Adding a middle-tier component that validates data prior to sending it to the database server will require that you update the middle-tier component every time a new experiment is started. The validation logic is best implemented on each lab sample table at the data services layer.

E: Data validation code should not be added to the application. This would make it very hard to maintain because the data validation logic will be different for every experiment, which means that the code will have to be rewritten every time a new experiment is started.

QUESTION 6

You are designing the database query strategy for Willow Bridge that will be used to analyze sample data. When writing the stored procedure you need to execute a query that will return:

1. The average quantity of each type of antioxidant found in each sample group
2. The average quantity all antioxidants found in each sample group
3. The average quantity of antioxidants found across all samples.
4. The average quantity each type of antioxidant found across all the samples.

What should you do?

A. Use the following code:

```
SELECT SampleGroup, Antioxidant, AVG(NumberFound)
FROM Test22BJ
GROUP BY SampleGroup, Antioxidant
```

WITH CUBE ;

B. Use the following code:

```
SELECT SampleGroup, Antioxidant, NumberFound, AVG(NumberFound)
FROM Test22BJ
GROUP BY Antioxidant, SampleGroup
WITH CUBE ;
```

C. Use the following code:

```
SELECT SampleGroup, Antioxidant, NumberFound,
SUM(NumberFound)/COUNT(NumberFound)
FROM Test22BJ
GROUP BY Antioxidant, SampleGroup
WITH ROLLUP ;
```

D. Use the following code:

```
SELECT SampleGroup, Antioxidant, AVG(NumberFound)
FROM Test22BJ
GROUP BY SampleGroup, Antioxidant
WITH ROLLUP ;
```

Answer: A

Explanation: The CUBE keyword adds rows that contain summary data for each category and for the total to the query results. Using this code will allow you the ability to meet the summary requirements.

Incorrect Answers:

B: This query as it is will result in a syntax error. You cannot include a non-aggregate column that is not listed in the GROUP BY clause in the SELECT list.

C: Using this code will result in the query having a syntax error because you cannot include a non-aggregate column that is not listed in the GROUP BY clause in the SELECT list.

D: The ROLLUP keyword only creates a summary of data by the first category, for each subcategory (second category), and for the whole. This code in a query will thus NOT include the average of each type of antioxidant type found across all samples.

QUESTION 7

You work as the database developer for Willow Bridge. You need to develop the queries that will be used by the researchers who do the follow-up calls for the clinical trials. To this end you decide to create a stored procedure that will retrieve the required participants. You need to select the correct code that will ensure that the query is limited to those who have participated in one study, but not in another. What should you do?

A. Use the following code:

```
SELECT ParticipantID
FROM @ClinicalTestA
EXCEPT
SELECT ParticipantID
```

FROM @ClinicalTestB

B. Use the following code:

```
SELECT ParticipantID
```

```
FROM @ClinicalTestA
```

```
INTERSECT
```

```
SELECT ParticipantID
```

```
FROM @ClinicalTestB
```

C. Use the following code:

```
SELECT a.ParticipantID
```

```
FROM @ClinicalTestA a
```

```
INNER JOIN @ClinicalTestB b
```

```
ON a.ParticipantID <> b.ParticipantID
```

D. Use the following code:

```
SELECT b.ParticipantID
```

```
FROM @ClinicalTestA a
```

```
LEFT JOIN @ClinicalTestB b
```

```
ON a.ParticipantID = b.ParticipantID
```

Answer: A

Explanation: The EXCEPT operator will cause all results from the first table that do not match the records in the second table to be returned. This will thus meet the requirement of generating a list of participants in @ClinicalTestA who did not participate in @ClinicalTestB.

1. Information regarding participants in clinical trials is stored in a separate table for each trial. The Willow Bridge researchers that run clinical trials often need to make follow-up telephone calls to participants. To this end they need to obtain a list of participants, who need a follow-up call; they need to select participants that have participated in a specific study or studies. They sometimes limit the query to those who have participated in one study, but not in another.

Incorrect Answers:

B: The INTERSECT operator will return all rows from the first query that match rows in the second query. This means that all participants who participated in both studies will be returned by this query. This is the opposite of what you want to achieve in this case.

C: An INNER JOIN will return all rows that match the join condition. This means that it will not return only the participants of @ClinicalTestA, but not in @ClinicalTestB.

D: A LEFT JOIN will return all the records in the table on the left of the ON clause and the matching records from the table on the right. This means that all participants who participated in ClinicalTestA and those who participated in ClinicalTestB will be returned. Since there are no columns from @ClinicalTestA in the Select list, actually only the participants who participated in both clinical tests will be returned.

QUESTION 8

You work as the database developer for Willow Bridge. You need to address the issues surrounding the performance of the lab trial application and database when retrieving sample data to be updated. To this end you need to design efficient access

to the SQL Server Service.
What should you do?

- A. The application should be modified to make use of Multiple Active Result Sets (MARS).
- B. The application should be modified to create a dynamic cursor.
- C. The application should be modified to call GetRows to retrieve 25 rows at a time.
- D. The application should be modified to make use of the default result set.
- E. Make use of Microsoft Data Access Components (MDAC) and enable Multiple Active Result Sets (MARS) on a single connection.

Answer: C

Explanation: When one makes use of a block cursor like in changing the application to call GetRows to retrieve 25 rows at a time, the cursor will provide better performance than fetching each row individually. This will result in fewer round trips to the server.

1. At present Lab Trial application uses a keyset-driven cursor to fetch a single row at a time.

Incorrect Answers:

A: MARS is used to allow multiple sets to be opened on a single connection. The lab technicians are working with a single set of results at a time. This will thus make MARS superfluous.

B: Changing the client application to create a dynamic cursor will improve the initial performance when the cursor is opened, but it will not decrease the fetch time.

D: The application is supposed to perform a positioned update. You cannot perform a positioned update with the default result set.

E: MDAC does not support MARS queries and furthermore MARS is not a viable option in this case.

QUESTION 9

You work as the database developer for Willow Bridge. While busy designing queries for the retrieval of data from XML sources, you are writing a script that will generate an XML file to be imported into the collections application. This application assumes XML data includes both elements and attributes.

You need to make a choice of the most appropriate type of Transact-SQL statement to use to retrieve the data from the database.

What should you do?

- A. Make use of the sp_xml_preparedocument Transact-SQL statement.
- B. Make use of the OPENXML Transact-SQL statement.
- C. Make use of the SELECT ... FOR XML AUTO Transact-SQL statement.
- D. Make use of the SELECT ... FOR XML PATH Transact-SQL statement.

Answer: D

Explanation: The FOR XML PATH clause of the SELECT statement will allow you to use XPath to define the structure for the XML data that is returned. This structure can include both elements and attributes.

1. Regularly at the end of each month, reports are generated manually and e-mailed to the Chicago office. Monthly sales and invoice aging information is included in these reports.
2. The Chicago research facility also operates as the head quarters and as such handles debt collection. WB-DB02 supports the sales and accounting operations. All customers with an outstanding balance over 60 days are also sent to the Chicago office. The accounting department makes use of an application that imports XML data.

Incorrect Answers:

A: The sp_xml_preparedocument system stored procedure is used prior to calling OPENXML.

B: The OPENXML statement is used to insert XML data into relational tables. Not to retrieve a resultant set formatted as XML from relational tables.

C: The FOR XML AUTO clause of the SELECT statement will generate a result set of nested elements and as such do not make allowance for a mix of elements and attributes.

QUESTION 10

You work as the database administrator at Willow Bridge. You need to collect information that will be required for the ongoing server management of WB-DB01. The information you need to collect includes hardware performance statistics. You are required to collect the statistics during average use and peak periods. You want to collect the hardware baseline performance statistics. What should you do?

- A. A Windows counter log should be created.
- B. A Windows Alerts should be created.
- C. A Windows trace log should be created.
- D. A SQL Server Profiler trace should be created.

Answer: A

Explanation: You should create a Windows counter log as this lets you specify the hardware resource performance counters that you want included in the baseline performance statistics.

Incorrect Answers:

B: This option should not be used in the scenario because alerts are created to trigger when counters fall above or below a specified threshold value.

C: This option should not be used in the scenario because the trace log is used to collect specific performance information and does not let you specify performance counters to monitor.

D: This option should not be used in the scenario as the counter itself does not collect hardware performance statistics.

QUESTION 11

Willow Bridge requires a database that will support the e-commerce application.

You need to find a way to provide the e-commerce application with information about the existing inventory. The database must determine the nearest research facility of sales office that has the products ordered in stock and send the order to that specific office to have it fulfilled.

You thus need to determine how the inventory information should be made available to the e-commerce application while providing the best possible performance when placing the order.

What should you do?

- A. Query the database server at each office using OPENROWSET from a stored procedure run on WB-DB02.
- B. At each office implement an HTTP endpoint on the database server and expose a Web method that checks inventory.
- C. Replicate the data from the database servers at the different factories to the WB-DB02 server using merge replication.
- D. Use a linked server for the database server at each office at WB-DB02.

Answer: C

Explanation: Using merge replication to replicate all data from the factories to WB-DB02 will make the data available locally and improve query performance when an order is placed. This is the best type of replication to use in this case because there could be conflicts by orders placed over the Web and order placed by the Sales staff at each office. Also in a case where the different sites are not well-connected making use of merge replication makes sense. And in this case the sales offices are connected to the Chicago office by demand-dial links.

1. The Willow Bridge management wants to expand business by selling their products and drugs on the company Web site.
2. WB-DB02 will host a database to support the e-commerce application that will be for the benefit of medical practitioners world-wide. The e-commerce application will make use of Simple Object Access Protocol (SOAP) to retrieve product information and submit orders.

Incorrect Answers:

A: Making use of OPENROWSET to query each database server at each factory from a stored procedure on WB-DB02 will result in one or more remote queries to be sent across the network for each product ordered.

B: Making use of an HTTP endpoint will result in more overhead than is required. There is already a demand-dial connection between the Chicago office and the different offices. Thus there is no firewall that restricts incoming protocol to only HTTP. Making use of a SOAP request will result in even more overhead than a remote query.

D: Making use of a linked server for the database server at each site and performing a remote query is not the best solution because each product that is ordered will cause at least one, and possible more, remote queries to be sent across the network.

Topic 8, TestLabs Inc., Scenario

BACKGROUND:

Company Overview

TestLabs Inc., is the leading manufacturer which produces specially brewed high octane racing fuels and prototype high performance accessories for several companies and customers in the racing industries. TestLabs Inc., mission is to provide the best racing fuels which meet precise manufacturing standards. TestLabs Inc., distributes the racing fuels all across the world. TestLabs Inc., has most of their customers located in the United States.

TestLabs Inc., makes use of fully automated manufacturing procedures which is used to create the prototype high performance accessories and high octane racing fuels. TestLabs Inc., has all their manufacturing equipment controlled by proprietary systems which can be programmed manually through a control interface program or automatically by uploading a process file to the control computer. TestLabs Inc., has the capability to switch production lines between components by loading a different XML process file.

TestLabs Inc., has four manufacturing lines which are responsible for operating at any time, each making a single component. TestLabs Inc., has recently decided to stop the production line for an hour each day for inspection, cleaning, and calibration. TestLabs Inc., makes use of a rolling schedule for taking down the production lines meaning no more than one production line is down for scheduled maintenance at a time. TestLabs Inc., shuts down each production line at the same time each day. TestLabs Inc., takes one day in a month to shutdown all production lines for major maintenance activities.

TestLabs Inc., has the manufacturing line running 24 x 7 except for the shut down.

TestLabs Inc., has recently decided to have all manufacturing tolerances for each component specified by the manufacturing contract. TestLabs Inc., use a unique serial number that includes an embedded batch number for identifying each component.

TestLabs Inc., decided to create the serial number as each component is started and applied during the manufacturing process. TestLabs Inc., has the batch numbers updated daily based on the year and Julian date. All serial numbers used by TestLabs Inc., conforms to a standard format including a series of numbers and letters.

TestLabs Inc., has decided to sell their prototype high performance accessories and high octane racing fuels only to contracting customer. The customers of TestLabs Inc., want to have the inventory on hand to minimize for the accessories to be scheduled for just-in-time delivery. TestLabs Inc., requires that the customers provide at least 48 hours warning when requesting a change in delivery schedule.

TestLabs Inc., has all the raw materials that they receive based on just-in-time delivery schedules. TestLabs Inc., also generates to the vendor purchasing the accessories 24-hours before expected delivery. TestLabs Inc., additionally generates daily internal-mail to the warehouse management listing orders expected for receipt that day.

TestLabs Inc., has recently decided to take the next step in the company evolution by automating most of its quality control (QC) procedures. TestLabs Inc's currently makes use of a proprietary network-based collection program named BillQuality with a limited reporting component. TestLabs Inc., knows that the proprietary network-based collection programs reporting component has the capability to generate a raw dump of sensor data to a comma-delimited file. TestLabs Inc., makes use of the latest sensors which are used to measure the exact weigh and diameter of the components verifying they are within manufacturing tolerance. TestLabs Inc., has the information reported to a specific IP

address.

Technical Environment

TestLabs Inc., currently makes use of a computer running Microsoft Windows Server 2003 Standard Edition for storing the XML process files. TestLabs Inc., has the Microsoft Windows Server 2003 Standard Edition computer running an XML Web service to support the process control computers.

TestLabs Inc., has loaded the Process files into the control systems as XML files containing the manufacturing steps and specifications. TestLabs Inc., has configured the process computer to submit an identifying code to the Web service when reset. TestLabs Inc., has configured the Web service to return the XML process file contains a seed value for the component serial number which is used as the starting serial number for the day's batch.

TestLabs Inc., was deployed originally with only three Microsoft SQL Server 2005 Enterprise Edition computers running Windows Server 2003 Enterprise Edition.

TestLabs Inc., has ensured that their deployment plan allows for more instances of SQL to be installed in the future if required to meet business or technical requirements.

TestLabs Inc., will deploy more servers in the future if it is required. TestLabs Inc., currently makes use of the following database servers:

1. TESTLABS-SR01 which maintains financial accounting records and contract information
2. TESTLABS-SR02 which maintain all personnel and employee-related records
3. TESTLABS-SR03 is used to handle the manufacturing line production, inventory, shipping records and XML process files
4. TESTLABS-SR04 which is the quality control tracking server

TestLabs Inc., requires having both the current and historic process files to be maintained on TESTLABS-SR03 which is configured with three databases named process for current process files, HistProc for the historic process files and Inventory for all the other records maintained. TestLabs Inc., will make use of the SQL Server named TESTLABS-SR04 to collect and monitor quality control information.

BUSINESS REQUIREMENTS:

TestLabs Inc., knows that this type of business has several time-sensitive reporting requirements. TestLabs Inc., requires having the Managers have daily reports on current inventory, daily output, and projected shipments. TestLabs Inc., wants to have all the reports e-mailed automatically to managers. The Managers have recently requested that they be able to recall previous day's reports. The Managers of TestLabs Inc., additionally requires the daily QC summary reports describing any variations from optimum tolerance and when they occurred. TestLabs Inc., considers it very critical that reporting should not be allowed to interfere with operational requirements.

1. TestLabs Inc., has identified other reporting requirements which include general production and sales reports, business-related expenses due to problems such as down time or raw materials delays, and general accounting reports.

1. TestLabs Inc., maintains all their purchase and sales records in US dollars. The Orders are generated automatically at TestLabs Inc., when the components are shipped and the sale posted at that time. TestLabs Inc., also generates invoices for customers of other countries in their local currency.

1. TestLabs Inc., requires having the customers receive inventory and shipment schedule

information by e-mail only. TestLabs Inc., wants to have their customers able to request shipment records for components already shipped throughout the life of the manufacturing contract. TestLabs Inc., wants to have the Managers to be able to request projected shipments for the next 30 days.

1. TestLabs Inc., wants to have their maintenance personnel generate detailed ad hoc QC tracking reports in chart and table report formats. The TestLabs Inc., maintenance staff should be able to filter the report by manufacturing line, specific QC sensors, and time of day range. TestLabs Inc., requires having the maintenance team to be able to generate these reports at any time accurate to the latest five minutes.

TECHNICAL REQUIREMENTS:

TestLabs Inc., makes use of the QC sensors to provide real-time data updated for each component manufactured. TestLabs Inc., has configured each component with a separate QC tracking table. TestLabs Inc., wants to have a new tracking table created before the start of each new batch run. TestLabs Inc., has the date encoded in the table name. A list of data to be tracked with each entry is shown in the table below:

Data	Description
Line	The line is identified by the process control computer serial number.
Sensor	The numeric identifier of the sensor providing the data.
Value	The numeric value. Potential range varies by sensor.
Time	The time recorded to millisecond. The Value is supplied by the sensor.
ComponentSerial	The component serial number. This is determined by the computer by comparing the serial number of the last component for which data was reported by the same sensor.

TestLabs Inc., wants to have each sensor return a single value for each component.

TestLabs Inc., considers it to be imperative that all QC data is collected and reported accurately.

1. TestLabs Inc., does not want the data to be changed manually, and any attempts to change the data must be tracked and reported.

1. TestLabs Inc., wants to have the QC records monitored on a real-time basis as any time QC values exceeding the tolerances specified on the contract are detected the manufacturing line must be shut down immediately.

1. TestLabs Inc., wants to have this initiated automatically by sending a coded message through the network. TestLabs Inc., will configure the manufacturing control system to monitor a specified TCP port for the message. TestLabs Inc., wants to be able to configure the port value individually for each control system through its internal maintenance application.

1. The initial testing of TestLabs Inc., has shown that QC data collection and monitoring will be resource intensive, with disk resources a special area of concern. TestLabs Inc., wants to have the extraneous activity on TESTLABS-SR04 kept to a minimum.

1. TestLabs Inc., will configure TESTLABS-SR04 with three read-write filegroups Write1 (the primary filegroup), Write2, and Back1. TestLabs Inc., will configure each read-write filegroup has its own RAID 5 disk set and consists of one data file.

1. TestLabs Inc., will keep the transaction log is kept on its own RAID 1 disk set whilst TESTLABS-SR04 data is maintained online for a minimum of one year ensuring that older TESTLABS-SR04 data is archived to offline storage, but must be retrievable for

reference.

1. TestLabs Inc., has configured all database servers for Windows authentication only. The TestLabs Inc., company security policy specifies that network server access require Active Directory user accounts.

Topic 8, TestLabs Inc (11 Questions)

QUESTION 1

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start designing efficient access to a SQL Server Service. The strategy you are designing will make use of the appropriate data access technologies. You have recently started designing access components to service client orders.

You are required to modify the current order processing procedures to meet the requirements of some of the larger TestLabs Inc., customers. TestLabs Inc., has recently allowed some clients to make use of a business-to-business application that generates orders as well-formed XML documents. You have recently received additional instruction to develop an application to manage receiving these documents whilst validating the schema and returning a response to the customer. You want to launch a stored procedure when the order is received to parse and process the order.

You are finally required to ensure that the orders are processed quickly as possible whilst ensuring that order processing does not require secure authentication.

TestLabs Inc., requires having the processing occur in proper order whilst continuing automatically after several restarts or network failure without loss of data whilst keeping components required to a minimum. You should choose what to use for the creation of the application.

What should you do?

- A. You should create a SQL Server Integration Services packet-based application.
- B. You should create an application based on SQL Server Agent jobs that execute hourly to check for orders.
- C. You should create a native XML Web service and HTTP endpoint.
- D. You should create a Service Broker application.

Answer: D

Explanation: You should always remember in the scenario that making use of a Service Broker application you ensure that you can prevent data loss in case of server reset or network failure as well as other serious problems.

1. TestLabs Inc., has identified other reporting requirements which include general production and sales reports, business-related expenses due to problems such as down time or raw materials delays, and general accounting reports.

Incorrect Answers:

A: You should not consider making use of SSIS in the scenario because SSIS is intended primarily for usage with data load and transfer applications.

B: You should not consider making use of SQL Server Agent jobs in the scenario even though it could be used because one of your requirements is to have the processing complete as quickly as possible.

C: You should not consider taking the choices presented by these options in the scenario as you would effectively not be including the functionality required to prevent data loss after server reset or network failure events.

QUESTION 2

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start a design strategy for error-handling routines. The solution strategy that you are designing will use code that validates input data and permissions. You have recently started designing the data access requirements for the various applications that are used by TestLabs Inc.

TestLabs Inc., knows that during the initial phase of contracting to produce a new component that specifications are submitted as well formed XML document.

TestLabs Inc., additionally knows that Submissions requirements state that the XML document should include an inline schema for which you provide four standard schemas for this purpose.

TestLabs Inc., will have the customers use one of the schemas and create the XML document based on that schema. TestLabs Inc., will have the specifications submitted through an XML web service hosted on a Web server running Microsoft Windows Server 2003. TestLabs Inc., wants to have the XML document validated against the included schema before it is loaded into a table named ComponentForReview which will hold the summary customer information and preliminary contract information.

You should remember that the document and schema are stored in a nvarchar(max) type column. You are required to ensure that the XML document is validated before loaded to a table whilst the processing load on the database servers hosting the ComponentForReview table is minimized.

What should you do?

A. You should have the Web service submit the document to the database server and use a table with a typed XML column for each schema to validate the document at the database server before loading the document into the ComponentForReview table.

B. You should parse the document at the Web service to remove the schema and do a string compare to compare the inline schema against the standard schema to see if one is a match.

C. You should parse the document at the Web service to remove the schema and pass the remaining document to the database server and use sp_xml_preparedocument to read

the XML document into memory.

D. You should make use of ADO.NET XmlValidatingReader to validate the document before it is submitted from the Web service to the database server.

Answer: D

Explanation:

You should remember in the scenario that we are required to validate the XML document schema to one of the four standard schema's that we provided and making use of the ADO.NET XmlValidatingReader which is designed to be able to parse the document, read the schema and validate the documents based on that schema you meet the required requirements.

1. TestLabs Inc., has loaded the Process files into the control systems as XML files containing the manufacturing steps and specifications. TestLabs Inc., has configured the process computer to submit an identifying code to the Web service when reset. TestLabs Inc., has configured the Web service to return the XML process file contains a seed value for the component serial number which is used as the starting serial number for the day's batch.

2. TestLabs Inc., currently makes use of a computer running Microsoft Windows Server 2003 Standard Edition for storing the XML process files. TestLabs Inc., has the Microsoft Windows Server 2003 Standard Edition computer running an XML Web service to support the process control computers.

Incorrect Answers:

A: You should not consider taking this action in the scenario because the document still contains the inline schema meaning validating will fail and testing against each column would work if the schema was removed first.

B: You should not consider taking the actions in this option in the scenario as you would only be checking the schema for errors and not validating the content document as required.

C: You should not consider taking the action ion this option in the scenario because this would effectively read the document into memory without validation to prepare for data extraction using the OPENXML function.

QUESTION 3

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start designing efficient access to a SQL Server Service. You are in the process of designing a cursor strategy for a data access component. You have recently worked with a database application programmer to create an application to read and process the QC records.

TestLabs Inc., has instructed the programmer to create the application that uses Open Database Connectivity (ODBC). TestLabs Inc., wants to have the application use a server side cursor for processing the database, the cursor will take a snapshot

of the data when it is retrieved and must be insensitive to data changes. TestLabs Inc., wants to have the application retrieve data from the cursor as quickly as possible. You are required to select the appropriate application programming interface (API) cursor type.

What should you do?

- A. You should make use of a static cursor.
- B. You should make use of a forward-only cursor.
- C. You should make use of a keyset-driven cursor.
- D. You should make use of a dynamic cursor.

Answer: A

Explanation: You should remember that there are four standard API cursor types supported in SQL Server: static, dynamic, keyset-driven and forward-only cursors. By making use of the static cursor you ensure that you provide high-speed access to data with minimal resource requirements as static cursors are read-only and insensitive to changes in the data.

1. The initial testing of TestLabs Inc., has shown that QC data collection and monitoring will be resource intensive, with disk resources a special area of concern. TestLabs Inc., wants to have the extraneous activity on TESTLABS-SR04 kept to a minimum.
2. TestLabs Inc., has identified other reporting requirements which include general production and sales reports, business-related expenses due to problems such as down time or raw materials delays, and general accounting reports.

Incorrect Answers:

- B: You should not consider taking this action in the scenario in context of the SQL server because ODBC recognizes this as a unique cursor but SQL Server does not.
- C: You should not consider making use of keyset-driven cursor in the scenario because the membership and orders is fixed but not column values as changes made by other applications will be visible as the application scrolls through the cursor.
- D: You should not consider the usage of dynamic cursors in the scenario because the dynamic cursor is sensitive to data changes and would not be the correct solution as required in the scenario.

QUESTION 4

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start designing a database query strategy. You are currently in the process of designing queries for retrieving data from XML sources. TestLabs Inc., knows that during the initial phase of contracting to produce a new component that specifications are submitted as well formed XML document. TestLabs Inc., additionally knows that Submissions requirements state that the XML document should include an inline schema for which you provide four standard schemas for this purpose.

TestLabs Inc., will have the customers use one of the schemas and create the XML document based on that schema. TestLabs Inc., will have the specifications submitted through an XML web service hosted on a Web server running Microsoft Windows Server 2003. TestLabs Inc., wants to have the data validated, parsed and the schema removed and written to a temporary holding table where it is stored in an XML type column.

The TestLabs Inc., product design personnel require portions of the document for review and analysis by providing the fragments needed based on the schema used to create the source document. You are required to design a query that can retrieve the XML fragment needed from the holding table.

What should you do?

- A. You should make use of the XML type exist () method.
- B. You should make use of OPENXML.
- C. You should make use of the XML type value () method.
- D. You should make use of the XML type nodes () method.
- E. You should make use of the XML type query () method.

Answer: E

Explanation: You should remember in the scenario that we are required to retrieve a specific section or element from an XML document and not data, by making use of the query () method you ensure that you meet the specific requirements from the scenario.

1. TestLabs Inc., currently makes use of a computer running Microsoft Windows Server 2003 Standard Edition for storing the XML process files. TestLabs Inc., has the Microsoft Windows Server 2003 Standard Edition computer running an XML Web service to support the process control computers.

2. TestLabs Inc., has loaded the Process files into the control systems as XML files containing the manufacturing steps and specifications. TestLabs Inc., has configured the process computer to submit an identifying code to the Web service when reset. TestLabs Inc., has configured the Web service to return the XML process file contains a seed value for the component serial number which is used as the starting serial number for the day's batch.

Incorrect Answers:

A: You should not consider taking the action in this option because making use of this method returns a true or false method based on evaluating a XQuery string to determine whether or not the node exists and this is not required in the scenario.

B, D: You should not consider making use of these options in the scenario because these options are both used to create a rowset when retrieving data in a relational format from an XML instance source and we are required to retrieve XML format data not relational data.

C: You should not consider making use of the value () XML type method in the scenario because you are not required to retrieve a document element value, you are required to retrieve the document fragment.

QUESTION 5

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to performance tune a database and database application. The solution you design will be used to resolve performance problems. TestLabs Inc., has recently informed you that the QC database is experiencing intermittent problems with lost updates during peak use times. TestLabs Inc., has additionally informed you that the problems occur during times when both read and write operations are required.

You have recently decided to start monitoring server hardware and discover that some of the counters are outside the values originally collected as the baseline when the QC database was deployed. The database server in question has four hard disks, one processor and two gigabytes of RAM. Disk 0 contains the operating system and SQL Server program files, Disk 1 contains active QC data, Disk 2 contains the archived QC data and Disk 3 contains the transaction log file and non-clustered indexes for the QC tables.

You have discovered that Disks 0, 2 and 3 are at below the baseline values and the following values have changed significantly from the baseline values:

1. SQL Server: Buffer Manager: Buffer Cache Hit Ratio has increased from 70% to over 90%.
2. PhysicalDisk: % Disk Time has increased from over 60 % to over 80% on disk 1.
3. PhysicalDisk: Avg. Disk Queue Length has increased from 0.4 to 1.1 on disk 1.
4. Processor: %Processor Time has increased from 60 % to over 90%.

You have recently collected a workload file during a peak period and use it with the Database Engine Tuning Advisor but the tuning advisor does not recommend any index changes. You are required to improve the system performance so that the database does not lose sensor updates.

What should you do?

- A. You should install a second processor.
- B. You should increase the amount of memory allocated to SQL Server.
- C. You should replace disk 1 with a faster hard disk.
- D. You should install additional memory.
- E. You should reduce the number of daily transaction log backups.

Answer: A

Explanation: You should remember in the scenario that we are presented with information that indicates that we should consider installing an additional processor to the system as values over 90% are considered excessive in this particular counter and indicates a potential bottleneck.

1. TestLabs Inc., wants to have the QC records monitored on a real-time basis as any time QC values exceeding the tolerances specified on the contract are detected the manufacturing line must be shut down immediately.

2. TestLabs Inc., wants to have each sensor return a single value for each component. TestLabs Inc., considers it to be imperative that all QC data is collected and reported accurately.
3. The initial testing of TestLabs Inc., has shown that QC data collection and monitoring will be resource intensive, with disk resources a special area of concern. TestLabs Inc., wants to have the extraneous activity on TESTLABS-SR04 kept to a minimum.

Incorrect Answers:

B, D: You should not consider taking these actions in the scenario because taking these actions are not required in the scenario as the Buffer Hit Cache Ratio increasing is a positive development and the information received indicates the need for an additional processor.

C: You should not consider taking this action yet in the scenario because the current values indicate that there is no concern of a potential bottle neck yet.

E: You should not consider reducing the amount of daily transaction logs in the scenario as the information we received indicates that the server only needs an additional server and disk 3 is still performing with values similar to those collected as baseline values.

QUESTION 6

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start designing a transaction strategy. The solution you are designing will be used to manage concurrency by selecting the appropriate transaction isolation levels. You have recently started designing a stored procedure to collect data from the QC server for creating the daily summary reports. You are required to collect data based on point-in-time QC information as of midnight each day. You decided to make use of cursors in the stored procedure and processing the data in context of a transaction to allow rollback in case of processing errors.

TestLabs Inc., will make use of the stored procedure to create a daily summary table that will contain demoralized data. TestLabs Inc., wants to have this be the only procedure that writes whilst the stored procedure's impact on normal operation on the QC database is minimized. You are required to identify the transaction isolation level that meets the requirements of the stored procedure whilst ensuring the isolation has repeatable reads and consistency at the transaction level. TestLabs Inc., has informed you that database-level isolation level configuration settings are at the SQL Server default settings.

What should you do?

- A. You should make use of Snapshot Isolation.
- B. You should make use if Read Committed Isolation Using Row Versioning.
- C. You should make used of Repeatable Read Isolation.
- D. You should make use of Serializable Isolation.

Answer: A

Explanation: You should remember in the scenario that we are required to use the appropriate isolation level and by making use of Snapshot Isolation you ensure that a snapshot of the database is taken in point-in-time when the transaction starts as required and remain consistent throughout.

1. TestLabs Inc., wants to have each sensor return a single value for each component. TestLabs Inc., considers it to be imperative that all QC data is collected and reported accurately.
2. The initial testing of TestLabs Inc., has shown that QC data collection and monitoring will be resource intensive, with disk resources a special area of concern. TestLabs Inc., wants to have the extraneous activity on TESTLABS-SR04 kept to a minimum.

Incorrect Answers:

B: You should not consider making use of the Read Committed Isolation Using Row Versioning in the scenario because this option is used and intended statement level rather than transaction level consistency.

C, D: You should not consider making use of Serializable Isolation or Repeatable Read Isolation in the scenario because both of these options are more locking requirements than either Snapshot Isolation or Read committed Isolation.

QUESTION 7

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to performance tune a database and database application. The solution you design will be used to optimize indexing strategies. You are in the process of designing the indexing strategy for the tables used to track QC data. TestLabs Inc., has informed you that the queries run usually retrieve data and time order. TestLabs Inc., additionally makes use of queries that run to retrieve data by sensor. You are required to create indexes for the QC tracking tables that:

1. Have minimal disk space requirements
2. Optimize data retrieval
3. Have minimal impact on data write operations

You are required to fulfill the primary requirement stating the solution should have minimal impact in write operations.

What should you do?

- A. You should create a clustered index on each QC tracking table and specify the Time and Sensor columns as key columns in that order.
- B. You should create a clustered index on each QC tracking table and specify the Sensor and Time columns as key columns in that order.
- C. You should create a clustered index on each QC tracking table and specify the Time column as the index key column and create a nonclustered index on each QC tracking table with the sensor column as the key column.
- D. You should create a nonclustered index on each QC tracking table and specify the

Time and Sensor columns as key columns in that order.

E. You should create a nonclustered index on each QC tracking table and specify the Sensor and Time columns as key columns in that order.

Answer: A

Explanation: You should remember in the scenario that taking the actions in the answer will ensure that the data will be organized based on Time column values as most of the queries return the data in this order and is the most efficient design.

1. TestLabs Inc., wants to have the QC records monitored on a real-time basis as any time QC values exceeding the tolerances specified on the contract are detected the manufacturing line must be shut down immediately.

2. TestLabs Inc., wants to have each sensor return a single value for each component. TestLabs Inc., considers it to be imperative that all QC data is collected and reported accurately.

3. The initial testing of TestLabs Inc., has shown that QC data collection and monitoring will be resource intensive, with disk resources a special area of concern. TestLabs Inc., wants to have the extraneous activity on TESTLABS-SR04 kept to a minimum.

Incorrect Answers:

B, C: You should not consider taking this action in the scenario because this is the most efficient index as most of the queries return data based on date and time but the solution organizes data based on the Sensor column values.

D, E:

You should not consider taking the options mentioned in these answers in the scenario because a nonclustered index required a separate structure which requires additional write operations impacting the performance and cannot be used for such reason even with a clustered index.

QUESTION 8

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start designing error-handling routines. You have recently started designing user-defined messages to communicate application events. You are in the process of using ADO.NET to access non-current data (historic data) from the QC database. You have decided to make use of classes within the System.Data.SqlClient namespace to manage connections, pass commands and retrieve data.

TestLabs Inc., knows that the application was just done developed and is ready for extensive beta testing. TestLabs Inc., has permitted you to use a copy of the database on the development server for testing. You have recently discovered that when errors occur during data access that the information is not collected. You are required to collect and document any errors as this is done as part of the test process and facilitates troubleshooting after the application is released into production.

What should you do?

- A. You should make use of the SqlError object to collect error information and write the information collected to the Windows Application log.
- B. You should try enclosing a TRY CATCH structure in any command strings passed to the database server through SqlCommand object.
- C. You should use SqlCommandBuilder to build all commands used for data retrieval by application.
- D. You should make use of a Try...Catch structure in the client application to retrieve values for ERROR_NUMBER(), and ERROR_STATE().

Answer: A

Explanation: You should remember in the scenario that the SqlError object is used to allow you to collect information relating to a warning or error returned by SQL Server. Making use of a .NET Framework application you would be capable of writing the information collected to the Windows Application log.

1. TestLabs Inc., requires having both the current and historic process files to be maintained on TESTLABS-SR03 which is configured with three databases named process for current process files, HistProc for the historic process files and Inventory for all the other records maintained. TestLabs Inc., will make use of the SQL Server named TESTLABS-SR04 to collect and monitor quality control information.

Incorrect Answers:

B: You should not consider taking this action in the scenario even though you could use this method the answer does not provide a way to retrieve error information.

C: You should not consider making use of the SqlCommandBuilder in the scenario because the functionality required is not supported by SqlCommandBuilder as the command builder is used to build single-table commands that reconcile changes made to DataSet objects with the associated SQL Server database.

D: You should not consider making use of the Try...Catch structure in the scenario because it is not possible to use the clause in context of a .NET Framework application it is only used in context to try and catch structures running as SQL Server batches, procedures or functions.

QUESTION 9

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start performance tuning a database and database application. The solution you are designing will be used to resolve the performance problems. You are in the process of evaluating hardware platforms to find a replacement server for the QC Server. You have recently acquired three servers for evaluation and you restore a recent copy of the QC database to each of the servers.

You are required to compare the server's performance with the servers operating under the same load whilst ensuring you directly compare performance counter

values from all of the servers. You should accomplish the evaluation task quickly as possible.

What should you do? (Choose TWO.)

- A. You should make use of SQL Server Profiler to capture a trace on TESTLABS-SR04 using the TSQL_Replay template and copy the trace to each evaluation server.
- B. You should run Windows System Monitor on a fourth computer to view the performance counters on each of the evaluation servers.
- C. You should make use of the Windows System Monitor on one of the evaluation servers and include the counters from each of the other evaluation servers.
- D. You should make use of SQL Server Profiler to capture a trace on TESTLABS-SR04 using the Tuning Template and copy the trace to all the evaluation servers.
- E. You should make use of the Database Engine Tuning Advisor to evaluate the trace data
- F. You should capture a counter log on each of the evaluation servers and compare the log results.

Answer: A, B

Explanation: You should make use of the SQL Server Profiler in the scenario to capture a trace on TESTLABS-SR04 with the TSQL_Replay template to replay the trace on the evaluation servers operating under the same load. By making use of Windows System Monitor on a fourth computer you ensure you directly compare the counters in an easy-to-read chart with real-time values.

1. TestLabs Inc., wants to have each sensor return a single value for each component. TestLabs Inc., considers it to be imperative that all QC data is collected and reported accurately.
2. The initial testing of TestLabs Inc., has shown that QC data collection and monitoring will be resource intensive, with disk resources a special area of concern. TestLabs Inc., wants to have the extraneous activity on TESTLABS-SR04 kept to a minimum.

Incorrect Answers:

- C: You should not consider taking this action in the scenario because the result would not be valid as the Windows System Monitor is additional loads even though it is not resource intensive.
- D: You should not consider making use of this template in the scenario because this template is used to create a workload file for evaluation by the Database Engine Tuning Advisor.
- E: You should not consider taking the action mentioned in the answer in the scenario because the only trace capture is from TESTLABS-SR04 and is not related to the evaluation servers.
- F: You should not consider capturing a counter log on each of the evaluation servers in the scenario because this is not the most efficient way to go about comparing the performance counter values.

QUESTION 10

You work as the network database administrator at TestLabs Inc. The TestLabs

Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start performance tuning a database and database application. The solution you are designing will be used to resolve the performance problems. You have recently received results of the QC tracking which you carefully monitor. TestLabs Inc., has recently reminded you that the company reputation for quality relies on the accuracy and completeness of the QC tracking data.

You have recently discovered that intermittent problems with failed updates when all four lines are running at full capacity. You decided to launch Windows Performance on a separate computer and discover that the PhysicalDisk: % Disk Time occasionally exceeds 90% with all lines running. You are required to prevent lost updates whilst remembering the budget allows you to make hardware changes as necessary to improve QC tracking performance without compromising data protection.

What should you do?

- A. You should replace the RAID 5 disk sets to RAID 0.
- B. You should create a second data file for write1 on the disk set with write2 and create a second data file for write2 on the disk set with write1.
- C. You should install a separate SQL Server 2005 instance and assign a RAID 5 disks set to each instance and finally track odd-number lines from the one instance and even numbered lines from the other instance.
- D. You should add two additional RAID 5 disk sets and create a new filegroup on each. You should then create each lines tracking table on a separate disk set.

Answer: D

Explanation: You should remember in the scenario that making use of two additional RAID 5 disk sets when creating the tracking tables you should create each line's tracking table on a separate disk set ensuring there is no contention when all four lines are running.

1. The initial testing of TestLabs Inc., has shown that QC data collection and monitoring will be resource intensive, with disk resources a special area of concern. TestLabs Inc., wants to have the extraneous activity on TESTLABS-SR04 kept to a minimum.
2. TestLabs Inc., will configure TESTLABS-SR04 with three read-write filegroups Write1 (the primary filegroup), Write2, and Back1. TestLabs Inc., will configure each read-write filegroup has its own RAID 5 disk set and consists of one data file.
3. TestLabs Inc., will keep the transaction log is kept on its own RAID 1 disk set whilst TESTLABS-SR04 data is maintained online for a minimum of one year ensuring that older TESTLABS-SR04 data is archived to offline storage, but must be retrievable for reference.

Incorrect Answers:

A: You should not consider making use of a RAID 0 disk set in the scenario even though it has better write performance you would lose the fault tolerance.

B: You should not consider taking the actions mentioned in this answer in the scenario because you would simply be moving the disk contention not reducing the disk contention.

C: You should not consider making use of an additional instance of SQL Server in the scenario as you would only be increasing the resource requirements which are what you should be avoiding.

QUESTION 11

You work as the network database administrator at TestLabs Inc. The TestLabs Inc., network consists of a single Active Directory domain named TestLabs Inc. All servers on the TestLabs Inc., network run Windows Server 2003 and all client computers run Windows XP Professional.

You have recently received instruction from the TestLabs.com network CIO to start performance tuning a database and database application. The solution you are designing will be used to optimize and tune queries for performance. You are in the process of optimizing a query that accesses tables in the QC database. The query used extracts values for specific sensors from the QC tracking tables. TestLabs Inc., knows that each query returns the values for one sensor in one tracking table. The values are returned in the order in which they were generated.

TestLabs Inc., additionally knows that the query optimizer consistently chooses less than optimum execution plan for the query. All the tracking tables used the same index structure and have the following indexes:

1. Ix_forPri - clustered index based on Time, Sensor and line.
2. Ix_byComp - nonclustered index based on ComponentSerial, Line and Time.

You are required to ensure that the query performance is optimized for retrieving the information.

What should you do?

- A. You should specify the INDEX(1) table hint when running the query.
- B. You should specify the INDEX(ix_byComp) table hint when running the query.
- C. You should recalculate table statistics and specify the TABLOCKX hint when running the query.
- D. You should specify the INDEX(0) table hint when running the query.
- E. You should specify the NOEXPAND and INDEX(ix_byComp) hints when running the query.

Answer: A

Explanation: You should remember that specifying the INDEX(1) table hint will force the query to use a clustered index seek to run the query as you are retrieving data from a specific sensor only and the data is physically sorted in time order which is the best solution.

1. TestLabs Inc., wants to have each sensor return a single value for each component. TestLabs Inc., considers it to be imperative that all QC data is collected and reported accurately.
2. TestLabs Inc., makes use of the QC sensors to provide real-time data updated for each

component manufactured. TestLabs Inc., has configured each component with a separate QC tracking table. TestLabs Inc., wants to have a new tracking table created before the start of each new batch run.

3. The initial testing of TestLabs Inc., has shown that QC data collection and monitoring will be resource intensive, with disk resources a special area of concern. TestLabs Inc., wants to have the extraneous activity on TESTLABS-SR04 kept to a minimum.

Incorrect Answers:

B: You should not consider making use of the INDEX(ix_byComp) hint used does not include the Sensor column as a key column and would not help retrieving the information.

C: You should not consider taking this action in the scenario because the TABLOCKX hint is used to force an exclusive lock on the table allowing no access to the table for the duration of the query.

D: You should not consider making use of the hint INDEX(0) in the scenario because this forces a clustered index scan when the query runs a clustered index seek is more appropriate in the scenario.

E: You should not consider making use of the NOEXPAND and INDEX(ix_byComp) hints when running the query in the scenario because you would generate a syntax error as the NOEXPAND option can only be used with indexed queries.

Topic 9, Certkiller .com, Scenario

OVERVIEW:

Company Overview

Certkiller .com provides self-learning course material for certification exams in the information and technology (IT) field as an online subscription. IT students subscribe to courses material on a particular certification exam.

When the objectives tested in the certification exam changes, or new course material is made available, Certkiller .com notifies the subscribers who have subscribed to that exam. Certkiller .com offers different subscription periods with different subscription prices.

Planned Changes

In the existing Web application, subscribers can receive only view course material online. Certkiller .com plans to allow subscribers to download a PDF version of the course material. These PDF files are generated from test that is stored in the database. At any given time, thousands of subscribers might download the same PDF file. These concurrent downloads must consume as little resources as possible from the Web server and the database servers.

Prospective subscribers must be able to search for information about exams by vendor based on the type of certification, minimum requirements, and other features. Prospective subscribers must be able to enter a word or phrase and SQL Server should return a match of synonyms and inflective forms of that word.

Prospective subscribers are provided with quotations for new subscriptions. All quotation data is stored in a database named CK_Quotes. The CK_Quotes database includes a SubscriberInfo and a QuotePrices table.

The prices quoted to prospective subscribers are quoted in their local currency.

Certkiller .com wants to implement a process that will run every hour to mark all price quotes that are older than 24 hours as invalid. This process must run on a per-subscriber basis with all price quotes for one subscriber being processed and then all price quotes for the next subscriber being processed, etc.

EXISTING ENVIRONMENT:

Existing Application Environment

The course material is made to subscribers as a Web application that runs on a Web server named Certkiller -SR04. The Web application retrieves the course material for each subscriber from a SQL Server 2005 database named CK_Courses that runs on a three database servers named Certkiller -DB01 and Certkiller -DB02.

Certkiller -DB01 and Certkiller -DB02 are defined as linked servers.

Certkiller .com employees use a Windows application to monitor exams for update information. This application displays a list of all certification exams and summary information about each exam.

The list of exams in the application can contain thousands of rows and must be updated immediately when an update to the Exams table in the CK_Courses database occurs. The network bandwidth required for the updating of exam information must be kept to a minimum at all times.

The price for subscriptions is quoted in the subscriber's local currency and is frequently updated to reflect changes in the currency exchange rate. When price quotes are sent to a subscriber, the price is retrieved several times from the table to ensure that quote is accurate. Inaccurate price quotes are not allowed.

BUSINESS REQUIREMENTS:

General Requirements

Certkiller .com has several partners that provide the company with course material. These partners must be able to use their own Web applications to upload course material written in Microsoft Word to the NewMaterials table in the CK_Courses database.

Certkiller .com employees use a Windows application to review and approve the new material. Once approved the new course material must be made available to subscribers and the Certkiller .com partner that provided the course material must be credited.

Performance

Certkiller .com wants the maximum performance for the Web application. All possible blocking caused by the reading and writing of data between connections to be kept at a minimum as the Web application has many users and processes.

Recoverability

When a subscriber takes a new subscription, a row is added to the Subscriptions table in the CK_Courses database. When this occurs, a row must be inserted into two separate tables named NewSubscriptions and SubscriptionLog in a database named CK_Subscriptions. The CK_Subscriptions database is hosted on Certkiller -DB02. If data fails to be inserted into the NewSubscriptions fails, then the corresponding data must not be inserted into the Subscriptions table. If data fails to be inserted into the SubscriptionLog fails, the corresponding data inserted into the NewSubscriptions table and the Subscriptions table.

Security

The Certkiller .com database servers must be protected against SQL Injection attacks.

TECHNICAL REQUIREMENTS:

Brief summaries about certification exams that are collected from the Internet are reviewed by Certkiller .com users and are inserted into a Web application. The summaries are then loaded into the CK_Courses database. However, some inserts might fail due to foreign key constraints and check constraints on the ExamSummaries table. When inserts fail, summaries should be considered as invalid and the corresponding rows should be ignored and the summaries should not be inserted. There could be many summaries that need to be loaded at any one time insert; therefore, the routine used to insert the summaries must be as efficient as possible.

A Windows application that is used to manage new subscribers has a form that calls a method named SubscriptionAdmin. The SubscriptionAdmin method issues either an update or an insert using a separate connection inside a loop that iterates through a System.Data.SqlClient.SqlDataReader object. This method is currently suffering from poor performance and experiences several blocking problems even when no other queries are performed against the databases.

The price for subscriptions that is quoted to the subscriber is frequently updated to reflect changes in the currency exchange rate. The stored procedure used to update the subscription price must use as little system resources as possible and must complete as quickly as possible. The stored procedure must only verify the update of the price. If it the update of the price is not successful, an error message must be issued. Blocking for this stored procedure is not a problem.

Certkiller .com has an intranet Web site that contains a Web page that displays the new subscriptions made the pervious day. Certkiller .com wants this page to display all active subscriptions rather than the new subscriptions of the previous day. Because there may be several thousand subscriptions, the data that is displayed on the page must not include dirty reads and must use the minimum memory on the Web server.

Topic 9, Certkiller .com (12 Questions)

QUESTION 1

You need to design the stored procedure that calculates the subscription price quoted to the subscriber. You need to ensure that your solution meets Certkiller .com's technical and business requirements. What should you do?

- A. Set the isolation level to SNAPSHOT and start a new transaction.
Issue the queries used to calculate the subscription price.
Commit the transaction.
- B. Set the isolation level to READ UNCOMMITTED and start a new transaction.
Issue the queries used to calculate the subscription price.
Commit the transaction.
- C. Set the isolation level to REPEATABLE READ.
Issue the queries used to calculate the subscription price.
Use the XLOCK table locking hint to lock the rows used in the calculation.
- D. Set the isolation level to READ UNCOMMITTED.
Issue the queries used to calculate the subscription price.
Use the XLOCK table locking hint to lock the rows used in the calculation.

Answer: A

Explanation: You should use the snapshot isolation level. Snapshot isolation prevents both uncommitted dependencies (dirty reads) and nonrepeatable reads. It uses row versioning by reading data into tempdb and accessing data from there. It is also less likely to result in blocking locks and deadlock conditions.

Incorrect Answers:

B: Read uncommitted isolation is the least likely to cause deadlock conditions, but prevents neither uncommitted dependencies (dirty reads) nor nonrepeatable reads. This may result in data inconsistencies.

C: Repeatable read isolation will take shared locks; therefore the transaction will be blocked if the Investments table is locked by another session.

D: Read uncommitted isolation is the least likely to cause deadlock conditions, but prevents neither uncommitted dependencies (dirty reads) nor nonrepeatable reads. This may result in data inconsistencies.

QUESTION 2

You need to design the process that will mark price quotes for prospective subscribers that are older than 24 hours as invalid. You need to ensure that your solution meets Certkiller .com's technical and business requirements.

What should you do? (Each correct answer presents part of the solution. Choose TWO.)

- A. Implement a dynamic cursor on the SubscriberInfo table.
- B. Implement a fast-forward cursor on the SubscriberInfo table.
- C. Implement a keyset cursor on the SubscriberInfo table.
- D. Run a correlated subquery against the QuotePrices table.
- E. Issue one update per subscriber.

Answer: C, E

Explanation: The process that marks all price quotes that are older than 24 hours as invalid will run every hour and must run on a per-subscriber basis with all price quotes for one subscriber being processed and then all price quotes for the next subscriber being processed, etc. Thus you need to retrieve one specific row at a time. This is called an absolute fetch and is only supported by keyset or static cursors. You also need to issue one update per subscriber to mark the price quote as invalid.

Incorrect Answers:

A: Dynamic cursors do not support absolute fetches. Dynamic cursors are also less efficient than keyset cursors.

B: Fast-forward cursors do not support absolute fetches.

D: There is not need to run correlated subqueries against the QuotePrices table as you only need to determine the date the quote was issued.

QUESTION 3**DRAG DROP**

You need to design a process that will refresh the list of exams in the application when new information or course material for an exam is made available. You need to ensure that your solution meets Certkiller .com's technical and business requirements.

What should you do? (To answer, drag the appropriate actions from the pane on the left to the pane on the right, and arrange them in the correct order.)

Actions, select from these

- Reissue the SELECT statement to repopulate the list when the event is raised
- Repopulate the exams list every 5 minutes.
- Issue multiple SELECT statements with WHERE clauses to populate the exams list.
- Use a SqlDependency object to subscribe to the OnChanged event.
- Repopulate the exams list every 15 minutes.
- Issue a single SELECT statement to populate the exam list.

Actions, place here

- Place first step here.
- Place second step here, if any.
- Place third step here, if any.
- Place fourth step here, if any.
- Place fifth step here, if any.
- Place sixth step here, if any.

Answer:

Actions, select from these

-
- Repopulate the exams list every 5 minutes.
-
-
- Repopulate the exams list every 15 minutes.
- Issue a single SELECT statement to populate the exam list.

Actions, place here

- Issue multiple SELECT statements with WHERE clauses to populate the exams list.
- Use a SqlDependency object to subscribe to the OnChanged event.
- Reissue the SELECT statement to repopulate the list when the event is raised.
- Place fourth step here, if any.
- Place fifth step here, if any.
- Place sixth step here, if any.

Explanation:

Running multiple SELECT statements with different WHERE clauses will allow you to refresh sections of the list at a time when a change occurs in records that match a particular WHERE clause. This is more efficient than refreshing the entire list. You need to use a SqlDependency object on the Web page to trigger the refresh when a change occurs in records that match a particular WHERE clause.

Incorrect Answers:

Repopulating the list after a preset time interval will be unnecessary if no changes to the database occur during that period. Issuing a single SELECT clause would also require that the entire list be refreshed when a few changes occur. It would be better to use multiple SELECT statements with different WHERE clauses will allow you to refresh sections of the list at a time when a change occurs in records that match a particular WHERE clause.

QUESTION 4

You need to design a stored procedure that allows Certkiller .com partners to upload new course material to the NewMaterials table. You need to ensure that your solution meets Certkiller .com's business and technical requirements. What should you do? (Each correct answer presents part of the solution. Choose TWO.)

- A. Use an HTTP endpoint that uses SSL ports to allow access to the NewMaterials table.
- B. Use a custom ADO.NET SQL Web application to allow access to the NewMaterials table.
- C. Write a stored procedure that will insert the data into the NewMaterials table.
- D. Create a user-defined function that will insert the data into the NewMaterials table.
- E. Allow ad hoc queries to insert the data into the NewMaterials table.

Answer: A, C

Explanation: HTTP endpoints can be used to allow .NET Framework Windows applications and ASP.NET applications access to the database. A stored procedure should be used to minimize database exposure.

Incorrect Answers:

B: You need to allow Certkiller .com's partners to use their own applications to upload data to the NewMaterials table. Creating a custom ADO.NET SQL Web application will not allow the partners to use their own applications.

D: It is not necessary to create a user-defined function as a simple stored procedure can be used to insert the data into the table.

E: You should not allow ad hoc queries to be run against the database as this could be open to abuse. It would be better to use stored procedures or user-defined functions to run the queries against the database.

QUESTION 5

You need to design the Transact-SQL query that inserts new subscriptions into the Subscriptions, NewSubscriptions, and SubscriptionLog tables. You need to ensure that your solution meets Certkiller .com's technical and business requirements. What should you do?

- A. Implement a distributed transaction and commit the transaction before the data is inserted into the SubscriptionLog table.
- B. Implement a distributed transaction and issue a transaction savepoint before the data is

inserted into the SubscriptionLog table.

C. Implement a local transaction and issue a transaction savepoint before the data is inserted into the SubscriptionLog table.

D. Implement a local transaction and commit the transaction after the data is inserted into the SubscriptionLog table.

Answer: A

Explanation: You need to perform inserts into two different databases that are hosted on different database servers. You therefore need to use a distributed transaction. You must also ensure that data inserted into the NewSubscriptions table and the Subscriptions table is saved even if the corresponding data insert into the SubscriptionLog table fails. Therefore you should commit the transaction before and not after the data is inserted into the SubscriptionLog table.

Incorrect Answers:

B: If part of a transaction fails, the transaction can be rolled back to a save point rather than to the start of the transaction. However, the transaction is not committed when the roll back to the save point occurs.

C: You cannot use a local transaction as local transactions do not support changes to different databases that are hosted on different database servers. Also, if part of a transaction fails, the transaction can be rolled back to a save point rather than to the start of the transaction. However, the transaction is not committed when the roll back to the save point occurs.

D: You cannot use a local transaction as local transactions do not support changes to different databases that are hosted on different database servers. You must also ensure that data inserted into the NewSubscriptions table and the Subscriptions table is saved even if the corresponding data insert into the SubscriptionLog table fails. Therefore you should commit the transaction before and not after the data is inserted into the SubscriptionLog table.

QUESTION 6

You need to design the query that populates the intranet Web page that displays active subscriptions. You need to ensure that your solution meets Certkiller .com's technical and business requirements.

What should you do?

A. Set the transaction isolation level to **SERIALIZABLE**.

B. Set the transaction isolation level to **SNAPSHOT**.

C. Set the transaction isolation level to **READ UNCOMMITTED**.

D. Set the transaction isolation level to **REPEATABLE READ**.

Answer: B

Explanation: You should use the snapshot isolation level. Snapshot isolation prevents both uncommitted dependencies (dirty reads) and nonrepeatable reads. It uses row versioning by reading data into tempdb and accessing data from there. It is

also a lower isolation level than Serializable and therefore less likely to result in blocking locks and deadlock conditions.

Incorrect Answers:

A: Serializable isolation prevents both uncommitted dependencies (dirty reads) and nonrepeatable reads but it serializable is also the highest isolation level and is the most likely to result in deadlocks.

C: Read uncommitted isolation is the least likely to cause deadlock conditions, but prevents neither uncommitted dependencies (dirty reads) nor nonrepeatable reads. This may result in data inconsistencies.

D: Repeatable read isolation will take shared locks; therefore the transaction will be blocked if the Investments table is locked by another session.

QUESTION 7

Certkiller .com users complain that the Windows application that is used to review and approve the new course material uploaded to the NewMaterials table takes a considerable amount of time to perform certain operations. You need to determine which queries that are used by the application are performing poorly. You intend to use SQL Server Profiler to replay the queries. You need to ensure that your attempt to resolve the problem has the least possible impact on network performance and the performance of database servers.

What should you do?

- A. Create a trace and save the data to a table in the tempdb.
- B. Create a trace and save the data to file on a shared folder on a remote file server.
- C. Create a trace and save the data to a file on the database server.
- D. Create a trace and save the data to a table in the user database.

Answer: C

Explanation: You should save the data as a file to minimize the performance impact on the database server.

Incorrect Answers:

A: To minimize the impact of saving the data, you should save the data as a file rather than as a table. Furthermore, saving the data to the tempdb will not ensure that you can replay the data at a later stage as the tempdb is rebuilt every time the database server is rebooted.

B: Saving the data to a network server and replaying the file at a later stage will increase network traffic and will impact on performance for users.

D: To minimize the impact of saving the data, you should save the data as a file rather than as a table.

QUESTION 8

You need to redesign the NewMaterials table. You want to implement full-text search to allow Certkiller .com users to search for new course material that have specific information. You need to ensure that the NewMaterials table supports Certkiller .com's business and technical requirements.

What Transact-SQL statement should you use?

A. CREATE TABLE NewMaterials

```
(  
NewMaterialID bigint UNIQUE,  
PartnerID int REFERENCES Partners.PartnerID,  
NewMaterialText varchar(max),  
DateSubmitted smalldatetime NOT NULL  
)
```

B. CREATE TABLE NewMaterials

```
(  
NewMaterialID bigint PRIMARY KEY CLUSTERED,  
PartnerID int REFERENCES Partners.PartnerID,  
NewMaterialText xml,  
DateSubmitted smalldatetime NOT NULL  
)
```

C. CREATE TABLE NewMaterials

```
(  
NewMaterialID bigint UNIQUE, PartnerID int REFERENCES Partners.PartnerID,  
NewMaterialText varbinary(max),  
FileType varchar(8),  
DateSubmitted smalldatetime NOT NULL  
)
```

D. CREATE TABLE NewMaterials

```
(  
NewMaterialID bigint PRIMARY KEY CLUSTERED,  
PartnerID int REFERENCES Partners.PartnerID,  
NewMaterialText image,  
FileExtension varchar(8),  
DateSubmitted smalldatetime NOT NULL  
)
```

Answer: D

Explanation: You should create the NewMaterialID column as the primary key so that it can be used as the key column for the full-text search. Because the Web master uses an application to copy the sample chapter text from Microsoft Word to the Samples table, you need to define the NewMaterialText column as either varbinary(max) or image. To enable Full-text indexing on a Microsoft Word document, you also need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

Incorrect Answers:

A: Full-text search requires a key column that contains unique, non-null values. A UNIQUE constraint does not ensure that the column has no null values; it limits null values to one row. Because the Web master uses an application to copy the sample chapter text from Microsoft Word to the Samples table, you need to define the

NewMaterialText column as either varbinary(max) or image and not varchar(max). Also, to enable Full-text indexing on a Microsoft Word document, you need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

B: Because the Web master uses an application to copy the sample chapter text from Microsoft Word to the Samples table, you need to define the NewMaterialText column as either varbinary(max) or image and not xml. Also, to enable Full-text indexing on a Microsoft Word document, you need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

C: Full-text search requires a key column that contains unique, non-null values. A UNIQUE constraint does not ensure that the column has no null values; it limits null values to one row.

QUESTION 9

You need to resolve the problem with the SubscriptionAdmin method. You want to determine which objects are involved in the deadlocks.

What should you do?

- A. Use SQL Server Profiler to run a trace.
- B. Run the sp_lock stored procedure.
- C. Query the sys.dm_tran_database_transactions dynamic management view.
- D. Query the sys.dm_tranlocks dynamic management view.

Answer: A

Explanation: You can use SQL Server Profiler to monitor locks and deadlocks.

Incorrect Answers:

B, D: The sp_lock stored procedure has been deprecated in SQL Server 2005 in favor of the sys.dm_tranlocks dynamic management view. However, the sys.dm_tranlocks dynamic management view allows you to view current locks and deadlocks. It does not allow you to view a history of deadlocks

C: The sys.dm_tran_database_transactions dynamic management view provides information about transactions rather than locks and deadlocks.

QUESTION 10

You need to design the functionality that allows Web site users to search for certification exams. You need to allow users to search for exams offered by a particular the vendor. You need to ensure that your solution meets Certkiller .com's technical and business requirements.

What should you do?

- A. Create a partitioned view.
- B. Create a Transact-SQL (T-SQL) user-defined function.
- C. Create a stored procedure with the WITH RECOMPILE option.
- D. Create a Common Language Runtime (CLR) user-defined function.

Answer: B

Explanation: A T-SQL user-defined table-type function returns a table to an application and protects against SQL Injection by allowing validation of the input parameters.

Incorrect Answers:

A: You cannot pass parameters to a view. Furthermore, partitioned views have been deprecated in SQL Server 2005.

C: A stored procedure with the WITH RECOMPILE option will be recompiled every time it is executed. This will degrade performance.

D: A CLR user-defined function should be created when T-SQL functions cannot be implemented to achieve the desired results. However, a T-SQL function can be used in this scenario.

QUESTION 11

You need to optimize the CK_Courses database to support a redesigned stored procedure. You need add a key column to an existing clustered index on the Subscriptions table but you must ensure that the table remains available to users while the index is built.

What should you do?

A. Run the CREATE INDEX statement with the DROP EXISTING and ONLINE options set to ON.

B. Disable the index. Then run the CREATE INDEX statement with the DROP EXISTING and ONLINE options set to ON

C. Run the ALTER INDEX statement with the REORGANIZE keyword and the ONLINE option set to ON.

D. Run the ALTER INDEX statement with the REBUILD keyword and the ONLINE option set to ON.

Answer: A

Explanation:

You must drop the clustered index and recreate it if you want to add a column to an existing clustered index. You can drop and recreate an existing index using the CREATE INDEX statement and setting the DROP EXISTING option to ON. To ensure that users can access the table while the index is being built, you must also set the ONLINE option to ON.

Incorrect Answers:

B: A clustered index cannot be rebuilt online if the index is disabled.

C, D: You cannot use the ALTER INDEX statement to add a column to a clustered index. You must drop the clustered index and recreate it if you want to add a column to an existing clustered index.

QUESTION 12

You need design a solution that resolves the problems with the SubscriptionAdmin

method. You need to ensure that your solution meets Certkiller .com's technical and business requirements.

What should you do?

- A. SET IMPLICIT_TRANSACTIONS OFF for the SubscriptionAdmin method.
- B. Set the transaction isolation level to SERIALIZABLE.
- C. Enable Multiple Active Result Sets (MARS) on a single connection.
- D. Implement Microsoft Data Access Components (MDAC).

Answer: C

Explanation: Multiple Active Result Sets (MARS) allows you to run multiple queries on a single connection and allows an application to have multiple default result sets open at one time. While the multiple default result sets are open the application can execute other Transact-SQL statements.

Incorrect Answers:

A: The implicit transactions options determines whether the connection used autocommit or not. However, this will not resolve a blocking issue experienced by the SubscriptionAdmin method.

B: Serializable isolation is the highest isolation level and is the most likely to result in deadlocks when other connection need to access the data.

D: Microsoft Data Access Components (MDAC) do not support Multiple Active Result Sets (MARS) connections.