

TestkingIT

Testking IT

> Contact Us

Login / Register

Search...



HOME

ALL VENDORS

★ GUARANTEE

? FAQ

TESTIMONIALS

CART (0)



Try **Desktop Test Engine** before you buy

We're not the only ones **happy** about TestKingsIT Practice Material ...

48236+ customers in 100+ countries use TestKingsIT Test Engine. Meet our customers.



<http://www.testkingit.com/>

Latest practice material - Exam Cram - TestKingIT

Exam : 070-454

Title : UPG: Transition MCITP SQL Serv
2005 to SQL Serv 2008 DBD

Vendors : Microsoft

Version : DEMO

NO.1 You are developing a new database. The database contains two tables named SalesOrderDetail and Product. You need to ensure that all products referenced in the SalesOrderDetail table have a corresponding record in the Product table. Which method should you use?

- A. JOIN
- B. DDL trigger
- C. Foreign key constraint
- D. Primary key constraint

Answer: C

NO.2 You have multiple tables that represent properties of the same kind of entities. The property values are comprised of text, geometry, varchar(max), and user-defined types specified as 'bit NOT NULL' data types. You plan to consolidate the data from multiple tables into a single table. The table will use semi-structured storage by taking advantage of the SPARSE option. You are tasked to identify the data types that are compatible with the SPARSE option. Which data type is compatible with the SPARSE option?

- A. text
- B. geometry
- C. varchar(max)
- D. A user-defined type defined as 'bit NOT NULL'

Answer: C

NO.3 You manage a SQL Server 2008 database that is located at your company's corporate headquarters. The database contains a table named dbo.Sales. You need to create different views of the dbo.Sales table that will be used by each region to insert, update, and delete rows. Each regional office must only be able to insert, update, and delete rows for their respective region. Which view should you create for Region1?

- A. CREATE VIEW dbo.Region1Sales AS SELECT SalesID, OrderQty, SalespersonID, RegionID FROM dbo.Sales WHERE RegionID = 1;

B.CREATE VIEW dbo.Region1Sales AS SELECT
SalesID,OrderQty,SalespersonID,RegionID FROM
dbo.Sales WHERE RegionID = 1 WITH CHECK OPTION;
C.CREATE VIEW dbo.Region1Sales WITH SCHEMABINDING AS SELECT
SalesID,OrderQty,SalespersonID,RegionID FROM dbo.Sales WHERE RegionID = 1;
D.CREATE VIEW dbo.Region1Sales WITH VIEW_METADATA AS SELECT
SalesID,OrderQty,SalespersonID,RegionID FROM dbo.Sales WHERE RegionID = 1;
Answer: B

NO.4 You need to create a column that allows you to create a unique constraint.Which two
column definitions

should you choose? (Each correct answer presents a complete solution.Choose two.)

- A.nvarchar(100) NULL
- B.nvarchar(max) NOT NULL
- C.nvarchar(100) NOT NULL
- D.nvarchar(100) SPARSE NULL

Answer: AC

NO.5 You are creating a new table in a database.Your business requires you to store data in
the table for

only seven days.You need to implement a partitioned table to meet this business
requirement.Which tasks

should you complete?

- A.Create the partition function Create the partition scheme Create the table
- B.Create the partition function Create the table Create a filtered index
- C.Add a secondary file to the primary filegroups Create the table Create the distributed
partitioned view
- D.Create the partition function Create the partition scheme Create the distributed partitioned
view

Answer: A

NO.6 You have a table named Customer.You need to ensure that customer data in the table
meets the

following requirements: credit limit must be zero unless customer identification has been
verified.credit

limit must be less than 10,000.Which constraint should you use?

- A.CHECK (CreditLimt BETWEEN 1 AND 10000)
- B.CHECK (Verified = 1 AND CreditLimt BETWEEN 1 AND 10000)

- C.CHECK ((CreditLimt = 0 AND Verified = 0) OR (CreditLimt BETWEEN 1 AND 10000 AND Verified = 1))
- D.CHECK ((CreditLimt = 0 AND Verified = 0) AND (CreditLimt BETWEEN 1 AND 10000 AND Verified = 1))

Answer: C

NO.7 You are creating a table that stores the GPS location of customers.You need to ensure that the table allows you to identify customers within a specified sales boundary and to calculate the distance between a customer and the nearest store.Which data type should you use?

- A.geometry
- B.geography
- C.nvarchar(max)
- D.varbinary(max) FILESTREAM

Answer: B

NO.8 You need to create a stored procedure that accepts a table-valued parameter named @Customers.Which code segment should you use?

- A.CREATE PROCEDURE AddCustomers (@Customers varchar(max))
- B.CREATE PROCEDURE AddCustomers (@Customers Customer READONLY)
- C.CREATE PROCEDURE AddCustomers (@Customers CustomerType OUTPUT)
- D.CREATE PROCEDURE ADDCUSTOMERS (@Customers varchar (max)) AS EXTERNAL

NAME
Customer.Add.NewCustomer

Answer: B

NO.9 Your database is 5GB and contains a table named SalesHistory.Sales information is frequently inserted and updated.You discover that excessive page splitting is occurring.You need to reduce the occurrence of

page splitting in the SalesHistory table.Which code segment should you use?.

- A.ALTER DATABASE Sales MODIFY FILE (NAME = Salesdat3, SIZE = 10GB);
- B.ALTER INDEX ALL ON Sales.SalesHistory REBUILD WITH (FILLFACTOR = 60);
- C.EXEC sys.sp_configure 'fill factor (%)', '60';
- D.UPDATE STATISTICS Sales.SalesHistory(Products) WITH FULLSCAN, NORECOMPUTE;

Answer: B

NO.10 You have a table named AccountsReceivable. The table has no indexes. There are 75,000 rows in the table. You have a partition function named FG_AccountData. The AccountsReceivable table is defined in the following Transact-SQL statement:

```
CREATE TABLE AccountsReceivable ( column_a INT NOT NULL, column_b VARCHAR(20)
NULL) ON
[PRIMARY];
```

You need to move the AccountsReceivable table from the PRIMARY file group to FG_AccountData. Which

Transact-SQL statement should you use?

A. CREATE CLUSTERED INDEX idx_AccountsReceivable ON AccountsReceivable(column_a) ON [FG_AccountData];

B. CREATE NONCLUSTERED INDEX idx_AccountsReceivable ON AccountsReceivable(column_a) ON [FG_AccountData];

C. CREATE CLUSTERED INDEX idx_AccountsReceivable ON AccountsReceivable(column_a) ON FG_AccountData(column_a);

D. CREATE NONCLUSTERED INDEX idx_AccountsReceivable ON AccountsReceivable(column_a) ON FG_AccountData(column_a);

Answer: C

NO.11 You administer a SQL Server 2008 database that contains a table name dbo.Sales, which contains the following table definition:

```
CREATE TABLE [dbo].[Sales]( [SalesID] [int] IDENTITY(1,1) NOT NULL PRIMARY KEY
CLUSTERED,
[OrderDate] [datetime] NOT NULL, [CustomerID] [int] NOT NULL, [SalesPersonID] [int]
NULL,
[CommentDate] [date] NULL);
```

This table contains millions of orders. You run the following query to determine when sales persons comment in the dbo.Sales table:

SELECT SalesID, CustomerID, SalesPersonID, CommentDate FROM dbo.Sales WHERE
CommentDate

IS NOT NULL AND SalesPersonID IS NOT NULL;

You discover that this query runs slow. After examining the data, you find only 1% of rows
have comment

dates and the SalesPersonID is null on 10% of the rows. You need to create an index to
optimize the

query. The index must conserve disk space while optimizing your query. Which index should
you create?

A. CREATE NONCLUSTERED INDEX idx1 ON dbo.Sales (CustomerID) INCLUDE
(CommentDate, SalesPersonID);

B. CREATE NONCLUSTERED INDEX idx1 ON dbo.Sales (SalesPersonID) INCLUDE
(CommentDate, CustomerID);

C. CREATE NONCLUSTERED INDEX idx1 ON dbo.Sales (CustomerID)
INCLUDE(CommentDate)

WHERE SalesPersonID IS NOT NULL;

D. CREATE NONCLUSTERED INDEX idx1 ON dbo.Sales (CommentDate, SalesPersonID)
INCLUDE(CustomerID) WHERE CommentDate IS NOT NULL;

Answer: D

NO.12 You have a SQL Server database. The database contains two schemas named
Marketing and

Sales. The Marketing schema is owned by a user named MarketingManager. The Sales
schema is owned

by a user named SalesManager. A user named John must be able to access the
Sales.Orders table by

using a stored procedure named Marketing.GetSalesSummary. John is not granted a
SELECT permission

on the Sales.Orders table. A user named SalesUser does have SELECT permission on the
Sales.Orders

table. You need to implement appropriate permissions for John and the stored procedure
Marketing.GetSalesSummary. What should you do?

A. Marketing.GetSalesSummary should be created by using the EXECUTE AS 'SalesUser'
clause. John

should be granted EXECUTE permission on Marketing.GetSalesSummary.

B. Marketing.GetSalesSummary should be created by using the EXECUTE AS OWNER
clause. John

should be granted EXECUTE WITH GRANT OPTION on Marketing.GetSalesSummary.

C. Marketing.GetSalesSummary should be created by using the EXECUTE AS CALLER clause. John

should be granted IMPERSONATE permission for the user named SalesUser.

D. Marketing.GetSalesSummary should be created without an EXECUTE AS clause. John should be

granted SELECT permission on the Sales.Orders table.

Answer: A

NO.13 You plan to add a new column named SmallKey to the Sales.Product table that will be used in a unique

constraint. You are required to ensure that the following information is applied when adding the new

column:

'a1' and 'A1' are treated as different values 'a' and 'A' sort before 'b' and 'B' in an ORDER BY clause You

need to select the collation that meets the requirements for the new column. Which collation should you

select?

A. Latin1_General_BIN

B. SQL_Latin1_General_CP1_CI_AI

C. SQL_Latin1_General_CP1_CI_AS

D. SQL_Latin1_General_CP1_CS_AS

Answer: D

NO.14 You have a computed column that is implemented with a user-defined function. The user-defined

function returns a formatted account number. The column must be indexed to provide adequate search

performance. You plan to create an index on the computed column. You need to identify the valid

combination of ObjectPropertyEX values for the user-defined function. Which combination should you

use?

A. IsDeterministic = True IsSystemVerified = True UserDataAccess = False

SystemDataAccess = False

B. IsDeterministic = True IsSystemVerified = True IsPrecise = True IsTableFunction = True

C. IsDeterministic = False IsSystemVerified = True UserDataAccess = False

SystemDataAccess = False

D.IsDeterministic = False IsSystemVerified = True IsPrecise = True SystemDataAccess = False

Answer: A

NO.15 You have a user named John.He has SELECT access to the Sales schema.You need to eliminate

John's SELECT access rights from the Sales.SalesOrder table without affecting his other permissions.Which Transact-SQL statement should you use?

A.DROP USER John;

B.DENY SELECT ON Sales.SalesOrder TO John;

C.GRANT DELETE ON Sales.SalesOrder TO John;

D.REVOKE SELECT ON Sales.SalesOrder FROM John;

Answer: B

NO.16 You currently store date information in two columns.One column contains the date in local time and

one column contains the difference between local time and UTC time.You need to store this data in a

single column.Which data type should you use?

A.time

B.datetime2

C.datetime2(5)

D.datetimeoffset

Answer: D

NO.17 You have two partitioned tables named Transaction and TransactionHistory.You need to archive one of

the partitions of the Transaction table to the TransactionHistory table.Which method should you use?

A.ALTER TABLE ...SWITCH ...

B.INSERT ...SELECT ...; TRUNCATE TABLE

C.ALTER PARTITION FUNCTION ...MERGE ...

D.ALTER PARTITION FUNCTION ...SPLIT ...

Answer: B

NO.18 You need to alter stored procedures to use the WITH RECOMPILE option.Which types of stored

procedures should you alter? (Each correct answer represents a complete solution.Choose

two.)

- A. Stored procedures implemented from CLR assemblies.
- B. Stored procedures that require the FOR REPLICATION option.
- C. Stored procedures that require the WITH ENCRYPTION option.
- D. Stored procedures that contain queries that use the OPTION (RECOMPILE) hint.

Answer: CD

NO.19 You have a SQL Server 2008 database named Contoso with a table named Invoice. The primary key of the table is InvoiceId, and it is populated by using the identity property. The Invoice table is related to the InvoiceLineItem table. You remove all constraints from the Invoice table during a data load to increase load speed. You notice that while the constraints were removed, a row with InvoiceId = 10 was removed from the database. You need to re-insert the row into the Invoice table with the same InvoiceId value. Which

Transact-SQL statement should you use?

- A. `INSERT INTO Invoice (InvoiceId, ...VALUES (10, ...`
- B. `SET IDENTITY_INSERT Invoice ON; INSERT INTO Invoice (InvoiceId, VALUES (10, SET IDENTITY_INSERT Invoice OFF;`
- C. `ALTER TABLE Invoice; ALTER COLUMN InvoiceId int; INSERT INTO Invoice (InvoiceId, ...VALUES (10, ...`
- D. `ALTER DATABASE Contoso SET SINGLE_USER; INSERT INTO Invoice (InvoiceId, ...VALUES (10, ...ALTER DATABASE Contoso SET MULTI_USER;`

Answer: B

NO.20 You have a table named dbo.Customers. The table was created by using the following Transact-SQL statement:

```
CREATE TABLE dbo.Customers ( CustomerID int IDENTITY(1,1) PRIMARY KEY CLUSTERED, AccountNumber nvarchar(25) NOT NULL, FirstName nvarchar(50) NOT NULL, LastName nvarchar(50) NOT NULL, AddressLine1 nvarchar(255) NOT NULL, AddressLine2 nvarchar(255) NOT NULL, City
```

nvarchar(50)
NOT NULL, StateProvince nvarchar(50) NOT NULL, Country nvarchar(50) NOT NULL,
PostalCode
nvarchar(50) NOT NULL, CreateDate datetime NOT NULL DEFAULT(GETDATE()),
ModifiedDate
datetime NOT NULL DEFAULT(GETDATE()))

You create a stored procedure that includes the AccountNumber, Country, and StateProvince columns

from the dbo.Customers table. The stored procedure accepts a parameter to filter the output on the

AccountNumber column. You need to optimize the performance of the stored procedure. You must not

change the existing structure of the table. Which Transact-SQL statement should you use?

A. CREATE STATISTICS ST_Customer_AccountNumber ON dbo.Customer
(AccountNumber) WITH
FULLSCAN;

B. CREATE CLUSTERED INDEX IX_Customer_AccountNumber ON dbo.Customer
(AccountNumber);

C. CREATE NONCLUSTERED INDEX IX_Customer_AccountNumber ON dbo.Customer
(AccountNumber) WHERE AccountNumber = ";

D. CREATE NONCLUSTERED INDEX IX_Customer_AccountNumber ON dbo.Customer
(AccountNumber) INCLUDE (Country, StateProvince);

Answer: D