

[Oct.-2016Free 191q 70-464 PDF Instant Download from Braindump2go[Q32-Q37

2016/10 New Microsoft 70-464: Developing Microsoft SQL Server 2012/2014 Databases Exam Questions Updated Today! Free Instant Download 70-464 Exam Dumps (PDF & VCE) 191Q&As from Braindump2go.com Today! 100% Real Exam Questions! 100% Exam Pass Guaranteed! 1. | 2016/10 Latest 70-464 Exam Dumps (PDF & VCE) 191Q&As Download:

<http://www.braindump2go.com/70-464.html> 2. | 2016/10 Latest 70-464 Exam Questions & Answers:

https://drive.google.com/folderview?id=0B9YP8B9sF_gNeUZnbHVMWDhuQzA&usp=sharing QUESTION 32 You have a text file that contains an XML Schema Definition (XSD). You have a table named Schemal.Table1. You have a stored procedure named Schemal.Proc1 that accepts an XML parameter named Param1. You need to store validated XML data in Schemal.Table1. The solution must ensure that only valid XML data is accepted by Param1. What should you do? (Each correct answer presents part of the solution. Choose all that apply.) A. Define an XML column in Table1 by using an XML schema collection. B. Create an XML schema collection in the database from the text file. C. Declare Param1 var1 as type XML and associate the variable to the XML schema collection. D. use the modify method to insert the XML schema into each row of the XML column in Table1.

Answer: ABDE Explanation: <http://msdn.microsoft.com/en-us/library/bb510420.aspx>

<http://msdn.microsoft.com/en-us/library/ms187856.aspx> <http://msdn.microsoft.com/en-us/library/ms176009.aspx>

<http://msdn.microsoft.com/en-us/library/hh403385.aspx> <http://msdn.microsoft.com/en-us/library/ms184277.aspx> QUESTION 33

Drag and Drop Question You plan to deploy two stored procedures named SP1 and SP2 that read data from the database. Your company identifies the following requirements for each stored procedure:- SP1 must allow dirty reads.- SP2 must place range locks on the data to ensure read consistency. You need to identify which isolation level you must set for each stored procedure. The solution must minimize the number of locks. Which isolation level should you identify? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area. (Answer choices may be used once, more than once, or not at all.)

Isolation Levels

read committed

read uncommitted

repeatable read

serializable

snapshot

Answer:

Isolation Levels	Answer Area
read committed	SP1 read uncommitted
repeatable read	SP2 serializable
snapshot	

QUESTION 34 You use SQL Server 2014. The physical server is a dedicated database server that has 120GB of RAM available. There is approximately 50GB of storage space available on a slow local disk. You create a new stored procedure. You decide you need to temporarily hold approximately 300,000 rows from two tables, from which you will compute two complex business scores. The stored procedure will use temporary storage defined as follows:

```
AccountNumber char(10) NOT NULL
YearToDateSalesTotal decimal(15,2) NULL
SalesScore int NULL
FutureSalesExpectationScore int NULL
```

The code will make several passes through the data, applying complex calculations before writing the data to a permanent disk-based table in the same database from which it reads the data. For this stored procedure, you need to deal with temporary data in

the most efficient way to minimize physical disk pressure. What should you do? More than one answer choice may achieve the goal.

Select the BEST answer.

```
○A CREATE TYPE dbo.AccountScoringModel as TABLE
(
    AccountNumber char(10) COLLATE Latin1_General_100_BIN2 NOT NULL ,
    YearToDateSalesTotal decimal(15,2) NULL,
    SalesScore int NULL,
    FutureSalesExpectationScore int NULL,
    INDEX AccountNumber HASH (AccountNumber) WITH (BUCKET_COUNT = 25000)
) WITH ( MEMORY_OPTIMIZED = ON )
GO
DECLARE @AccountScoring as dbo.AccountScoringModel

○B DECLARE @AccountScoring as TABLE
(
    AccountNumber char(10) NOT NULL,
    YearToDateSalesTotal decimal(15,2) NULL,
    SalesScore int NULL,
    FutureSalesExpectationScore int NULL
)

○C CREATE TABLE #AccountScoring
(
    AccountNumber char(10) NOT NULL,
    YearToDateSalesTotal decimal(15,2) NULL,
    SalesScore int NULL,
    FutureSalesExpectationScore int NULL
)

○D CREATE TYPE dbo.AccountScoringModel as TABLE
(
    AccountNumber char(10) COLLATE Latin1_General_100_BIN2 NOT NULL ,
    YearToDateSalesTotal decimal(15,2) NULL,
    SalesScore int NULL,
    FutureSalesExpectationScore int NULL,
    INDEX AccountNumber HASH (AccountNumber) WITH (BUCKET_COUNT = 120)
) WITH ( MEMORY_OPTIMIZED = ON )
GO
DECLARE @AccountScoring as dbo.AccountScoringModel
```

A. Option AB. Option BC. Option CD. Option D Answer: A Explanation: - You must specify a value for the BUCKET_COUNT parameter when you create the memory-optimized table. In most cases the bucket count should be between 1 and 2 times the number of distinct values in the index key. If the index key contains a lot of duplicate values, on average there are more than 10 rows for each index key value, use a nonclustered index instead You may not always be able to predict how many values a particular index key may have or will have. Performance should be acceptable if the BUCKET_COUNT value is within 5 times of the actual number of key values. QUESTION 35 You have a SQL Server 2012 database named Database1. You execute the following code: You insert 3 million rows into Sales. You need to reduce the amount of time it takes to execute Proc1. What should you do?

```
CREATE TABLE Sales
(
    ID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
    OrderDate char(10) NOT NULL,
    Amount decimal
);
GO

CREATE INDEX IX_Sales_OrderDate
ON Sales(OrderDate)
INCLUDE (ID, Amount);

CREATE PROC usp_Proc1
(
    @date1 datetime,
    @date2 datetime
)
AS
SELECT ID, OrderDate, Amount
FROM Sales
WHERE CAST(OrderDate AS datetime)
BETWEEN @date1 AND @date2
ORDER BY ID;
GO
```

A. ProductType varchar(11) '@ProductType', B. ProductType varchar(11) 'ProductType/ID', C. ProductType varchar(11) 'ProductType/@ID', D. ProductType varchar(11) 'ProductType1'. Answer: D QUESTION 36 You have a SQL Azure database. You execute the following script: You add 1 million rows to Table1. Approximately 85 percent of all the rows have a null value for Column2. You plan to deploy an application that will search Column2. You need to create an index on Table1 to support the planned deployment. The solution must minimize the storage requirements. Which code segment should you execute?

```
CREATE TABLE dbo.Tab
(
    Column2 varchar(50)
)
```

A. CREATE INDEX IX_Table1 ON Table1 (Column2)WITH FILLFACTOR=0B. CREATE INDEX IX_Table1 ON Table1 (Column1)INCLUDE (Column2)C. CREATE INDEX IX_Table1 ON Table1 (Column2)WHERE Column2 IS NULLD. CREATE INDEX IX_Table1 ON Table1 (Column2)WHERE Column2 IS NOT NULL Answer: D Explanation: <http://msdn.microsoft.com/en-us/library/ms188783.aspx><http://msdn.microsoft.com/en-us/library/cc280372.aspx> QUESTION 37 You are creating a table named Orders. You need to ensure that every time a new row is added to the Orders table, a table that is used for auditing is updated. What should you use? More than one answer choice may achieve the goal. Select the BEST answer. A. a CHECK constraint B. a FOREIGN KEY constraint C. a DEFAULT constraint D. a data manipulation language (DML) trigger E. a Data Definition Language (DDL) trigger Answer: D Explanation: <http://www.techrepublic.com/blog/programming-and-development/comparing-sql-server-constraints-and-dmltriggers/402> <http://msdn.microsoft.com/en-us/library/ms178110.aspx> !!!RECOMMEND!!! 1.|2016/10 Latest 70-464 Exam Dumps (PDF & VCE) 191 Q&As Download: <http://www.braindump2go.com/70-464.html> 2.|2016/10 Latest 70-464 Exam Questions & Answers: https://drive.google.com/folderview?id=0B9YP8B9sF_gNeUZnbHVMWDhuQzA&usp=sharing